LUISS

Department of Business and Management

Chair of International Financial Economics

Federal Reserve’s Communication: a Latent Dirichlet Allocation Analysis with Application to FOMC Minutes

Guido Traficante

SUPERVISOR

Mauro Visaggio

CO-SUPERVISOR

Daisy Martina Laganà
713391

CANDIDATE

Academic Year 2019/2020

*Alla piccola Elena Sofia, che tu possa avere sempre il coraggio
di inseguire i tuoi sogni e le tue ambizioni.*

Abstract

Over the last thirty years Central Banks have become remarkably more transparent about their policy-making, drifting apart from the tradition of “never explain, never excuse” and using communication as a key tool of monetary policy. This revolution in thinking was motivated by different reasons but mostly by a considerable amount of evidence of enhancement of the effectiveness of monetary policy. In the aftermath of the Global Financial Crisis much attention has been posed to Central Banks’ forward-looking communication related to the monetary policy stance, what is known as forward guidance. In this thesis I apply Latent Dirichlet Allocation (LDA), a text mining technique, to the FOMC minutes from 2002 to 2013 to analyze if forward guidance has transformed the contents of FOMC minutes over time.

Table of Contents

I. Introduction

Thirty years ago there was a popular belief that secrecy and ambiguity were part of Central Banks' culture. Many central bankers believed that their actions should be shrouded in mystery to enhance the effectiveness of monetary policy. However, over time Central Banks have become remarkably more transparent about their policy-making, drifting apart from the tradition of *"never explain, never excuse"*, and also started to rely more heavily on communication for monetary policy-making. This revolution in thinking was motivated by different reasons but mostly by a considerable amount of empirical evidence of enhancement of the effectiveness of monetary policy via the management of the future short-term rate expectations. As illustrated by many, such as Bernanke (2004) and Woodford (2005), although the public tends to focus on short term rates such as the fed funds rate, long-term ones are much more important to economic activities. Such long-term rates reflect market expectations of the future evolution of overnight rates and other related short-term rates. Thus, many economists argue that better communication about Central Bank's next moves contributes to the management of those expectations and thus in the effectiveness of monetary policy. Eventually, central banking practices changed, embracing a new era of openness and greater communication. As a result, Central Bank communication has gained increasing attention by market participants, media, academics and the public in general and nowadays can boast a vast literature. Within such, there is also a small but growing branch that focuses on the application of text mining techniques on Central Banks' documents for content analysis. Text mining refers to the wide set of computational tools and statistical techniques that are able to quantify text data, that is, to treat it quantitatively to extract valuable insights (Bholat et al., 2015, p.1). Within central banking those techniques are mostly applied to infer the tone or the content describing a set of documents. For my analysis I make use of a popular text mining technique, called LDA, to infer the proportion of FOMC minutes devoted to the topic of monetary policy over time. Thus, this thesis contributes to a small but growing literature devoted to the use of text mining techniques to Central Banks' documents (and in particular, to the Federal Reserve's ones) for content analysis. As I illustrate later, the very first application of LDA technique in economics is very recent, in 2017, when Hansen et al. studied how policy deliberations have evolved following FOMC members' knowledge of their discussion being public. Following such work, Jagadeesh and Wu (2017) applied LDA to the FOMC minutes to extract the different topics and compute their proportions over time so to study their informativeness for the stocks and bonds markets. Lastly, the most recent application is the one of Edison and Carcel (2020) who analyzed the evolution of the contents of FOMC transcripts.

Created in 1936, the Federal Open Market Committee (FOMC), often referred to as the Fed, is the body in charge of monetary policy decisions. It is made up of twelve members, seven being the members the Board of Governors, one the President of the Federal Reserve Bank of New York and four the Presidents of the remaining eleven Federal Reserve Banks, who serve for a one-year term on a rotating basis. Although not all entitled to vote, all Federal Reserve Banks Presidents attend FOMC meetings and participate in the discussions (known as "participants"). The Chairman of the Board of Governors is also the Chairman of the FOMC. The FOMC typically meets eight times a year and more frequently if necessary. During the meetings each of the twelve members presents its view on the economic outlook and vote on the fed funds rate. The FOMC marked significant milestones toward more transparent policy-making over time (see table 3 of Appendix 1: Additional Tables). Although with some initial reluctance, such road started with Alan Greenspan, the Maestro who prided himself for his "mumbling with great incoherence", and still continues under the current Chairman Jerome Powell. Nowadays, the FOMC communicates with the public with different tools, such as post-meeting statements, minutes, transcripts, frequent speeches, FOMC participants' economic projections, memos, the so-called "Beige Books", semi-annual testimonies before the Government and recently with regular post-meeting press conferences. Three of such documents gain particular attention from the market and academic world, transcripts, minutes and post-meeting statements. The transcripts are the most detailed record of the FOMC meeting proceedings as they report the members' full conversations. They are released after 5 years. The minutes are a comprehensive summary of the transcripts and differently from the latter are released only after a three-week lag from the meeting. The most important document is the post-meeting statement, a short document released shortly after the end of the meeting which reports the FOMC's view of the economy, its latest policy decisions and reasons as well as an indication of the future path of monetary policy (although not always). Given the importance of the contents contained in the statements, minutes and transcripts of meetings, it is not surprising that these are heavily scrutinized. Although today the FOMC makes use of such documents to communicate with the public about its latest policy decisions as well as the future path of monetary policy, such use was not perfect from the start. Indeed, only in 1994 the Fed started to issue its post-meeting statements, which, however, were much shorter[1] and looser in content with respect to today's ones. Perhaps the most significant change in communication policy was in August 2003 when the FOMC started to provide some indications in their post-meeting statements of the expected or intended future path of policy rates, what it is known today as forward guidance. Although nowadays the major Central Banks make amply use of such language, at that time forward guidance was a fundamentally new idea in the monetary policy-making. Providing guidance on the likely future path of the benchmark rate was a breakthrough also for the Federal Reserve, who has been characterized for a long time by the "old school" of mystery and secrecy. Woodford (2005) would later describe such openness about policy decisions and policy future likely course as one of the most significant changes in the Fed during the Greenspan era. The reason why the FOMC was historically skeptical to provide such guidance is that most of its members, including Chairman Greenspan, feared the public might confuse such indications with commitment and so that such guidance would have limited their flexibility to adjust policy according to the evolving conditions of the economy (Poole, 2005). The Maestro's view may also have been buttressed by his long-term belief that the always-evolving

[1]For example, see Wynne, M. (2014). *How the FOMC Talks*. Presentation to the 2014 Economic Summit "The Challenge: Keeping Up". Retrieved from https://www.dallasfed.org/~/media/documents/educate/events/2014/14summitwynne.pdf, to observe how the FOMC statements' word count significantly increased over time.

conditions of the economy require flexible and adaptable monetary policy. For this reason, from its very first adoption, the FOMC remarks the conditionality of its forward guidance statements on the information available to the Central Bank at the time of the guidance. In reality, the FOMC already provided looser policy guidance since May 1999, in the form of "tilt" or "bias" and then in the form of "balance of risks". The latter indicated the members' perceived risks in the foreseeable future to the Fed's long-run goals of price stability and sustainable economic growth. Such statements did not provide a direct indication of the FOMC next moves but rather a language from which such information could be implicitly inferred. In August 2003, worried about a possible emergence of deflation, after lowering the fed funds rate almost to the zero lower bound, the FOMC tried to manage expectations with a firmer and time-dependent form of forward guidance: *"policy accommodation can be maintained for a considerable period"*. Afterward, the FOMC used a sequence of changes in its statement language to guide the market on the likely future course of the policy. This direct interest rate guidance that started in August 2003 ended in June 2006. The forward guidance language was then resurrected once the FOMC had clearer ideas regarding the economic outlook. In December 2008, at the peak of the Global Financial Crisis, the FOMC used a purely qualitative language to indicate that policy accommodation would have continued for *"some time"*. Such forward guidance language evolved over time, according to the needs of the Committee, from pure qualitative to the much more specific and direct date-based forward guidance of August 2011 - October 2012 and state-contingent forward guidance of December 2012 - December 2013. Table 2 in Appendix 1: Additional Tables reports a detailed record of FOMC use of direct forward guidance between 2003-2013. Overall, the FOMC increase in policy transparency, especially during and after the Global Financial Crisis, was noteworthy and unthinkable before 1994. It is not surprising that the Chairman of the crisis and post-crisis years, Ben Bernanke, was a long-term proponent of policy transparency and of the use of forward guidance for shaping market expectations, especially in the zero lower bound condition, when future rate cuts are no longer possible. From the very beginning of his chairmanship, Bernanke was strongly intentioned to make *"monetary policy as transparent and open as reasonably as possible"* (Bernanke, 2013). Although Greenspan set the stage, Bernanke steered the Federal Reserve in the direction of greater policy transparency, especially thought greater use of forward guidance language. In conclusion, over the course of the last twenty years the Fed marked significant steps toward a more transparent policy-making. Such steps were welcomed not only because they demystified the "Temple" but also because they helped market participants to better understand the likely evolution of monetary policy and so helped the FOMC to better manage the economy.

For this thesis I applied Latent Dirichlet Allocation (LDA), a text mining technique belonging to unsupervised machine learning, to the FOMC minutes from 2002 to 2013 to analyze if increased policy transparency via forward guidance and greater use of the latter in the FOMC post-meeting statements have transformed the contents of their minutes over time. In particular, as the FOMC increased its transparency about monetary policy via the use of forward guidance in its post-meeting statement, can we see also see a greater proportion of contents devoted to monetary policy (defined as references to current and future setting of the fed funds rate) in the minutes over time? And if so, is this greater proportion more visible during and after the Global Financial Crisis, under Bernanke's chairmanship, when the FOMC used a firmer and more explicit form of forward guidance language? The FOMC discusses a wide variety of topics in its meetings, such as inflation, labor market, economic projections, financial markets, among others, including monetary policy. The FOMC meetings (and so the minutes) follow a particular structure such that most of the topics are covered in each of them. Taking into consideration this aspect, before conducting the analysis I didn't expect significant variation in the proportion of topics over time, including the topic devoted to monetary policy, the one under analysis. However, I expected to see a slightly greater proportion of content devoted to such topic starting from the end of 2008, when the FOMC, under Bernanke's guidance, moved one step further in its peregrination toward policy transparency with greater use of its forward guidance tool in the post-meeting statements. That is, I expected that with increased transparency via greater use of forward guidance in the post-meeting statements during and post-crisis the FOMC members also devoted more time to discuss the current and future stance of the fed funds rate ("monetary policy").

The technique I applied for my analysis is able to algorithmically identify the topics discussed in the FOMC minutes and their proportions over time. Thus, another second but important goal of this thesis is being able to identify properly, via the use of LDA, such topics and their mixture over time. The LDA technique is particularly suitable for the text documents, such as FOMC minutes, that cover a large number of topics and that have a consistent structure over time. Although LDA, like other text mining techniques, has been widely applied for topic modeling in different fields, it has been less frequently applied in the economics literature, especially to Central Banks' documents. However, considering the advantages of its application, as I illustrate later, I believe LDA and similar text mining techniques will be largely employed in the future. To the best of my knowledge, the LDA technique has been applied to the FOMC minutes only in a recent paper by Jegadeesh and Wu (2017), although to a different period and within a different research question. This thesis thus contributes to the small but growing literature devoted to the use of text mining techniques to Central Bank's documents (and in particular, to the Federal Reserve's ones) for content analysis.

The rest of the thesis is organized as follows: in section II I present the Federal Reserve's milestones toward greater transparency as well as the main reasons underneath this global revolution in thinking and the literature concerning the application of text mining techniques to Central Banks' documents for content analysis; section III describes in detail the sample subject of analysis and the methodology applied (LDA); section IV presents some initial data exploration as well as the results of my analysis along with a commentary; section V provides a discussion of the main results; in the last sections I report the conclusion (VI), the references (VII) and the appendix (VIII).

II. Theory and Literature Review

i. The Opening of the Temple

Thirty-three years ago, William Greider (1987) wrote a very popular book, "The Secrets of the Temple: How the Federal Reserve Runs the Country", in which he described the Federal Reserve as a very powerful and yet mysterious and secretive institution.

> *"The Governors of the Federal Reserve decided the largest questions of the political economy, including who shall prosper and who shall fail, yet their role remained opaque and mysterious"* (Greiders, 1987, p. 2)

There was a popular belief that central banking was an esoteric art in the hands of a community of experts with skills and knowledge not available to the general public. During those years, secrecy and ambiguity were commonly known as part of Central Banks' culture. On the other hand, also central bankers believed their activities should be shrouded in mystery. Historically, many of the Fed officials opposed greater transparency. They opposed for long the release of the FOMC transcripts, mostly because they feared to cause unwanted consequences, such as volatility in the markets. This secrecy was not just a characteristic of the United States' Central Bank but also of other major Central Banks all over the world. For example, the 1921-1944 Governor of the Bank of England Montagu Norman's very famous motto was "never explain, never excuse" and he was not the only one with the same view on the matter. Many Central Bankers of those times believed that openness could damage the effectiveness of monetary policy. Indeed, secrecy was almost seen as monetary policy tool. Alan Greenspan, Chairman of the Federal Reserve from 1987 to 2006, called the "Maestro" because of his forecasting abilities, was very famous for his long, wordy and vague statements, practicing what at that time was called "constructive ambiguity"[2] and today "Fedspeak"[3] or "Greenspeak" and for his reluctance in moving toward greater policy transparency. *"Since I've become a central banker, I've learned to mumble with great incoherence. If I seem unduly clear to you, you must have misunderstood what I said"*, he once joked in a 1987 speech as quoted in the Wall Street Journal[4]. Later Greenspan (2007) admitted that he was aware of the influence his words had and he used this intentional strategy to prevent overreactions to his remarks by financial markets, and so to avoid a self-fulfilling prophecy. *"I took care, naturally, to couch any discussion of possible future moves in Fedspeak to keep from rolling the markets",* Greenspan (2007) said. He was referring to a time he was asked in 1993 about interest rates at one of the semi-annual reports to Congress. However, despite some concerns, starting

2 This term was first coined by Henry Kissinger, a former U.S. politician, for political matters.
3 Recently Fedspeak is also used to describe general Central Banks' talk.
4 Please, see Hookway, J. (2020). *In the Eternal Quest to Decode Fedspeak, Here Come the Computers.* The Wall Street Journal. Retrieved from https://www.wsj.com/articles/in-the-eternal-quest-to-decode-fedspeak-here-come-the-computers-11581287394

from the early 90s there was a growing tendency within the Fed toward greater openness[5]. In 1994 the FOMC not only began to publish the meeting transcripts, although with a five-year lag, but also to release immediately following a scheduled meeting the post-meeting statements reporting the members' latest policy decision. Indeed, before that year the Fed did not issue any statement regarding changes in the fed funds rate and the market had to infer such information closely monitoring the actions of the New York Fed trading desk. Afterward, in 1999, the FOMC also started to provide some indications of the expected or intended future path of policy rates in terms of "tilt" or "bias" and then "balance of risks" in their post-meeting statements. The Federal Reserve wasn't the only Central Bank moving in such direction. At that time, some form of forward guidance was already adopted by the Central Bank of New Zealand (from 1997) and Japan (from 1999). The latter was the very first Central Bank to have used forward guidance in the zero lower-bound condition (Contessi and Li, 2013). That is, welcomed by many, over the last thirty years many Central Banks drifted apart from the long-term tradition of secrecy and not only became noticeably more transparent but also started to rely more heavily on communication for monetary policy-making. The communication channels differ from Central Bank to Central Bank although generally include the release of timely monetary policy announcements, minutes and transcripts, reports, holding of post-meeting press conferences, speeches, among others. This revolution in thinking was motivated by different reasons. First of all, given that their actions do not merely affect financial markets but rather the life of every citizen, Central Banks act as public servants and so have the duty to be more open to the public about their decision-making process (Bernanke, 2004). Second, independent Central Banks should be accountable for their actions to their respective Government and, more generally, to the public. Third, as the notion of the importance of managing market expectations became increasingly widespread, many Central Banks started to use communication as a key tool to make monetary policy more effective and improve economic results, especially under exceptional conditions. Two of the first pundits to express their view about Central Banks communication were Blinder (1998, pp. 70-72) and Woodford (2001, p. 307) who advocated that higher transparency and communication can improve monetary policy by managing market expectations of Central Banks' setting of future short-term interest rates. As illustrated by many, such as Bernanke (2004) and Woodford (2005, p. 402), although the public tends to focus on short-term rates, such as the fed funds rate, long-term ones are much more important to economic activities as most investment and borrowing decisions depend on their value. However, there is a link between the two: long-term rates reflect market expectations of the future evolution of overnight rates (such as fed funds rate) and other related short-term rates. This is because long-term rates can be decomposed into two parts: the expected future short-term

[5] Especially after the many critics received by Henry B. Gonzalez. At that time, he was Head of the Housing Banking Committee and started a series of initiatives aimed at improving public accountability of the Fed. Among other things, he accused the Fed of keeping secret from the public for 17-years the existence of the transcripts of the meetings (as well as tapes). Alan Greenspan, the Chairman of those years, eventually decided to release the transcripts as he didn't want the Fed to be perceived as a Temple. For more information, please see: Todd, T. (2016). A corollary of accountability: A history of FOMC policy communications. Federal Reserve Bank of Kansas City. Retrieved from https://www.kansascityfed.org/~/media/files/publicat/acorollaryofaccountability/acorollaryofaccountability.pdf.

rates and a term premium[6]. Thus, as many economists argue, better communication about Central Bank's future short-rate setting contributes to the management of long-term ones and thus in the effectiveness of monetary policy. One of those advocates was Ben Bernanke. Bernanke et al. (2004) published an empirical analysis presenting some evidence that Central Bank communication can help manage public future short-term rates expectations, especially in the zero lower bound condition, when there is little or no room for further easing. Another empirical study of those years worth mentioning is the one of Gürkaynak et al. (2005) whose findings suggest that the FOMC can be able to shape longer-term rates and stimulate the economy by effectively communicating its future policy intentions. Eventually, central banking practices changed, embracing a new era of openness and greater communication.

ii. Text Mining Techniques to Central Banks' Documents for Content Analysis: Review of Literature

As described previously, over the last thirty years Central Banks have become remarkably more transparent about their policy-making, drifting apart from the tradition of "never explain, never excuse" and using communication as a key tool of monetary policy. Given the importance of the contents contained in the Central Banks' documents it is not surprising that these are heavily scrutinized by market participants as well as by academics. As a result, Central Banks' communication nowadays can boast a vast literature, the majority of which focuses on the impact of communication on financial markets or the link between different communication strategies and the economic performance (Blinder et al., 2008). Most of the empirical works in economics tend to focus on quantitative data and on the application of data mining, techniques that enable to discover patterns in large data sets. More recently there is also a small but growing branch that focuses on the application of text mining techniques on Central Bank's documents (also in conjunction with the application of data mining ones). As I explain better later, text mining refers to the wide set of computational tools and statistical techniques that are able to quantify text data, that is, to treat it quantitatively to extract valuable insights (Bholat et al., 2015, p.1). Within central banking, those techniques are mostly applied to documents to infer the tone (for instance, positive or negative and/or dovish or hawkish[7]) and/or the content. For my analysis I make use of a popular text mining technique, called LDA, to infer the proportion of FOMC minutes devoted to the topic of monetary policy over time. Thus, this thesis contributes to a small but growing literature devoted to the use of text mining techniques to Central Banks' documents (and in particular, to the Federal Reserve's ones) for content analysis. In the following paragraph I present the major works belonging to such literature[8], with a focus on the ones most similar to mine.

[6] Such premium reflects the remuneration investors ask for holding assets with longer maturities.

[7] See, for instance, Lucca and Trebbi (2009) who classified Fed statements as hawkish or dovish, Bligh and Hess (2013) who measured certainty, pessimism, and macroeconomic language of Alan Greenspan's speeches, testimonies and FOMC statements under his chairmanship and Kahveci and Odabas (2016) who observed the change in tone pre and post-crisis of monetary policy statements of the Federal Reserve, European Central Bank, and Central Bank of Republic of Turkey.

[8] The majority of which use in conjunction data mining techniques (mostly regressions), usually to measure the effects on the market.

One of the very first applications of text mining techniques to Central Banks' documents for content analysis can be found in Boukus and Rosenberg (2006). The authors applied another popular technique, similar to LDA, called Latent Semantic Analysis (LSA) to extract the different topics of the minutes from 1987 to 2005. After observing that the minutes release moves Treasury yields, they used LSA technique to analyze if such reaction depends also on the specific themes expressed[9] (combined with other factors), as they found out. Although they did not specifically label the topics identified by the LSA (they just mentioned the keywords per topic cluster) their work shows how text mining techniques are able to extract valuable information from central banks' documents. Indeed, this work was the very first application of LSA to the Federal Reserve's documents and set the stage for its future use. Hendry and Madeley (2010) performed a similar analysis to the statements of the Bank of Canada from 2002 to 2008, trying to identify the most valuable topics by financial markets. The authors found that such topics were the ones devoted to major shocks in the Canadian economy, the balance of risks to the economic projection and some forward-looking statements. One recent and well-known application of LSA is a paper by Acosta (2015). The author applied such technique (along with other text mining and data mining ones) to both the FOMC minutes and the transcripts, from 1976 to 2008, to analyze how accurately the minutes convey the information discussed during the meetings (and so, the transcripts[10]). Moreover, Acosta studied if such accuracy changed as the transcripts became public starting from 1994, that is, as the members knew their discussion would have become public, as well as how such change shaped the contents of the meetings. In order to analyze such research questions, he computed the "similarity" (mathematically, the cosine) between the transcripts and minutes via LSA (thus, avoiding labeling the topics[11]). The author's findings suggest an increase in transparency (proved by an increase in conformity) after the publication of the transcripts, partially driven by a changing of the contents of the transcripts. The most recent paper, to the best of my knowledge, which uses the technique of LSA for topic modeling is the one of Mazis and Tsekrekos (2017) which analyses the contents of the FOMC post-meeting statements (from 2003 to 2014) and their impact on the U.S. Treasury market (through a time series regression). The authors identified six recurring topics, being inflation, labor market, monetary policy accommodation (both in terms of interest rates and Large-Scale Asset Purchases), housing market, consumer spending, output growth, industrial production and financial markets (some of which, similar to the ones identified in my analysis). Mazis and Tsekrekos (2017) found the first four to have higher importance on the market. Their results suggest that the information content of the FOMC statements is more significant starting from the years of the Global Financial Crisis, also in line with the intention of those years of the FOMC, under Bernanke's guidance, to enhance its forward guidance policy

9 They used first order autoregressions for such analysis.

10 I recall that the minutes can be considered as a summary of the transcripts. The latter is the most detailed record of the FOMC meeting proceedings as they report the members' full conversations. Those, however, were not public until 1994. They are released only after a 5-year lag. The minutes, on the other hand, were published after a six-to-eight-week lag until 2004, when they started to be published after a three-week lag.

11 As I explain better later LSA, as well as LDA, does not provide topic labels, which must be instead identified by the user in accordance with the subject being studied. This aspect introduces some subjectivity in the analysis.

and transparency. As Boukus and Rosenberg (2006), the authors found that the U.S. Treasury yields' reaction also depends on the specific topic expressed in the FOMC statements.

The papers described so far all apply the Latent Semantic Analysis (LSA) technique. Very recently, also Latent Dirichlet Allocation (LDA), the method I chose for my analysis, has been applied to Central Banks' documents for content analysis. The very first application of such methodology in economics is the one of Hansen et al. (2017), who performed a similar study to the one of Acosta (2015), trying to infer how policy deliberations have evolved following FOMC members' knowledge of their discussion being public (within the period of 1987-2009). The authors distinguished between two types of discussion during the meeting, the FOMC1 (economic situation discussion, prepared in advance by each member) and FOMC2 (monetary policy strategy discussion) and computed the distribution of topics (forty different) to the discussions of each of the members (instead of at a document level), concluding that career concerns also shape policy-markers' transparency[12]. Another relevant work within LDA application to Central Banks' documents for content analysis is the one of Jagadeesh and Wu (2017). They studied the informativeness of the contents of the FOMC minutes from 1991 to 2015 for the stocks and bonds market. They observed that the most valuable topics for the market to be policy stance, inflation and employment, as they are strongly related to the magnitude of price change of bonds and stocks. They also computed the tone of the entire minutes and of each topic so to examine their directional impact and found that the latter depends both on the tone and on the topic. For my analysis, as Jagadeesh and Wu (2017), I applied LDA to the FOMC minutes to extract the different topics and compute their proportions over time. However, my work is different from theirs as I addressed a different research question and a different period. Moreover, I made different choices in the application of LDA. For instance, differently from them, I did not include the section of the minutes devoted to the staff overview and outlook of the economy and while they considered eight to be the appropriate topic number for their sample I privileged nine for mine. The most recent application of LDA to Central Banks' documents for content analysis is the one of Edison and Carcel (2020). The authors applied such technique to the FOMC transcripts from 2003 to 2013 to simply analyze the evolution of their contents over time. Their results show an increase of the topic devoted to economic modeling as well as the one devoted to the banking system during the Global Financial Crisis and that discussion devoted to communication increased in the last three years of the sample. Although the minutes can be considered as a summary of the transcripts, and that both my analysis and Edison and Carcel's one address the same period, the results are difficult to compare because of the different choice of the number of topics and of the respective labeling.

[12] Besides LSA, which was used to extract the topics from the FOMC minutes and transcripts, the authors used a wide set of data mining tools (such as LASSO) for their analysis.

Another work worth mentioning within such literature is the one of Vallès and Schonhardt-Bailey (2015), who used the Alceste Software[13], which, as LSA and LDA, is able to detect the latent topics of a set of documents. They studied how the contents discussed by the Bank of England's Monetary Policy Committee (MPC) during their monthly meetings and by their members were shaped by the adoption of forward guidance in 2013. In particular, they focused on the minutes and speeches of the final year of Mervyn King's Governorship, when forward guidance language was not yet used, and the first year of Mark Carney's, when such language was used for the first time. Their results show that the adoption of forward guidance did not change significantly the contents discussed by the committee as a whole (As also emerges in my analysis. This is because, as for the FOMC minutes, the BoE minutes have consistent contents over time) whereas there is a divergence between internal members, who also spoke to the topic of forward guidance, and external ones, who avoided such discussion. My analysis is similar to Vallès and Schonhardt-Bailey's one as I also focus on how the adoption of forward guidance shaped the contents of the FOMC discussions.

In conclusion, the literature here reported shows the benefits from the application of text mining techniques to Central Banks' documents, as they offer valuable tools for the extraction of the contents from text documents and for their transformation into quantitative data that can be used to extract valuable insights.

[13] Alceste is a textual statistic software capable of describing, classifying, and automatically synthesizing text. However, differently from LSA and LDA who can be freely applied, Alceste software is paid.

III. Methods

This section is composed of three parts: in the first one I provide a brief description of the data subject of analysis (FOMC minutes); in the second I present the pre-processing passages required before LDA implementation; in the last one I explain in better detail text mining, LDA algorithm and Gibbs sampling (the method chosen to apply LDA) and then I present the LDA application and the creation and visualization of the topic proportions over time. Below a chart that summarizes all steps implemented for my analysis that I describe along this section. A list of all functions used can be found in Appendix 3: Coding Used for Analysis.

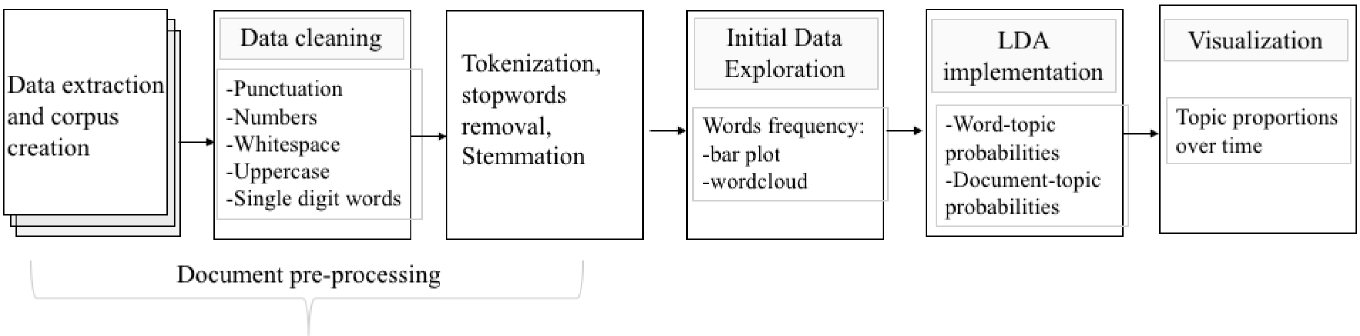

i. Description of the Data

To address my research question, I analyzed the minutes of all the scheduled FOMC meetings from 2002 to 2013. The minutes of the Federal Open Market Committee are an accurate summary of the contents discussed by the participants during the so-called "FOMC meetings". They are used as one of the key communication tools for informing Congress, markets and public about the monetary policy decision process. In particular, the minutes convey a detailed insight about all the relevant topics members consider during policy-making, such as their views on the economy; their policy decisions and supporting reasons; as well as the voting process and dissenting views. The FOMC started to draft the minutes since its earlier days, in 1936, although in different forms (and names) over time[14]. The minutes were initially kept confidential until 1967 and only in 1993 they started to provide more detailed content (Danker and Luecke, 2005). From 1993 to 2004 the minutes were published after the subsequent meeting, so after a lag of six to eight weeks. However, to satisfy the public's desire for a more-timely communication, in December 2004, the Committee decided to reduce the time lag of the release to three-weeks, as it is also today. The minutes are characterized by a common language and by a clear and recurring structure made up of four major sections. In the first section we can find administrative matters, including the list of participants, annual organizational matters if it the first meeting of the year, and a review of foreign currency and previous open market operations. The second section, which is prepared before the meeting, covers the staff overview and outlook of the economic and financial situation. In the third section all meeting participants, that is both the seven Governors and the twelve Presidents of the Regional Banks, convey their view on economic

[14] For additional information about the history and structure of the FOMC minutes you can refer to Danker and Luecke (2005).

conditions and outlook, common topics typically being inflation, consumer spending, business investment, economic growth, labor market, among others. The last section summarizes the Governors' discussion on policy decision, outlook for future monetary policy and on the contents of the post-meeting statement language. It also includes the vote as well as dissenting views on the policy action and on the post-meeting statement that will be released shortly after the meeting and the directive to the Fed of NY (Danker and Luecke, 2005), as well as, although infrequently, one last additional discussion topic of some changes in the communication policy. Starting from October 2007, the minutes also provide four times each year a Summary of Economic Projections (SEP), an addendum conveying a summary of economic projections and assessment of future monetary policy by FOMC meeting participants. I have decided to focus my analysis on the FOMC minutes instead of the post-meeting statements or transcripts because of the minutes' relatively better balance between information content and timing. Although the minutes are released after a three-week lag, whereas the statements immediately after the FOMC meeting, the minutes provide a much richer information content about the meeting than the statements. The transcripts, on the other hand, although provide the most detailed record of the meetings, are released after a five-year lag, and thus lose power in terms of communication relative to the minutes. Moreover, I decided to examine the minutes instead of the transcripts because the formers are characterized by a recurring structure and use of words, making them, more appropriate text documents candidates for an LDA analysis. The FOMC holds eight scheduled meetings during each year and other meetings, such as conference calls, if needed[15]. For my analysis, I used only the scheduled meetings, for a total sample of 96 meeting minutes. I chose to investigate the period of 2002-2013 as it was marked by greater policy transparency via the use of the forward guidance tool as well as by the leadership of two Fed Chairmen with a different perspective on transparency. In particular, I included also year 2002 to analyze the difference in contents between 2002, the year still characterized by the "balance of risks" language and 2003, the starting year of the forward guidance language.

ii. Document Pre-processing and Initial Data Exploration

I freely downloaded each minutes between 2002 and 2013 from the Federal Reserve Board of Governors' official website [16] and then transformed each file into a txt format, and saved each with the corresponding date. Before starting to process the data, I removed manually[17] from all the minutes the first and second sections, together with the voting part of the last section as I did not deem them to be useful for my analysis. In particular, I decided to remove the first section since unlikely to provide any relevant information while

[15] Not scheduled meetings and conference calls minutes are much lower in length and do not have a consistent structure. For this reason, they were not used for my analysis.

[16] You can find the historical meeting minutes and the most recent ones at the following links, respectively:
https://www.federalreserve.gov/monetarypolicy/fomc_historical_year.htm and
https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

[17] I preferred not to use any parsing algorithm to remove the unneeded parts since not all the minutes respected the same entering quotes for each section.

I removed the second one, since, like Edison and Carcel (2020), I was only interested to analyze the contents discussed by the Governors (and so, not the data and reports provided by the staff). Lastly, I did not include the last part of the fourth section, that is, the voting actions, the directive to the Fed of NY and the voting for the approval of the post-meeting statement since not relevant. Generally, my text analysis starts with the line "In the Committee's discussion of current and prospective economic conditions…" or "In their review of current and prospective economic developments, members commented…" or "In their discussion of the economic…". I then broke each minutes into different paragraphs, keeping record of the specific minutes where a paragraph belongs[18]. A total sample of 818 documents (paragraphs) was created. To perform my analysis I used Rstudio, a free programming language. The first steps consisted of the download of the required packages[19], load of the paragraph documents, and creation of the corpus, that is, a collection of the text documents under mining. Then, as suggested by the text mining guide provided by Bholat et al. (2015, p. 5-8), I performed some pre-processing passages very common in any text mining analysis to prepare the documents in a format suitable for analysis. That is, I removed the numbers, punctuation, excess whitespace, single-digit words and converted uppercase to lowercase. I proceeded by breaking the corpus into individual tokens or words[20], a process known as "tokenization", so that one token-per-document-per-row table was created. The latter format, known as "tidy text" format, is particularly useful for further text processing as it allows to use some common text mining packages. Afterward, I removed from the documents the so-called "stopwords", very frequent words that do not provide any insight about the contents of documents, such as *"the", "and", "of"* as well as articles, pronouns, conjunctions, among others (for a total of 1,149 stopwords), as well as an additional list of stopwords ("mystopwords") with meaningless words commonly seen in the FOMC language, such as *"committee"* and *"participants"* (the complete list of "mystopwords" is in table 4 of Appendix 1: Additional Tables). I then stemmed the words[21], that is, removed the suffix or prefixed of words and maintained just their root forms so that the algorithm can recognize them as the same word. I then removed the "stopwords" and "mystopwords" a second time to eliminate unwanted words that may have been surfaced after stemming (an example of the entire cleaning process is provided in Appendix 2: Example of a Text Cleaning Process). At the end of such pre-processing steps the number of total words was more than halved, from an initial of 185,249 to 84,404[22]. To initially explore the data and so investigate the documents' most frequent words I created a Term Document Matrix (TDM), a matrix where each row represents one unique word, each column one document and each value the frequency of a word in a document. The matrix created is characterized by 2,437 unique words and 818 documents (that is, the number of total paragraphs created). I used this matrix to generate a frequency table

18 Splitting the documents into different paragraphs improves the performance of the LDA in the topic identification as paragraphs (although not all of them) are devoted to a lesser number of topics.
19 You can find a list of the packages used in Appendix 3: Coding Used for Analysis
20 A token is defined as any meaningful unit of a text and this can be a single word, or sentence or paragraph.
21 I used Martin Porter's stemming algorithm and the Snowball's C libstemmer library
22 Note that this number refers to the total number of words present in the corpus after the pre-processing steps which is different from the number of unique words of the corpus (that is, number of words distinct from each other. As explained later, the unique words will make up the rows and columns of the Term Document Matrix and the Document Term Matrix, respectively)

and a word cloud whose results are presented later in the Exploratory Results: Words Frequency section of this thesis. The next steps consist of the creation of the Document Term Matrix and of the application of the Latent Dirichlet Allocation algorithm that I illustrate after providing an overview of the text mining and LDA techniques.

iii. Latent Dirichlet Allocation (LDA): a Text Mining Technique

To examine the impact of greater policy transparency on the FOMC minutes over time I decided to apply the Latent Dirichlet Allocation (LDA), an advanced topic modeling algorithm first developed by Blei, Ng and Jordan (2003). Before explaining the remaining passages of my analysis, I will first give a brief overview of text mining, LDA, and Gibbs sampling (the method used to implement LDA) and my choice of the number of topics.

a) An Overview of Text Mining and its Application to Central Banks Documents

Text mining, also known as Natural Language Processing (NLP), refers to the wide set of computational tools and statistical techniques that are able to quantify text data, that is, to treat it quantitatively to extract valuable insights (Bholat et al., 2015, p.1). They have been widely applied in many social science literatures, such as politics and marketing, but have been less frequently applied in economics, especially within the Central Bank studies (Bholat et al., 2015, p.1), although it recently captured increasing interest in such field. Indeed, most of the empirical works in economics tend to focus on quantitative data and to apply data mining, techniques that enable to discover patterns in large data sets. However, there are many gains also from the application of text mining techniques. First of all, they enable to analyze large sets of text data in a relatively shorter time and also enable to extract valuable information that may be disregarded by a human reader (Bholat et al., 2015, p.1). Second, they can be used for a wide range of purposes, the most common being information retrieval, content comparison, topic classification and modeling, among others. In such a context, observations, instead of being numbers, are text data, also known as documents. Examples of documents eligible for text mining analysis can be newspapers, or just newspaper titles, emails, scientific papers, books, documents, Twitter archives, among others. Candidate text data for central banking analysis are Chairmen's and Governors' speeches, meeting materials, reports, presentations, etc. The text data could be a single document or a collection of documents, known as corpus, depending on the analysis undertaken. Text mining techniques can be distinguished between supervised and unsupervised. In the context of content analysis, the former requires the user to classify the documents so to train the algorithm while the latter do not require classified documents. Another distinction is between abductive techniques, which draw conclusions for a particular case from the data under analysis without then generalizing to other cases, and deductive, which try to validate a general theory testing a particular dataset. Latent Dirichlet Allocation (LDA), the technique I applied for my analysis, belongs to the unsupervised and abductive techniques. Bholat et al. (2015) provide a useful guide with explanations of the most common text mining techniques, including LDA, applied in the most recent Central Banks' studies, with a focus on unsupervised ones, as well as the common pre-processing text mining steps, as the ones I applied.

b) The Latent Dirichlet Allocation (LDA) Algorithm and Gibbs sampling

First developed by Blei et al. (2003), Latent Dirichlet Allocation (LDA) is one of the most popular algorithms for topic modeling. Topic modeling refers to the unsupervised machine learning techniques that classify documents into a user-chosen number of topics by identifying clusters of similar word patterns of such documents[23]. Topic modeling assumes words are not independent of one another but rather are linked together by latent topics based on their appearance alongside each other. LDA is an unsupervised generative probabilistic topic model algorithm that is able to detect underlying topics (also known as latent variables) in a collection of text documents (that is, a corpus). The word "Latent" stands for hidden, in this case the topics; "Dirichlet" refers to the distribution of distribution, in this case each document can be described as a probabilistic distribution over latent topics, each of which is a probabilistic distribution over words; and "Allocation" refers to the action of allocating topics to the documents and words to the topics. LDA is a powerful tool, able to analyze the content of a large number of unclassified documents. It requires the number of topics T to be defined a priori and to be fixed for the entire process (I later explain my choice of the number of topics). As described by Schwarz (2018), LDA relies on two baseline assumptions. The first is that documents with similar topics will use a similar group of words. The second assumption is that each document (in our case paragraph) can be described as a probabilistic distribution over latent topics and each topic as a probabilistic distribution over words. In other words, each document is a mixture of topics and each topic is a mixture of words. Words identified by the model with a high probability of belonging to a certain topic usually give a good insight into the content of such topic. Therefore, if a word has a high probability of belonging to a certain topic, the documents with such word will more probably belong to such topics as well. Before going into the detail of the application of LDA for my analysis I will make use Schwarz (2018) paper to illustrate how the probability distribution of LDA works as well as its generative process.

Given a corpus consisting of D paragraphs (also referred to as documents) and T topics (T chosen in advance), each paragraph d of corpus D is a probability distribution over T topics. Paragraph vector $\theta_d = [\theta_{d,1}, \cdots, \theta_{d,T}]$ of length T contains such probabilities. The output of LDA will be a matrix θ of dimension $D \times T$, where $\theta_1, \ldots, \theta_D$ are row vectors $1 \times T$, and whose entries ($\theta_{d,t}$) represent the probability of paragraph d to belong to topic *t:*

[23] Within the data mining field it can be compared to the clustering techniques, such as k-means clustering.

$$\theta = \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_D \end{pmatrix} = \begin{pmatrix} P(t_1 \mid d_1) & \cdots & P(t_T \mid d_1) \\ \vdots & \ddots & \vdots \\ P(t_1 \mid d_D) & \cdots & P(t_T \mid d_D) \end{pmatrix}$$

Given a vocabulary consisting of V unique words from the corpus, each topic t belonging to T is a probabilistic distribution over the vocabulary *V*. Topic vector $\emptyset_t$, of length *V*, contains such probabilities. The second output of LDA will be matrix $\emptyset$ of dimension $V \times T$, where $\emptyset_1, \ldots, \theta_T$ are column vectors $V \times 1$, and whose entries ($\emptyset_{v,t}$) represent the probability of observing word *v* from the vocabulary conditional on topic *t*:

$$\emptyset = (\emptyset_1, \cdots, \emptyset_T) = \begin{pmatrix} P(w_1 \mid t_1) & \cdots & P(w_1 \mid t_T) \\ \vdots & \ddots & \vdots \\ P(w_V \mid t_1) & \cdots & P(w_V \mid t_T) \end{pmatrix}$$

Given the parameters θ and $\emptyset$, the LDA will infer the T topics assuming a generative process whereby paragraphs of N_n words are generated. In simple words, LDA works backward as if it were to construct the paragraphs by itself. More precisely, first a word probability distribution $\emptyset$ is sampled from a Dirichlet distribution[24] ($\emptyset \sim Dir\ (\beta)$) . Then, for each paragraph *d* of the corpus topic proportions are sampled from a Dirichlet distribution ($\theta \sim Dir\ (\alpha)$). For each of the N_d words w_d in paragraph *d* a topic is sampled from a Multinomial distribution ($z_{d,n} \sim Mult\ (\theta_d)$) and a word $w_{d,n}$ is drawn, conditioned on the $z_{d,n}$ topic, from a multinomial distribution[25] ($p\ (w_{d,n} \mid z_{d,n}, \emptyset)$)[26]. α and β are hyperparameters which determine the sparsity of the Dirichlet, and their prior values are required for the Gibbs sampling process, the method I chose to apply among others for the LDA process[27]. Given this probabilistic model, the following is the likelihood of the corpus of paragraphs with respect to the parameters θ and $\emptyset$:

$$\prod_{d=1}^{D} P\ (\theta_d|\alpha)\ \left\{ \prod_{n=1}^{N_d} \sum_{z_{d,n}} P\ (z_{d,n} | \theta_d)\ \ P\left(w_{d,n}|z_{d,n}, \emptyset\right) \right\},$$

where:

- $P\ (\theta_d|\alpha)$ represents the probability of observing the topic distribution of θ_d of document d conditional on α;

[24] A Dirichlet Distribution is a distribution of positive probabilities whose sum is equal to one. It is the generalization of a Beta distribution.
[25] A Multinomial Distribution is the generalization of the Binomial distribution (i.e. when there are more than two outcomes).
[26] In simpler words, the LDA first randomly assigns for each word in each document a topic. Then, for each document, the LDA draws a distribution over topics; for each word in the document it draws a topic from the distribution over topics and a word from such corresponding topics.
[27] In other words, LDA uses two Dirichlet distributions, each with its own hyperparameter. α is the hyperparameter for the distribution of topics over documents: the higher its value the more likely each document will contain a mixture of most of the topics (instead of just one or few). β is the hyperparameter for the distribution of words over topics: the higher its value the more likely each topic will contain a mixture of most of the words (instead of just few specific words).

- $P\left(z_{d,n} | \theta_d\right)$ represents the probability of the topic assignment $z_{d,n}$ of word n in document d conditional on the topic distribution of the document;
- $P\left(w_{d,n} | z_{d,n}, \emptyset\right)$ represents the probability to observe a specific word d conditional on both the topic assignment of such word and on the word probability of the topic.

LDA tries to find the optimal topic assignment $z_{d,n}$ per word per document and the optimal word probability $\emptyset$ per topic that maximizes this likelihood, which however is computationally unfeasible. Thus, some other methods for the LDA such as the Gibbs sampling have been developed[28]. Gibbs sampling is a Monte Carlo Markov-chain algorithm, where Markov-chain refers to the action of sampling from each variable one at a time, keeping fixed the current values of the other variables. Within LDA application, without going into the math, the Gibbs sampling iteratively changes the topic assignment of one word, conditional of the topic assignment of all the other words. In particular, the method will first randomly assign a topic to each word in the corpus of documents (as said before, the number of topics must be chosen beforehand). As a consequence, each document will be randomly assigned to some topics. Afterward, the Gibbs sampling will assume that all the word assignments are correct but the current one in document d. The Gibbs sampling will assign a new topic to such word following a probabilistic model: it multiplies the proportion of words in document d that are currently assigned to topic t (P $(topic\ t \mid document\ d)$) to the proportion of assignments to topic t over all documents coming from that word (P $(word\ w \mid topic\ t)$). That is, in assigning the new topic to the current word, the Gibbs sampling will take into account both the proportion of words of document d assigned to topic t (excluding the current word) and the proportion of assignments to topic t over all documents that came from this word w. This process will be repeated for each word w of each document d of corpus D until a steady-state of word topic assignments is reached (i.e., when the assignments will stop changing). The outputs of such process are the word-topic probabilities and the document-topic probabilities. As already mentioned before, one of the two assumptions of LDA is that each document can be described as a probabilistic distribution over latent topics and each topic as a probabilistic distribution over words. Thus, if a word will have a high probability of being assigned to a topic and if a document contains such word, then also the document will have a higher probability of belonging to such topic as well.

In summary, the power of LDA consists in its capability of clustering words belonging to the same topic without using any user pre-defined clusters of words. LDA only requires the number of topics T to be defined a priori. This advantage is what makes popular such method with respect to Boolean and dictionary-based techniques[29] (Bholat et al., 2015, p.11). LDA isn't the only method with such advantage as also Latent Semantic Analysis (LSA) does not make use of any pre-defined cluster of words. As I illustrate in the literature review, LSA is another popular algorithm used for topic modeling which uses Singular Value

28 Gibbs sampling has been developed by Xuan-Hieu Phan et al.
29 Dictionary methods require the user to pre-define a set of words of interest to identify the clusters of topics.

Decomposition (SVD) on the Document Term Matrix. Both LDA and LSA assume words are not independent of one another but rather are linked together by latent topics based on their appearance alongside each other, differently from other text mining techniques (such as Vector Space Models[30]), and both do not require pre-defined clusters of words. However, the LSA works better when documents and words are centered on a single topic while LDA has a distinguishing feature of being a probabilistic method, in which words and documents can be assigned to multiple topics (Bholat et al., 2015, p.12). Given the multi-topic nature of both the words[31] and documents being analyzed (that is, the FOMC minutes), the LDA was the most suitable method.

c) The Choice of the Number of Topics

As explained before, the LDA requires the number of latent variables, that is the topics *T*, to be defined ex-ante by the user. This choice represents one of the main challenges of LDA use as there is a trade-off between choosing few topics, with the risk of mixing different topics and loose interpretability, and choosing too many, with the risk of creating too specific topics and losing the general picture. Generally, the number of topics is chosen arbitrarily by the researcher based on the context, although there are some statistical methods (such as k-means clustering) that are able to guide the choice. As Edison and Carcel (2020), I chose the number of topics arbitrarily, after a careful analysis of the content of the minutes over time. Identifying a precise number of topics for the FOMC minutes is not simple as it strongly depends on the research's judgment and knowledge on the subject (as well as on the ability of the user of correctly applying the LDA). In general, the FOMC minutes touch the following macro topics[32]: inflation, economic outlook, current and prospective economic developments, economic projections, labor market, monetary policy, financial markets, consumer spending, business investments, housing market, fiscal policy, communication, international economy, asset purchase program[33]. To choose the correct number of topics for my analysis, I performed the LDA process using a different number of topics, from eight to twelve. The aim of this step was to find the number of topics such that the top words from each topic cluster were as distinct as possible so that I was able to identify without doubts a topic labeling from such words. I identified nine as the most suitable number of topics for my analysis.

[30] They are models that represent text documents as vectors whose entries are the words of the documents. Such models are used to measure the similarity of topics between documents. One of their limitations is that they are not able to recognize documents with similar context that use a different term vocabulary. This is because, differently from LSA and LDA, they rely on the assumption that words are independent of one another.

[31] For example, the word "policy" can either be associated with the topic of fiscal policy or the one about monetary policy.

[32] This doesn't mean that every single meeting touch all the topics but just that the meetings touch more or less consistent content over time.

[33] The difference (although slightly) between economic outlook and current and prospective economic developments consists in that the former reports general conditions of the economy while the latter links such economic conditions to the other topics (such as inflation, monetary policy, fiscal policy, etc).The housing market topic includes also house construction.

d) Application of Latent Dirichlet Allocation (LDA) on FOMC Minutes

As explained before, the first steps of my analysis consisted of applying some of the most common pre-processing text mining functions to clean the corpus, tokenize, and stem the words. Those passages enabled to prepare the documents in a format suitable for analysis. Along this section, I continue illustrating the remaining part of the analysis that I conducted following the passages (and functions) illustrated by a popular free text mining book (Silge and Robinson, 2020, ch. 6). Before applying the LDA I needed to transform my text data into a Document Term Matrix (DTM) as this is the object of the function `LDA()` from the *topicsmodel package*, the one I decided to use for LDA implementation. A Document Term Matrix (DTM) is a matrix where each row represents a document (in this case, paragraph) of the corpus, each column a unique term (that is, in the columns there are all the unique words present in all the documents) and where the entries represent the frequency of the terms that occur per each document. That is, a DTM is just the transpose of the Term Document Matrix (TDM) illustrated before. The Term Frequency (TF) weighting is not the only possible method for term-weighting. Another common method is the one used by Boukus and Rosenberg (2006), Hansen (2017) and Acosta (2015) in their researches, the Term Frequency Inverse Document Frequency (TF-IDF) which is the product of word's Term Frequency (TF), defined as the frequency of a word divided by the total number of words in the document, and the Inverse Document Matrix (IDM), defined as the logarithm of the total number of documents divided by the number of documents containing that term[34]. TF-IDF provides a measure of the importance of a word in a document. In simple words, this weighting-scheme places more weight to words that appear more frequently in a document and less weight to words very frequent to all documents. However, I decided to not use this weighting scheme because the content (and so the words used) of the documents subject of analysis are very similar (indeed, each document has a high number of topics associated with it) and the aim of the analysis doesn't consist of finding unique topics per document but rather topic proportions over time. The dimension of the Document Term Matrix obtained is 818 paragraphs times 2,437 unique words. Afterward, I implemented the LDA algorithm using as object the obtained matrix and as method the Gibbs sampling, the one I illustrated before, to create a 9 topic LDA model[35]. Then, I used the function `tidy()` from the *tidytext package* to calculate what the function calls "betas", the expected word-topic-probabilities, that is the expected probabilities of words of belonging to a certain topic (what I previously denoted as $\emptyset_{v,t}$). For each topic *t*, the model will compute the probability of each word being generated from that topic. Of course, as previously illustrated, the higher the probability for a given word the more likely such word characterizes topic *t*. Considering that the corpus under analysis has 2,437 unique words and the number of topics chosen is 9, the model generated 21,933 probabilities, most of which with a value close to zero. Table 1 presented in Topics Labeling section reports the top 20 words for each topic (that is, the ones with the highest probabilities).

34 The product will be zero or very close to zero for very frequent words among the documents of the corpus.
35 For this study I used 1,000 burn–in periods and 1,000 iterations.

e) The Estimation of the Proportions of Topics Over Time

Besides modeling the "betas", the `tidy()` function also estimates what it calls the "gammas", which represent the expected document-topic-probabilities, that is the expected probabilities of documents (in this case, paragraphs) of belonging to a certain topic (what I previously denoted as $\theta_{d,t}$). For each paragraph *d*, the model will compute the probability of belonging to each of the topic *t*. The sum of such probabilities for each paragraph is one and so such probabilities can be interpreted as the proportion devoted to a certain topic. Thus, a higher probability means a higher proportion devoted to a certain topic. Considering that the corpus under analysis is composed of 818 paragraphs and the number of topics chosen is 9, the model generated 7,362 probabilities. I used the obtained probabilities ("gammas") to identify the topic proportions per document over the years 2002-2013. To do so I aggregated the gammas per year, normalizing by the different number of paragraphs per year. Then I obtained a second figure which illustrates only the topic of Monetary Policy & Communication, the one of interest. The two figures are shown in the next section.

IV. Results

In the following section I first present some pre-LDA data exploration results and then the ones obtained with the LDA application.

i. Exploratory Results: Words Frequency

As explained before, the first steps of my analysis consisted of applying some of the most common pre-processing text mining functions in order to clean the corpus, tokenize, and stem the words. Those passages enabled to prepare the documents in a format suitable for analysis. After these pre-processing steps, before starting the LDA process, I proceeded by creating a Term Document Matrix, made up of 2,437 unique words and 818 documents, with the aim of initially exploring the data. Therefore, I used the matrix to generate a frequency plot and a word cloud, a popular text mining representation of words, where the frequency of each word is shown with font size and color.

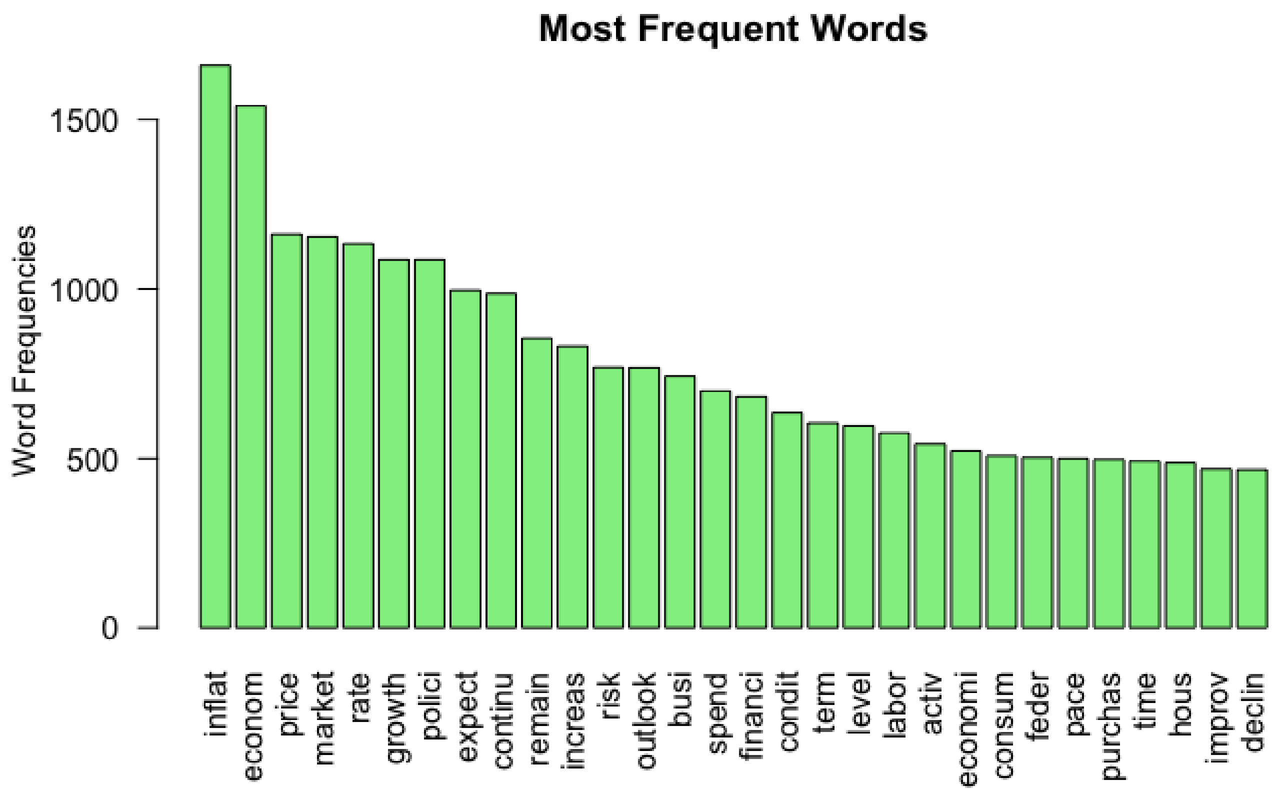

Figure 1: the 30 most frequent words

Figure 2: Word Cloud

The size and color of each word indicate the frequency

As can be seen from the two figures, the seven most common words in the FOMC minutes are "inflation", "economy", "price", "market", "rate", "growth" and "policy", words closely related to some of the most discussed topics of Central Banks, that is, inflation, economy, and monetary policy. Other very frequent words are "business", "spending", "financial" and "labor", the first two related to key sectors, the third to financial markets, and the last to labor market topics. It is not surprising that the most frequent word in the minutes in the considered period is "inflation" (and the third one "price"). Indeed, price stability, generally interpreted as low and stable inflation, is one of Federal Reserve's two mandates (together with maximum employment)[36]. Another reason why it is not surprising is that Alan Greenspan and Ben Bernanke, the two Chairmen during the period analyzed, were inflation fighters first and foremost as they both accredited the benefits of low and stable inflation.

ii. Topics Labeling

As illustrated before, I implemented the LDA algorithm with the aim of splitting the corpus into nine distinguishing topics. The number of topics (in this case, nine) must be pre-defined before the application. As we can see from table 1, LDA was able to cluster words belonging to the same topics without using any user pre-defined clusters of words. This characteristic is what makes this method so powerful and popular.

[36] The 1913 Act that established the Federal Reserve System did not explicitly provide with a monetary policy mandate but rather directed the Fed to generally promote economic and financial stability in the country. Following the Great Depression and World War II, Congress passed in 1946 the *Employment Act* specifying as goals "maximum employment, production, and purchasing power". During the 1970s, the U.S. economy entered into a period characterized by high inflation and unemployment known as "Stagflation", calling for a change in the conduct of monetary policy. Congress then passed in 1977 the *Reform Act* which promoted the goals of "maximum employment, stable prices, and moderate long-term interest rates". Although it specifies three goals, only the first two are explicitly discussed. In particular, maximum sustainable employment and price stability are commonly known as the Fed's dual mandate. More on Fed's history at the following website: www.federalreservehistory.org

LDA manages to do so by assuming that words are linked together based upon their appearance alongside each other.

Topic 1: Asset Purchase Program		Topic 2: Inflation		Topic 3: Labor Market		Topic 4: Fiscal Policy & International Economy		Topic 5: Current and Prospective Economic Developments & Economic Projections	
Word	**Beta**	**Word**	**Beta**	**Word**	**Beta**	**Word**	**Beta**	**Word**	**Beta**
purchas	0.0564	inflat	0.1451	labor	0.0585	growth	0.0798	econom	0.0700
addit	0.0353	price	0.0906	market	0.0379	economi	0.0452	rate	0.0437
asset	0.0319	expect	0.0558	employ	0.0348	fiscal	0.0346	expect	0.0392
term	0.0274	increas	0.0337	rate	0.0306	econom	0.0341	outlook	0.0382
secur	0.0230	pressur	0.0327	declin	0.0271	activ	0.0315	anticip	0.0287
agre	0.0223	energi	0.0247	unemploy	0.0264	recoveri	0.0202	level	0.0278
program	0.0218	cost	0.0245	continu	0.0261	demand	0.0197	run	0.0270
feder	0.0204	measur	0.0227	level	0.0261	support	0.0197	percent	0.0257
reserv	0.0199	remain	0.0216	suggest	0.0217	export	0.0185	continu	0.0253
outlook	0.0187	core	0.0212	low	0.0202	data	0.0158	pace	0.0250
pace	0.0186	resourc	0.0194	increas	0.0201	effect	0.0157	growth	0.0239
time	0.0165	risk	0.0171	remain	0.0176	contribut	0.0153	moder	0.0222
accommod	0.0159	rise	0.0155	product	0.0173	foreign	0.0149	gradual	0.0207
polici	0.0153	commod	0.0151	indic	0.0167	incom	0.0143	inform	0.0194
effect	0.0153	concern	0.0129	growth	0.0158	govern	0.0143	project	0.0191
monetari	0.0143	moder	0.0126	slow	0.0148	restrain	0.0134	consist	0.0191
stabil	0.0134	term	0.0116	improv	0.0137	uncertainti	0.0124	term	0.0179
potenti	0.0119	slack	0.0105	output	0.0128	stimulu	0.0119	remain	0.0158
hold	0.0117	read	0.0103	factor	0.0123	reduc	0.0117	judg	0.0154
condit	0.0114	util	0.0100	gain	0.0120	balanc	0.0110	inflat	0.0143

Topic 6: Financial Markets		Topic 7: Economic Outlook		Topic 8:Key Sectors		Topic 9: Monetary Policy & Communication	
Word	**Beta**	**Word**	**Beta**	**Word**	**Beta**	**Word**	**Beta**
market	0.0859	outlook	0.0251	spend	0.0536	polici	0.0917
financi	0.0771	econom	0.0248	busi	0.0446	econom	0.0401
condit	0.0431	expans	0.0247	consum	0.0375	rate	0.0393
credit	0.0372	uncertainti	0.0187	hous	0.0374	feder	0.0300
bank	0.0254	economi	0.0182	sector	0.0307	fund	0,0299
risk	0.0212	prospect	0.0171	household	0.0258	monetari	0.0256
increas	0.0186	includ	0.0171	invest	0.0246	risk	0.0239
remain	0.0171	signific	0.0171	report	0.0235	statement	0.0223
improv	0.0169	factor	0.0168	sale	0.0206	agre	0.0220
loan	0.0152	evid	0.0166	continu	0.0187	futur	0.0164
eas	0.0136	strengthen	0.0159	incom	0.0160	stanc	0.0156
declin	0.0127	persist	0.0151	price	0.0158	accommod	0.0150
tighten	0.0101	anticip	0.0146	inventori	0.0148	forward	0.0150
concern	0.0099	develop	0.0145	home	0.0147	target	0.0142
real	0.0086	comment	0.0142	capit	0.0142	assess	0.0139
strain	0.0086	product	0.0137	contact	0.0140	commun	0.0136
continu	0.0084	effect	0.0135	activ	0.0117	action	0.0114
deterior	0.0084	posit	0.0128	construct	0.0111	guidanc	0.0109
liquid	0.0083	continu	0.0125	mortgag	0.0109	decis	0.0106
investor	0.0083	provid	0.0125	firm	0.0104	view	0.0104

Table 1: Distribution of the Top 20 Words per Topic
The beta measures the probability that a word belongs to topic *t*

For each topic, table 1 reports the top twenty words with the highest "betas", the probabilities of words of belonging to a certain topic. As illustrated previously, the higher the probability for a given word the more likely such word characterizes topic *t*. While on one hand LDA manages to cluster words belonging to the same topics, on the other it does not provide topic labels. The latter must be instead identified by the user in accordance with the subject being studied. As we can see from the table, the top words from each topic cluster are mostly distinct, enabling us to recognize without doubts the relevant topic labels for each of them. The keywords of the first cluster consist of *purchase, assets, security, program, federal, reserve, accommodation, policy, monetary,* all words clearly associated with the Asset Purchase Program. In the second cluster we can see words clearly identified with the topic of Inflation, such as *inflation, price, expectations, pressure, energy, cost, core, commodity, slack*. *Labor, market, employment, rate, unemployment* are all words commonly associated with the topic of Labor Market (third cluster). The keywords of the fourth cluster, *fiscal, support, government, stimulus,* and *economy, demand, export, foreign* indicate the topic of Fiscal Policy and International Economy, respectively. In the fifth cluster we can see words clearly identified with the topics of Current and Prospective Economic Developments and Economic Projections, *economy, expected, outlook, growth* and *rate, percentage, projection*, respectively. The keywords of the sixth are *market, financial, conditions, credit, bank, loan, tightened/tightening, eased/easing, liquidity, investors*, strongly indicating the topic of Financial Markets, while the ones of the seventh cluster, *outlook, economic, economy, expansion, uncertainty, prospect, strengthened* indicating the topic of Economic Outlook. The keywords of the eight cluster, *spending, business, consumption, house/housing, sector, household, investment, sales, income, price, inventory, home, capital, contract, construction, mortgages, firms,* strongly indicate the topic of Key Sectors (which comprise Consumer Spending, Business Investments, and Housing Market). Lastly, the ninth cluster words, *policy, economic, rate, federal, funds, monetary, statement, future, stance, accommodative, forward, target, action, guidance* and *communication*, are strongly associated with the topics of Monetary Policy and Communication, the one subject of my analysis[37]. Identifying the correct labeling of such topics is not simple as it depends on the researcher's interpretability and knowledge of the words associated with each topic. Therefore, to mitigate such subjectivity, before starting to label the topics a careful analysis of the data texts (FOMC minutes) is crucial to understand which words the participants typically use during the discussion of a certain topic.

Another aspect worth noticing is that the probability characteristic of LDA is important, especially within this analysis, as it lets words to be associated with different topics. For example, the word *price* has a high probability of being associated with the topic of Inflation but also has a modest probability of being associated with the topic of Key sectors. Indeed, the level of prices is also relevant to the discussion of

[37] The cluster of interest englobes also the topic of Communication, which refers mostly to discussions on how to improve public understanding of the FOMC monetary policy decisions (current and future) and the ways in which those decisions depend on the members' assessments of the economic and financial conditions. Those discussions were not very frequent. The inclusion of this topic to the one of Monetary Policy do not much influence the results of my analysis that I show later.

business investment and consumer spending. Moreover, usually, when the members discuss a certain topic they frequently mention also others because they are related (for instance, economic conditions are strongly related to the level of inflation, unemployment, fiscal and monetary policy, etc).

iii. Proportion of Topics Over Time

As described in the Methods section, I used LDA also to calculate the "gammas", the document-topic probabilities. For each paragraph *d* the model computed the probability of belonging to each of the topic *t*. The sum of such probabilities for each paragraph is one and so such probabilities can be interpreted as the proportion devoted to a certain topic. Thus, a higher probability means a higher proportion devoted to a certain topic. I used the obtained gammas to identify the topic proportions per document over the years 2002-2013. Figure 3 shows the evolution over time, from 2002 to 2013, of the proportion of the FOMC minutes devoted to the topics previously identified with the LDA algorithm.

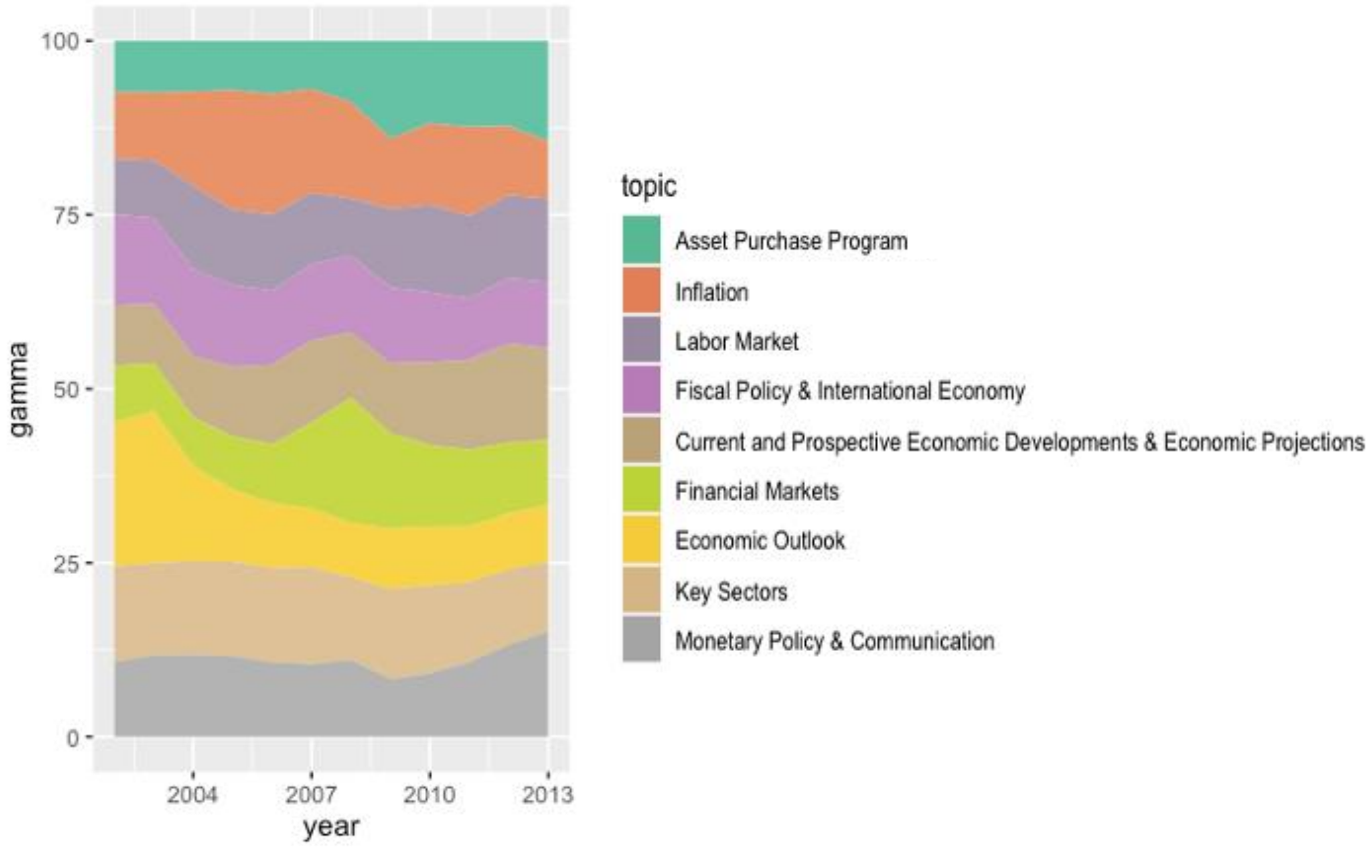

Figure 3: Proportion of FOMC Minutes Topics Over Time

As I expected, the figure doesn't show significant variations in the proportion of the topics discussed during the meetings throughout time. This is because the meetings (and so the minutes) follow a particular structure such that most of the topics are covered in each of them. Despite this aspect, we can still infer some relevant information. Inflation is amply discussed throughout the entire period, but especially between 2004-2006. This is coherent to the fact that during that period inflation pressures were a major concern for the FOMC. Also the topic of Key Sectors is significant throughout the entire period. Indeed, it encloses consumer spending, business investments, and housing market, information crucially important for the assessment of the economic conditions. As expected, due to the serious deterioration of financial markets, the proportion of the minutes devoted to such topic increased substantially during the years of the Global Financial Crisis. Following the same reasoning, also the topic of the Asset Purchase Program increased in the last part of the figure due to the large-scale program initiated in late 2008 (ended in 2014). In the next section I will analyze in detail the evolution of the topic regarding Monetary policy & Communication, the one subject of my analysis.

iv. A Focus on the Topic of Monetary Policy & Communication

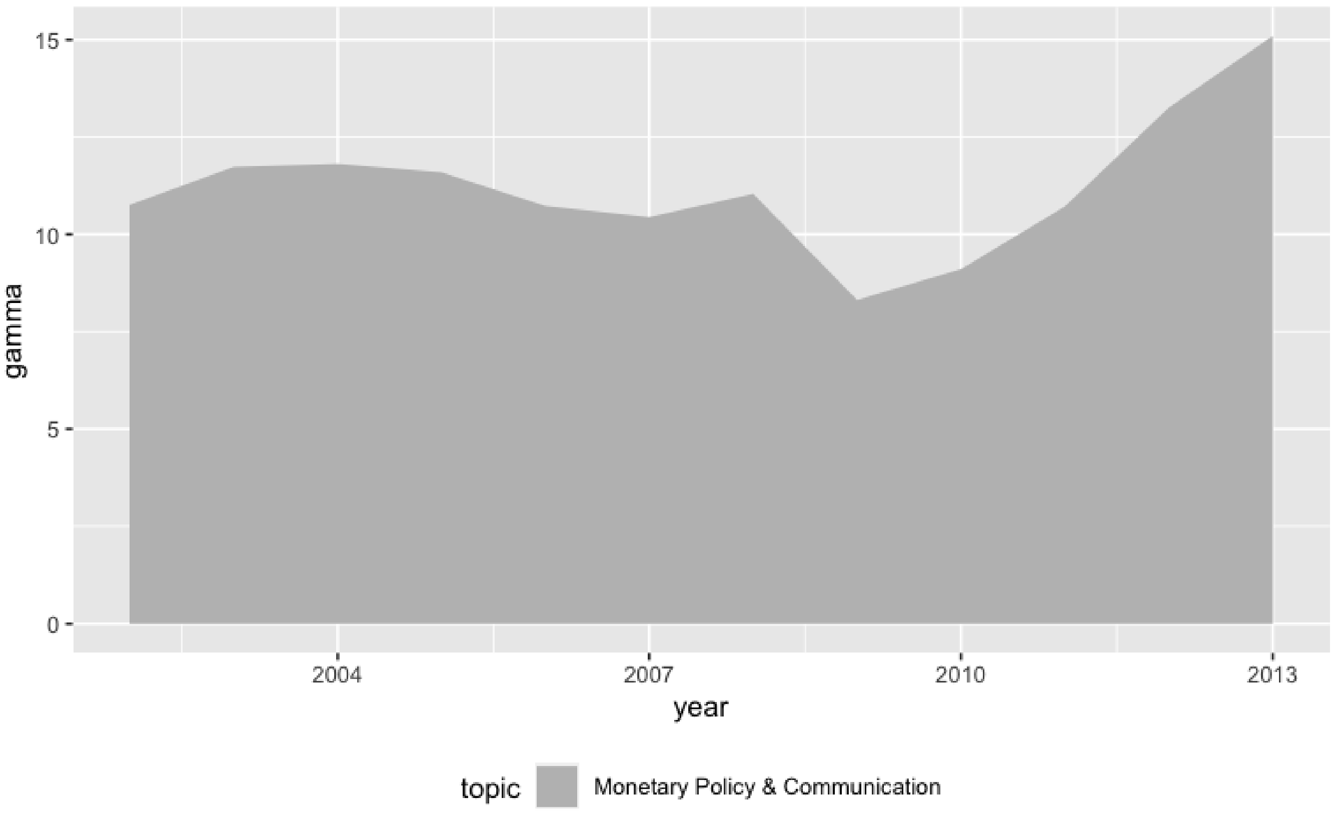

Figure 4: Proportion of FOMC Minutes Topic “Monetary Policy & Communication” Over Time

As illustrated in the introductory section, the goal of this thesis is to analyze if the increased policy transparency via greater use of forward guidance in the FOMC post-meeting statements has transformed the contents of their minutes between 2002-2013. In particular, I wanted to observe if this greater transparency in the post-meeting statements meant also a greater proportion of contents devoted to monetary policy (defined as references to current and future setting of the fed funds rate) in the FOMC minutes and if so, if this increase is more visible during and after the Global Financial Crisis, under Bernanke’s chairmanship, when the FOMC used a firmer and more explicit form of forward guidance language. Before conducting the analysis, I didn’t’ expect significant variation in the proportion of topics over time, including the topic devoted to monetary policy. However, I expected to see an increase in the proportion of the topic under analysis starting from the end of 2008, when the FOMC, under Bernanke’s guidance, moved one step further in its peregrination toward policy transparency with greater use of its forward guidance tool. In other words, I expected that with greater use of forward guidance the FOMC members also devoted more time to discuss the current and future stance of monetary policy. In figure 4 we can better view the evolution throughout 2002-2013 of the topic subject of analysis, that is, the one labeled before as Monetary Policy & Communication. As the figure shows, there is a slight increase in the proportion of the topic under analysis between 2002, the year still characterized by the "balance of risks" language and 2003, the starting year of the forward guidance language. Such proportion remains stable until mid-2005, when it starts to slightly decrease and then slight increase. Overall, there is not a significant variation in the proportion of the topic between 2002 and the end of 2007. Afterward, there is a decrease in the proportion until the end of 2008. If we consider the last part of the figure, we observe a visible increase in the topic proportion starting from the end of 2008-beginning of 2009 until the end of the sample. Overall the results support my initial hypothesis as there is not a significant change in the proportion of the topic over time (as figure 3 shows) but still there is a slight increase starting when the FOMC made greater use of its forward guidance tool (as figure 4 shows). In the following paragraph I will present and comment on the results in more detail.

The period under analysis is marked by the leadership of two Fed Chairmen, Alan Greenspan, at the helm at the beginning of the sample until January 2006 and Ben Bernanke, Chairman from February 2006 until the end of the sample. Alan Greenspan and Ben Bernanke share a lot of points of interest. Both were inflation fighters first and foremost and both were committed to maintain the independence of the Federal Reserve from short-term political pressures. They also had the same opinion regarding many aspects, one for example being the role of Central Banks in bursting asset bubbles. On the other hand, they were different in many other ways. One of these was that while Greenspan was very well known for his "Constructive Ambiguity" language, that is the use of long, wordy and vague statements and for his initial reluctance in moving toward greater policy transparency, Bernanke, from the early days as Governor in 2002, was a strong supporter of transparency and policy plans communications. From the very beginning of his chairmanship, Bernanke was strongly intentioned to make *"monetary policy as transparent and open as reasonably as possible"* (Bernanke, 2013). Although Greenspan set the stage, Bernanke steered the Federal Reserve in the direction of greater policy transparency, especially through greater use of the forward guidance language during and after the crisis. The sample period starts in 2002 when the FOMC didn't yet employ a direct form of forward language in its post-meeting statement. Until August 2003, the statement did not provide a direct indication of the next FOMC move but rather a language from which such information could be implicitly inferred, the "balance of risks". The latter indicated the members' perceived risks in the foreseeable future to the Fed's long-run goals of price stability and sustainable economic growth. In January 2001 the FOMC began a series of interest rate cuts to stabilize the economy after the burst of the dot.com bubble as well as other subsequent major events such as the terrorist attacks, U.S. invasion in Afghanistan and accounting scandals. In 2003 the FOMC was mostly worried about the lower growth of the rate of inflation and decreased the rate to the low level of 1 percent, where further reduction was very limited. By August of that year the balance of risks was not able to capture the Committee's fear of deflation and introduced the following qualitative time-dependent form of forward guidance: *"policy accommodation can be maintained for a considerable period"*. Although the definition of what constitutes a "considerable period" is flexible and the Committee, as always, remarked the dependence of future policy to the future conditions of the economy, such guidance was a direct and not vague indication of future maintenance of policy accommodation, far from the loose guidance of the "balance of risks". The *"considerable period"* language is today commonly recognized as the very first true form of FOMC forward guidance. Although such language started only at half year, we can see from the figure a slight increase in the proportion of the topic in the minutes between 2002 and 2003. The *"considerable period"* language was used for the subsequent meetings until the statement of January 2004 when *"can be patient in removing its policy accommodation"* was introduced, which was then substituted in May by *"accommodation can be removed at a pace that is likely to be measured"*, signaling that as the economic conditions were changing, the fed funds rate would have likely been raised soon. Indeed, at the meeting of June 2004, the FOMC increased the fed funds rate by 25 basis points to 1.25 percent and also repeated the same language of May, indicating that similar moderate increases would likely be followed in the other meetings of that year, as actually happened. At every subsequent meeting the Fed increased its rate by 25 basis points until the last meeting of Chairman Greenspan, January 2006, where the rate reached 4.5 percent. The FOMC kept the May 2004 forward guidance language until November 2005. It was then changed with looser statements, *"some further policy firming is likely to be needed to keep the risks to the attainment of both sustainable economic growth and price stability roughly in balance"* in December of the same year and in *"some further policy firming may yet be needed to address inflation risks,* […], *such firming will depend importantly on the evolution of the economic outlook as implied by incoming information"* in May 2006. The direct interest rate guidance that started in August 2003 ended in June 2006 when the FOMC released an indirect statement indicating that *"inflation risk remains"* and that the extent and timing of additional tightening were conditional on future information[38]. Such language continued until March 2007 when also the indication of likely direction was dropped, although the statement still indicated inflation as the primary concern. Such a drop in the language is also reflected in the minutes (although not from the beginning) as we can see a decrease in the topic proportion until the end of 2008-beginning of 2009. The Forward guidance language was resurrected only in December 2008 with *"likely to warrant exceptionally low levels of the federal funds rate for some time".* As the U.S. economy faced deep difficulties due to the subprime crisis, the Federal Reserve started a series of cuts in the fed funds rate until it reached the range of zero to 0.25 percent in December 2008[39]. That is, in case the economy needed more accommodation, as it happened, the Fed could not have decreased more the rate as it reached the zero lower bound. Given the severity of the recession and the problems in financial markets, the FOMC, as other major Central Banks' monetary policy bodies, significantly expanded its toolkit available for policy accommodation with

[38] At that date the fed funds rate reached the higher level of 5.25 percent.
[39] Just one year before the rate was 4.25 percent. The fed funds rate remained at the zero lower bound until December 2015.

unconventional programs, implementing the Large-Scale Asset Purchases (LSAPs) program, the open market purchase in large-scale of long-term U.S. Treasuries (along with securities issued by Government-sponsored agencies, such as MBS)[40]. Besides that, the Committee also started to make greater use of its forward guidance tool. During those years, Ben Bernanke was at the helm of the Federal Reserve. As highlighted before, he was a proponent of the use of forward guidance, especially in the zero lower bound condition. His view was buttressed by his 2004 empirical research (previously presented in the theory review) where he and the other authors argued that policy accommodation is not only provided by the current value of the benchmark rate but also by the public expectation of future path of such rate (Bernanke et al., 2004). Thus, the Committee, under the guidance of Bernanke, with the December 2008 statement, started to provide some further easing by assuring the markets that low level of the rate would likely be maintained for a long time (Bernanke, 2013). The reason why such guidance was not provided also in the initial stages of the crisis was that the rapidly changing conditions of financial markets and the economy made it difficult to announce in advance the likely FOMC future actions[41]. The December 2008 statement was slightly modified in March 2009, when *"for some time"* was replaced with *"for an extended period"*. Such language was maintained until June 2011, after which the Committee significantly changed the statement with a less qualitative and more explicit language introducing a specific date. In August 2011 the statement indicated that *"likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013"*. This type of language is also referred to as date-based forward guidance. Such date was then pushed out twice, *"at least through late 2014"* (in January 2012) and *"at least through mid-2015"* (in September 2012). Afterward, the language underwent a second important change as the future path of the policy rate was tied to the Fed's economic objectives. As Bernanke (2013) stated, the FOMC believed the date-based statements to be limited and opted for a language that explained how future policy would be affected by future economic conditions, the so-called "state-contingent language". The December 2012 statement indicated that accommodation would be maintained at least as long the *"the unemployment rate remains above 6-1/2 percent, inflation between one and two years ahead is projected to be no more than a half percentage point above the Committee's 2 percent longer-run goal, and longer-term inflation expectations continue to be well anchored"*. Also the statement of December 2013, the second to last meeting of Chairman Bernanke, in a similar way reported a state-contingent language (Table 1 in Appendix 1: Additional Tables reports a detailed record of FOMC use of direct forward guidance between 2003-2013). That is, during and after the Global Financial Crisis, the FOMC under Bernanke used a firmer and more explicit form of forward guidance language. If we consider the last part of figure 4, we can infer that this increase in policy transparency via greater use of forward guidance in the post-meeting statement is also reflected in the minutes, as we observe an increase in the proportion of topic devoted to monetary policy starting from the end of 2008-beginning of 2009 until the end of the sample.

[40] More on the program at the following: https://www.newyorkfed.org/markets/programs-archive/large-scale-asset-purchases
[41] As it can be inferred from the post-meeting statements of those months.

Another aspect worth discussing regards the Large Asset Purchase Program (also known as Quantitative Easing) mentioned before. As figure 3 shows, the Asset Purchase Program topic proportion increased starting from late 2008. As the economic recession and the problems in financial markets became severe, besides lowering to zero the fed funds rate and providing guidance that such rate would have been kept low for a long period, starting from late 2008 (until October 2014) the FOMC provided additional stimulus via the unconventional tool of Large-Scale Asset Purchase Program (LSAPs). As Bernanke (2013) illustrated, both forward guidance and LSAPs affect long-term interest rates but in different ways: forward guidance affects long-term rates by shaping the expectations of the future short-term rates whereas LSAPs by influencing the term premium. However, the LSAPs can also be seen as a commitment of a Central Bank to maintain the target rate to zero for a certain period[42] and thus it may also affect short-term market expectations. Thus, in a certain way, also the topic devoted to the Asset Purchase Program, captures, although in a small part, also the topic of monetary policy (relative to future interest rate setting). Therefore, the increase of the proportion of such topic visible in the last part of figure 4 can be partly considered as a contribution to greater monetary policy transparency in the minutes.

[42] This is because, by implementing a Large-Scale Asset Purchase Program a Central Bank commits itself to keep reserves well above those needed to maintain the benchmark rate to zero (Bernanke and Reinhart, 2004)

V. Discussion

The goal of this thesis is to analyze if the increased policy transparency via greater use of forward guidance in the FOMC post-meeting statements has transformed the contents of their minutes over 2002-2013. In particular, I wanted to observe if this greater transparency in the post-meeting statements means also a greater proportion of contents devoted to monetary policy (defined as references to current and future setting of the fed funds rate) in the FOMC minutes and if this increase is more visible during and after the Global Financial Crisis, under Bernanke's chairmanship, when the FOMC used a firmer and more explicit form of forward guidance language. The findings support my initial hypothesis. Overall there isn't a significant variation in the proportion of topics over time as the minutes are characterized by consistent contents over time. The results also show an increase in the proportion of the topic under analysis starting with the firmer forward guidance during and post-crisis.

The same research question of this thesis can also be addressed to the transcripts[43]. If the LDA process is applied in an analogous way the results would probably be the same, as the minutes can be considered as a summary of the transcripts. It would also be interesting to apply the same research question to Greenspan's and Bernanke's speeches to observe the difference in the proportion of contents devoted to monetary policy transparency of the two Chairmen. As illustrated before, I chose to investigate the period of 2002-2013 as it was marked by greater policy transparency via the use forward guidance tool as well as by the leadership of two Fed Chairmen with a different perspective on policy transparency. However, it would also be interesting to apply the same research question to a different period, for example to the more recent years, and to different chairmanships. To address my research question, I focused on the proportion over time of the particular topic devoted to monetary policy (defined as references to current and future setting of the fed funds rate), trying to link the results to FOMC use of forward guidance over time. In the same way, through the use of topic modeling techniques such as LDA, future researchers can investigate in more detail the evolution over time of other contents, linking the results, for instance, to the FOMC use of different monetary policy rules over time.

A second but important goal of this thesis is also being able to correctly identify, through the use of LDA, the topic discussed by the FOMC during their meetings. Word clustering is the most crucial step of LDA process and of my analysis. Without an appropriate word clustering each document could have been assigned to the topics with improper proportions, thus invalidating figure 3 and so the entire analysis. As we can see from table 1, the top words from each topic cluster are mostly distinct, enabling us to recognize without doubts the relevant topic labels for each of them. Although its application is not simple, the LDA is a very powerful tool as, if properly applied, is capable of accurately clustering words belonging to the same topic without the use of any user pre-defined clusters of words. This is an important advantage with respect to other text mining techniques and so I expect to see in the future a greater application of it in

[43] Also the transcripts can be analyzed via the use of LDA, as the work of Edison and Carcel (2020) demonstrates.

different fields and hopefully in economics. More in general, as seen in the literature review, considering that text mining techniques are capable of addressing a wide variety of research questions within central banking documents, I expect such techniques to be largely employed in the future.

Although topic modeling studies such as mine overcome the subjective scoring of the manual approach, still their process and so the output is influenced by the researcher's judgment and knowledge on the subject. Thus, for instance, studies that make use of the same topic modeling technique to the same or similar documents still may produce different results. As I mentioned before, the technique I applied, LDA, requires the number of topics to be defined ex-ante by the user. Furthermore, LDA does not provide topic labels, which must be instead identified by the user in accordance with the subject being studied. Both topic identification and topic labeling strongly depend on the research's judgment and knowledge on the subject. These aspects introduce some subjectivity in the analysis. For instance, although Jagadeesh and Wu (2017) and Edison and Carcel (2020) performed a similar analysis to mine the respective results are difficult to compare because of the different judgments used in the choice of the number of topics and on the topic labeling.

VI. Conclusion

In this thesis I apply LDA, a popular text mining technique, to the FOMC minutes from 2002 to 2013 to quantify the content discussed during their meetings. I implemented such algorithm intending to split the corpus into nine distinguishing topics and identify the topic proportions per document over the years 2002-2013. I wanted to observe if this greater transparency in the post-meeting statement meant also greater proportion of contents devoted to monetary policy in the FOMC minutes and if so, if this increase is more visible during and after the Global Financial Crisis when the FOMC used a firmer and more explicit form of forward guidance language. The results support my initial hypothesis as overall there is not a significant variation in the proportion of topics over time but still, we can observe an increase in the proportion of the topic under analysis starting with the use of firmer forward guidance during and post-crisis.

In the past thirty years we have seen a revolution in the Federal Reserve's communication. Before such openness the Fed provided little information to the public about its policy decisions, fearing the market would overreact or that communication about the future course of fed funds rate would have reduced FOMC discretion to adjust it according to new economic conditions. That is, there was a widespread belief among central bankers that openness could damage the effectiveness of monetary policy. During those years, indeed, the market was accustomed to look how the FOMC systematically responded in the past to certain economic conditions to infer its next policy moves and to a monetary policy that relied on changes to the fed funds rate. The first publication of the post-meeting statement in 1994 marked a turning point. That revolution in thinking was motivated by different reasons but mostly by a considerable amount of evidence of enhancement of the effectiveness of monetary policy via the management of the future fed funds rate expectations. Although with initial reluctance, starting in 1999, the FOMC provided some form of future rate guidance in terms of "tilt" or "bias" and then "balance of risks". Afterward, in August 2003, fearing a possible deflation, the FOMC issued what is today recognized as its first forward guidance statement. As Bernanke's successor Janett Yellen (2013) commented, *"The FOMC had journeyed from "never explain" to a point where sometimes the explanation is the policy".* The use of forward guidance accelerated during and after the Global Financial Crisis when the FOMC couldn't rely anymore neither on its past behavior for communicating its future moves nor on its conventional tools to stimulate the economy. Bernanke's successors, Janet Yellen and Jerome Powell also embraced the benefits of clear communication with the public, for instance, supplementing post-meeting statements with more frequent press conferences. To conclude, the Federal Reserve's communication policy evolved over time according to the evolving needs of the Committee. As the coronavirus demonstrates, surely new challenges lie ahead of the Fed. Those will inevitably lead to an evolution in the Federal Reserve's conduct of monetary policy as well as communication practices to preserve its maximum employment and stable price goals. As Bernanke (2015) said, although economics is very good at illustrating ex-post to policy-makers why their choices turned out to be wrong, it is not so good at forecasting the effects of future choices.

VII. References

Acosta, M. (2015). FOMC responses to calls for transparency. Finance and Economics Discussion Series. 2015-60. Board of Governors of the Federal Reserve System, Washington D.C. DOI: http://dx.doi.org/10.17016/FEDS.2015.060.

Bernanke, B. (2004). *Central bank talk and monetary policy.* Remarks at the Japan Society Corporate Luncheon, New York, New York. Retrieved from https://www.federalreserve.gov/boarddocs/speeches/2004/200410072/default.htm.

Bernanke, B. (2013). *Communication and monetary policy*. Speech at the National Economists Club Annual Dinner Herbert Stein Memorial Lecture, Washington, D.C. Retrieved from https://www.federalreserve.gov/newsevents/speech/bernanke20131119a.htm.

Bernanke, B. (2015). *The courage to act: A memoir of a crisis and its aftermath*. New York, NY: W. W. Norton & Company.

Bernanke B., & Reinhart, V. (2004). Conducting monetary policy at very low short-term interest rates. *American Economic Review.* vol. 94(2), pages 85-90, May.

Bernanke, B., Reinhart, V., & Sack, B. (2004). Monetary policy alternatives at the zero bound: An empirical assessment. *Brookings Papers on Economic Activity*, 2004(2), 1-100.

Bholat, D., Hansen S., Santos, P., & Schonhardt-Bailey, C. (2015). *Text mining for central banks.* Centre for Central Banking Studies Handbook (33), 1–19.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3: 993–1022.

Blinder, A. S. (1998). *Central banking in theory and practice*. Cambridge: MIT Press.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., & Jansen, D. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature, 46*(4), 910-945.

Bligh, M. C., & Hess, G. (2013). *Deconstructing Alan: A quantitative assessment of the qualitative aspects of Chairman Greenspan's communication.* in central bank communication, decision making, and governance: issues, challenges, and case studies, P. L. Siklos and J.-E. Sturm, eds. Cambridge, MA: MIT Press, 2013, 123–147.

Boukus, E., & Rosenberg, J. V. (2006). The information content of fomc minutes. Mimeo, Federal Reserve Bank of New York.

Contessi, S. & Li, L. (2013). Forward guidance 101b: A roadmap of the international experience. Economic Synopses. Federal Reserve Bank of St. Louis, Vol 28.

Danker, D. J., & Luecke, M. M. (2005). Background on FOMC meeting minutes. Federal Reserve Bulletin 175-179.

Edison, H. & Carcel, H. (2020). Text data analysis using Latent Dirichlet Allocation: an application to FOMC transcripts. *Applied Economics Letters.* DOI: 10.1080/13504851.2020.1730748.

Greider, W. (1987). *Secrets of the temple: How the Federal Reserve runs the country*. New York: Simon & Schuster.

Greenspan, A. (2007). *Alan Greenspan: The age of turbulence.* New York: The Penguin Press.

Gürkaynak, R. S., Sack, B. P., & Swanson, E. T. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1): 55-93.

Hansen, S., McMahon, M., & Prat, A. (2017). Transparency and deliberation within the FOMC: A computational linguistics approach. *The Quarterly Journal of Economics*, 133 (2): 801–870. doi:10.1093/qje/qjx045.

Hendry, S., & Madeley, A. (2010). Text mining and the information content of Bank of Canada communications. Bank of Canada Working Paper.

Jegadeesh, N., & Wu, D. A. (2017). Deciphering Fedspeak: The information content of FOMC meetings.

Kahveci, E., & Odabaş, A. (2016). Central banks' communication strategy and content analysis of monetary policy statements: The case of Fed, ECB and CBRT. *Procedia - Social and Behavioral Sciences*, vol. 235, 2016, pp. 618–629.

Lucca, D. O., & Trebbi, F. (2009). Measuring central bank communication: An automated approach with application to FOMC statements. Cambridge, MA: National Bureau of Economic Research.

Mazis, P. & Tsekrekos, A. (2017). Latent semantic analysis of the FOMC statements. *Review of Accounting and Finance,* vol. 16, no. 2, 2017, pp. 179-217.

Poole, W. (2005). After Greenspan: Whither Fed policy? Presentation to the Western Economic Association International Conference (WEAI). San Francisco, California.

Schwarz, C. (2018). ldagibbs: A command for topic modelling in Stata using Latent Dirichlet Allocation. *The Stata Journal* 18 (1): 101–117. doi:10.1177/ 1536867X1801800107.

Silge, J. & Robinson, D. (2020). *Text mining with R: a tidy approach*. Sebastopol, CA: OReilly Media.

Vallès, D. W. & Schonhardt-Bailey, C. (2015). Forward guidance as a new idea: Changes in the MPC Discourse of the MPC under King and Carney. Presented at the Political Leadership and Economic Crisis Symposium, Yale University.

Woodford, M. (2001). *Monetary policy in the information economy.* Cambridge, MA: National Bureau of Economic Research.

Woodford, M. (2005). Central bank communication and policy effectiveness. In the Greenspan Era: Lessons for the Future. Kansas City: Federal Reserve Bank of Kansas City, 399- 474.

Yellen, J. (2013). Communication in monetary policy. Speech at the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, D.C. Retrieved from https://www.federalreserve.gov/newsevents/speech/yellen20130404a.htm.

VIII. Appendix

i. Appendix 1: Additional Tables

Chairman	Meeting Date	Fed funds rate (in percentage)	Forward Guidance Language
	Pre-crisis		
A. Greenspan	Aug-03	1	"accommodation can be maintained for a considerable period"
A. Greenspan	Jan-04	1	"can be patient in removing its policy accommodation"
A. Greenspan	May-04	1	"accommodation can be removed at a pace that is likely to be measured"
A. Greenspan	Dec-05	4.25	"some further policy firming is "likely to be needed to keep the risks to the attainment of both sustainable economic growth and price stability roughly in balance."
B. Bernanke	May-06	5	"some further policy firming may yet be needed to address inflation risks, [..], such firming will depend importantly on the evolution of the economic outlook as implied by incoming information"
	During and Post-crisis		
B. Bernanke	Dec-08	0-0.25	"likely to warrant exceptionally low levels of the federal funds rate for some time"
B. Bernanke	Mar-09	0-0.25	"likely to warrant exceptionally low levels of the federal funds rate for an extended period"
B. Bernanke	Aug-11	0-0.25	"likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013"
B. Bernanke	Jan-12	0-0.25	"are likelt to be warrant at least through late 2014"
B. Bernanke	Sep-12	0-0.25	"are likely to be warranted at least through mid-2015"
B. Bernanke	Dec-12	0-0.25	"exceptionally low range for the federal funds rate will be appropriate at least as long as the unemployment rate remains above 6-1/2 percent, inflation between one and two years ahead is projected to be no more than a half percentage point above the Committee's 2 percent longer-run goal, and longer-term inflation expectations continue to be well anchored"
B. Bernanke	Dec-13	0-0.25	"likely will be appropriate to maintain the current target range for the federal funds rate well past the time that the unemployment rate declines below 6-1/2 percent, especially if projected inflation continues to run below the Committee's 2 percent longer-run goal"

Table 2: Direct use of forward guidance language by the FOMC as communicated in its post-meeting statements between 2003-2013

This table was constructed using the forward guidance language contained in the post-meeting statements that you can find at the following links:
https://www.federalreserve.gov/monetarypolicy/fomc_historical_year.htm and
https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

Chairman	Date	Milestone
Alan Greenspan	February 1994	Publication of the FOMC transcripts Release of the first post-FOMC meeting statement but only in case of policy action. It informed about the decisions on the intended fed funds rate
	May 1999	Release of post-FOMC meeting statement of decisions on the intended fed funds rate, regardless of the change, including Committee expected future direction of the policy (**Policy Tilt**)
	February 2000	Replacement of the Policy Tilt with the **Balance of Risk Assessment** to the dual mandated objectives
	March 2002	Addition of roll-call vote to the post- FOMC meeting statement
	August 2003	Addition of the first **forward guidance statement** in the post-FOMC meeting statements
	December 2004	Acceleration in the release of the minutes from six-week lag to three
Ben Bernanke	October 2007	Addition of the **Summary of Economic Projections** (SEP) to the minutes (attached as an addendum to the minutes)
	November 2007	Release of the annual economic projections quarterly instead of semiannually
	September 2010	Addition of non-voting FOMC members' monetary policy view in the minutes
	April 2011	First post-FOMC meeting press conference
	January 2012	Addition of Dot Plot (members' projections of future fed funds rate) in the SEP Release of the annual Statement of Longer-run Goals and Policy Strategy, which included an explicit 2-percent inflation target

Table 3: Federal Reserve major transparency milestones during Greenspan's and Bernanke's chairmanships

This table was constructed using the information contained in the minutes that you can find at the following links: https://www.federalreserve.gov/monetarypolicy/fomc_historical_year.htm and https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm.

also	eight	governor	many	mrs.	one	seventh
although	eleven	governors	march	ms	participants	six
april	eleventh	however	meeting	ms.	period	ten
augst	february	in may	meetings	much	prospective	tenth
chairman	fifth	intermeeting	member	nine	quarter	third
committee	first	jackson	members	ninth	recent	three
committees	five	january	month	nonethless	second	twelve
current	four	july	mr	noted	secretary	two
december	fourth	june	mr.	november	september	word
discussion	generally	last	mrs	october	seven	year

Table 4: List of additional stopwords (“mystopwords”) for the data cleaning process

ii. Appendix 2: Example of a Text Cleaning Process

Original text, from the 30th June 2003 FOMC Minutes

```
“In the Committee's discussion of policy for the intermeeting period, all of the members indic
ated that they could support an upward adjustment in the target for the federal funds rate fro
m a     level of 1 percent to 1-1/4 percent. Recent developments, notably the persistence of s
olid gains in output and employment along with indications of some increase in inflation, were
seen as           warranting a first step in the process of removing policy accommodation. The
timing and pace of    further policy moves would depend, of course, on the members' reading of
the incoming economic information and their interpretation of its implications for economic ac
tivity and inflation.”
```

Text cleaned from special characters, punctuation, numbers, excess whitespace and uppercase converted into lowercase

```
“in the committees discussion of policy for the intermeeting period all of the members indicat
ed that they could support an upward adjustment in the target for the federal funds rate from a
level of percent to percent recent developments notably the persistence of solid gains in outp
ut and employment along with indications of some increase in inflation were seen as warranting
a    first step in the process of removing policy accommodation the timing and pace of further
policy moves would depend of course on the members reading of the incoming economic informatio
n and their interpretation of its implications for economic activity and inflation”
```

Text cleaned from the most common English “stopwords” and from “mystopwords” list

```
“policy indicated support  upward adjustment target funds rate  level  percent  perce
nt  recent developments notably  persistence  solid gains  output  employment along
indications   increase  inflation  seen  warranting  first step   process  removing
policy    accommodation   timing   pace    policy moves   depend   course     reading
incoming economic information    interpretation    implications   economic activity
inflation”
```

Stemming of the words

```
"polici  indic  support  upward  adjust  target  fund  rate  level  percent  percent  recent
develop notabl persist solid gain output employ along indic increas inflat seen warra
nt first step process remov polici accommod time pace polici move depend cours read i
ncom    econom inform interpret implic econom activ inflat"
```

In order to perform the above passages, I used the *tm package*[44] and the following functions in Rstudio:

- `tolower()`: Convert uppercase characters to lowercase
- `removePunctuation()`: Remove the punctuation
- `removeNumbers()`: Remove numbers
- `stripWhitespace()`: Remove excess whitespace
- `removeWords()`: Remove “stopwords” and “mystopwords”
- `StemDocument()`: Stem the words

[44] This package is preferred for a single text document.

iii. Appendix 3: Coding Used for Analysis

```
#Packages needed to perform the analysis
library(tm) #needed for different functions
library(wordcloud) #needed to create the wordcloud
library(topicmodels) #interface to the C code for Latent Dirichlet Allocation (LDA)
library(tidytext) #allows the use of tidy principles for text mining
library(ggplot2) #needed to create graphics
library(dplyr) #needed for different functions (such as mutate, anti_join, etc)
library(tidyr) #needed for different functions (such as spread)
library(lubridate) #allows certain data manipulation
library(tidyverse) #contains needed data anlysis packages
library(SnowballC) #needed for stemming

#The hereby codes refer to the part after the upload of the documents and creation of the corpus ("text").

#initial data cleaning
text1=gsub(pattern = "\\W", replace=" ", text) #to remove the punctuation
text1=gsub(pattern = "\\d", replace=" ", text1) #to remove the numbers
text1=tolower(text1) #to convert uppercase to lowercase
text1=gsub(pattern="\\b[A-z]\\b{1}", replace=" ", text1) #to remove single-digit words (like "s")
text1=stripWhitespace(text1) #to remove the excess whitespace

#In order to tokenize the words with the unnest_tokens function I transform the corpus into a data frame.
text_df = tibble(line = 1:818, text = text1) #818 refers to number of documents
#tokenization
minutes_tidy = unnest_tokens(text_df, word, text) #one token per each row in one column. The other column indicates
                                                   the document where the word belongs.

#The most common words, such as "the", "of" and "in" are not meaningful. Thus, they must be removed before
proceeding further. I clean the documents from stopwords together with a customized list of words ("mystopwords")
with meaningless words commonly seen in the Federal Reserve language, such as “committee” and “participants”.
mystopwords=tibble(word=c("also", "although", "april", "augst", "chairman", "committee", "committees", "december",
"eight", "eleven", "eleventh","february", "fifth", "first", "five", "four", "fourth", "generally", "governors",
"governor", "in may", "jackson", "january", "july", "june", "however", "last", "many", "march", "meeting",
"meetings", "much", "member", "members", "mr", "mr.", "mrs", "-", "mrs.", "ms", "ms.", "nine", "ninth",
"nonethless", "november", "october", "one", "participants", "quarter", "second", "secretary", "september", "seven",
"seventh", "six", "ten", "tenth", "third", "three", "twelve", "two", "year", "intermeeting", "discussion", "noted",
"recent", "prospective", "current", "period", "month"))

minutes_cleaned=minutes_tidy %>%
  anti_join(stop_words) # it cut the number of words almost by half (from 185,249 to 92,950)
minutes_cleaned = minutes_cleaned %>%
  anti_join(mystopwords) # (86,172)

#Let's apply stemming, removal of the suffix or prefixed of words to maintain just their root forms. After stemming
the algorithm can recognize such words as the same word. I then re-apply “stopword” and “mystopwords” a second time
to remove unwanted words that may have been surfaced after stemming.
minutes_stemmed= minutes_cleaned %>%
  mutate(word = wordStem(word))
minutes_cleaned2 = minutes_stemmed[!str_detect(minutes_stemmed$word, '[[:digit:]]'),]
minutes_cleaned2= minutes_cleaned2 %>%
  anti_join(stop_words) %>%
  anti_join(mystopwords) # the number of words slightly decreased (84,404)

#Visualization of the most frequent words after cleaning, tokenizing and stemming (a TermDocumentMatrix is needed)
minutes_tdm= minutes_cleaned2 %>%
  count(word, title) %>%
  cast_tdm(word, title, n) #creation of a Term Document Matrix
minutes_tdm #2,437 unique words (or "terms")

minutes_tdm_matrix=as.matrix(minutes_tdm) #to visualize the matrix

print(findFreqTerms(minutes_tdm, lowfreq = 500)) # list of the most frequent terms
v1=sort(rowSums(minutes_tdm_matrix), decreasing = TRUE) #needed for the bar plot
d1=data.frame(word=names(v1), frequency=v1) #needed for the bar plot
barplot(d1[1:30,]$frequency, las=2, names.arg=d1[1:30, ]$word, col="light green", main="Most Frequent Words",
ylab="Word Frequencies") #bar plot

set.seed(1234)
wordcloud(words = d1$word, freq = d1$freq, min.freq = 1,
          max.words=100, random.order=FALSE, rot.per=0.15,
          colors=brewer.pal(8, "Dark2")) #wordcloud


#In order to apply the LDA i need to create a Document Term Matrix: each row represents a document d of the corpus
D, each column a unique term (that is, in the columns there are all the unique words present in all the documents)
and where the entries represent the frequency of the terms that occur per each document. The weighting is Term
Frequency.
minutes_dtm=minutes_cleaned2%>%
```

```
  count(title, word) %>%
  cast_dtm(title, word, n)

#Application of the LDA. The Document Term Matrix is used as object of the LDA. The method used is the Gibbs
sampling. The number of topics chosen is nine.
minutes_lda9 = LDA(minutes_dtm, k = 9, method = "Gibbs", control = list(seed = 1, burnin = 1000, thin = 100, iter =
1000)) #lda process
topics_beta9 = tidy(minutes_lda9, matrix = "beta") #to find the betas
top_terms9 = topics_beta9 %>% group_by(topic) %>% top_n(20, beta) %>% ungroup() %>% arrange(topic, -beta) #to find
                                                                                  the top 20 words for each topic
topics_gamma9 = tidy(minutes_lda9, matrix = "gamma") #to find the gammas

#assign a topic labeling to the cluster of words
topics_gamma9$topic = factor(topics_gamma9$topic, levels = 1:9, labels = c("Asset Purchase
Programme","Inflation","Labor Market", "Fiscal Policy and International Economy", "Current and Prospective Economic
Developments and Projections", "Financial Markets", "Economic Outlook", "Key Sectors", "Monetary Policy &
Communication"))


#Visualization of the proportions of topics over time using the gammas.
new_topics_gamma9= new_topics_gamma9 %>%
  group_by (date, topic) %>%
  summarise(gamma=sum(gamma)) #group by year and sum of the topic gammas with the same year

new_topics_gamma9_wide=spread(new_topics_gamma9, date, gamma)
new_topics_gamma9_wide #change it to a table: in the rows there are topics and in the columns the years. The entires
represent the sum of the gammas

#convert to percentage to take into account the difference in the number of paragraphs per year (normalization)
new_topics_gamma9_wide_new = mutate(new_topics_gamma9_wide,
                                    `2002` = 100 * new_topics_gamma9_wide$`2002` /
sum(new_topics_gamma9_wide$`2002`),
                                    `2003` = 100 * new_topics_gamma9_wide$`2003` /
sum(new_topics_gamma9_wide$`2003`),
                                    `2004` = 100 * new_topics_gamma9_wide$`2004` /
sum(new_topics_gamma9_wide$`2004`),
                                    `2005` = 100 * new_topics_gamma9_wide$`2005` /
sum(new_topics_gamma9_wide$`2005`),
                                    `2006` = 100 * new_topics_gamma9_wide$`2006` /
sum(new_topics_gamma9_wide$`2006`),
                                    `2007` = 100 * new_topics_gamma9_wide$`2007` /
sum(new_topics_gamma9_wide$`2007`),
                                    `2008` = 100 * new_topics_gamma9_wide$`2008` /
sum(new_topics_gamma9_wide$`2008`),
                                    `2009` = 100 * new_topics_gamma9_wide$`2009` /
sum(new_topics_gamma9_wide$`2009`),
                                    `2010` = 100 * new_topics_gamma9_wide$`2010` /
sum(new_topics_gamma9_wide$`2010`),
                                    `2011` = 100 * new_topics_gamma9_wide$`2011` /
sum(new_topics_gamma9_wide$`2011`),
                                    `2012` = 100 * new_topics_gamma9_wide$`2012` /
sum(new_topics_gamma9_wide$`2012`),
                                    `2013` = 100 * new_topics_gamma9_wide$`2013` /
sum(new_topics_gamma9_wide$`2013`)
)

data_long=gather(new_topics_gamma9_wide_new, year, gamma, `2002`:`2013`, factor_key=TRUE)
data_long=data_long %>%
  mutate(year = as.integer(as.character(year))) #table with topic, year and gamma

#plot of the proportions of topics over time
nb.cols=9 #as the number of topics
mycolors=colorRampPalette(brewer.pal(8, "Set2"))(nb.cols) #select the colors
ggplot(data_long, aes(x=year, y=gamma, fill=topic)) + geom_area() + scale_fill_manual(values = mycolors) +
theme(axis.text.x = element_text()) #plot of distribution of topics over time

# Extract only "Monetary Policy & Communication" topic
monetary_policy = new_topics_gamma9_wide_new[9,]
monetary_policy_long = gather(monetary_policy, year, gamma, `2002`:`2013`, factor_key=TRUE)
monetary_policy_long = monetary_policy_long %>%
  mutate(year = as.integer(as.character(year)))
monetary_policy_long

#plot of the proportion of topic over time
ggplot(monetary_policy_long, aes(x=year, y=gamma, fill=topic)) + geom_area() + scale_fill_manual(values = "grey") +
theme(legend.position="bottom")
```

For further information regarding the codes used and to receive the text data subject of analysis please contact me at daisy.lagana@studenti.luiss.it.

Acknowledgements

I would like to thank LUISS University for the good education, the incredible opportunities and the positive learning environment it offered me throughout these five years.

I would like to thank all my professors and teaching assistants (Bachelor’s and Master’s) for the quality of their courses and for the time they dedicated to me during office hours. I enjoyed every project and exam.

I would like to thank my supervisor prof. Guido Traficante for the good advice and the availability with which he supported me during the drafting of this thesis.

Page intentionally left blank

LUISS

Department of Business and Management

Chair of International Financial Economics

Federal Reserve’s Communication: a Latent Dirichlet Allocation Analysis with Application to FOMC Minutes

(Summary)

Guido Traficante

SUPERVISOR

Mauro Visaggio

CO-SUPERVISOR

Daisy Martina Laganà
713391

CANDIDATE

Academic Year 2019/2020

Table of Contents

Abstract

Over the last thirty years Central Banks have become remarkably more transparent about their policy-making, drifting apart from the tradition of "never explain, never excuse" and using communication as a key tool of monetary policy. This revolution in thinking was motivated by different reasons but mostly by a considerable amount of evidence of enhancement of the effectiveness of monetary policy. In the aftermath of the Global Financial Crisis much attention has been posed to Central Banks' forward-looking communication related to the monetary policy stance, what is known as forward guidance. In this thesis I apply Latent Dirichlet Allocation (LDA), a text mining technique, to the FOMC minutes from 2002 to 2013 to analyze if forward guidance has transformed the contents of FOMC minutes over time.

I. Introduction

Over the last thirty years Central Banks have become remarkably more transparent about their policy-making, drifting apart from the tradition of secrecy and also started to rely more heavily on communication for monetary policy-making. This revolution in thinking was motivated by different reasons but mostly by a considerable amount of empirical evidence of enhancement of the effectiveness of monetary policy driven by communication. As a result, Central Bank communication has gained increasing attention by market participants, media, academics, and the public in general and nowadays can boast a vast literature. Within such, there is also a small but growing branch that focuses on the application of text mining techniques, computational and statistical tools able to quantify text data, on Central Bank's documents for content analysis. This thesis contributes to such literature as I make use of a popular text mining technique, called LDA, to infer the proportion of FOMC minutes devoted to the topic of monetary policy over time. To the best of my knowledge, the LDA technique has been applied to the FOMC minutes only in a recent paper by Jegadeesh and Wu (2017), although to a different period and within a different research question.

The Federal Open Market Committee (FOMC), the Federal Reserve's body in charge of monetary policy decisions, marked significant milestones toward more transparent policy-making over time (see table 3 of Appendix 1: Additional Tables). Nowadays, the FOMC communicates with the public about its latest and future policy decisions with different tools although the transcripts, minutes, and post-meeting statements contain the most important contents for the market and the academic world. The transcripts report the members' full conversations during their meetings, the minutes are a summary of the transcripts whereas the post-meeting statements provide a short summary of the FOMC's latest policy decisions and reasons as well as an indication of the future path of monetary policy (although not always). Before such openness, the Fed provided little information to the public about its policy decisions. Indeed, only in 1994 the Fed started to issue its post-meeting statements, which, however, were much shorter and looser in content with respect to today's ones. Perhaps the most significant change in communication policy was in August 2003 when the FOMC started to provide some indications in their post-meeting statements of the expected or intended future path of policy rates, what it is known today as forward guidance. In reality, the FOMC already provided looser policy guidance since May 1999, in the form of "tilt" or "bias" and then in the form of "balance of risks". Such statements did not provide a direct indication of the FOMC next moves but rather a language from which such information could be implicitly inferred. In August 2003, worried about a possible emergence of deflation, after lowering the fed funds rate almost to the zero-lower bound, the FOMC tried to manage expectations with a firmer and time-dependent form of forward guidance. Such direct guidance ended in June 2006 and then was resurrected once the FOMC had clearer ideas regarding the economic outlook. In December 2008, at the peak of the Global Financial Crisis, the FOMC used a purely qualitative language to indicate that policy accommodation would have continued for *"some time"*. Such forward guidance language evolved over time, according to the needs of the Committee, from pure qualitative to the much more specific and direct date-based forward guidance of August 2011 - October 2012 and state-contingent forward guidance of December 2012 - December 2013 (see Table 2 in Appendix 1: Additional Tables). It is not surprising that the Chairman of the crisis and post-crisis years, Ben Bernanke, was a long-term proponent of policy transparency and of the use of forward guidance for shaping market expectations, especially in the zero lower bound condition, when future rate cuts are no longer possible.

For this thesis I applied Latent Dirichlet Allocation (LDA) to the FOMC minutes from 2002 to 2013 in order to analyze if increased policy transparency via forward guidance and greater use of the latter in the FOMC post-meeting statements have transformed the contents of their minutes over 2002-2013. In particular, I wanted to observe if this greater transparency in the post-meeting statement meant also a greater proportion of contents devoted to monetary policy (defined as references to current and future setting of the fed funds rate) in the FOMC minutes and if this increase is more visible during and after the Global Financial Crisis, under Bernanke's chairmanship, when the FOMC used a firmer and more explicit form of forward guidance language. Considering that the FOMC meetings (and so the minutes) follow a particular structure such that most of the topics are covered in each of them I didn't expect significant variation in the proportion of topics over time, including the topic devoted to monetary policy, the one under analysis. However, I expected to see a slightly greater proportion of content devoted to such topic starting from the end of 2008, when the FOMC, under Bernanke's guidance made greater use of its forward guidance tool in the post-meeting statements.

The technique I applied for my analysis is able to algorithmically identify the topics discussed in the FOMC minutes and their proportions over time. Thus, another second but important goal of this thesis is being able to identify properly, via the use of LDA, such topics and their mixture over time.

II. Theory and Literature Review

i. The Opening of the Temple

Thirty years ago central bankers believed their activities should be shrouded in mystery as this would enhance the effectiveness of monetary policy. Indeed, secrecy was almost seen as monetary policy tool. Alan Greenspan, Chairman of the Fed from 1987 to 2006, was very famous for his long, wordy and vague statements, practicing what at that time was called "constructive ambiguity" and today "Fedspeak" or "Greenspeak", and for his reluctance in moving toward greater policy transparency. As he later admitted, Greenspan (2007) was aware of the influence his words had and that he used this intentional strategy to prevent overreactions to his remarks by financial markets, and so to avoid a self-fulfilling prophecy. However, despite some concerns, starting from the early 90s there was a growing tendency within the Fed toward greater openness. In 1994 the FOMC not only began to publish the meeting transcripts (with a five-year lag) but also to release immediately following a scheduled meeting the post-meeting statements reporting the members' latest policy decision. Afterward in 1999, the FOMC also started to provide some indications of the expected or intended future path of policy rates in terms of "tilt" or "bias" and then "balance of risks" in their post-meeting statements. The Federal Reserve wasn't the only Central Bank moving in such direction as some form of forward guidance was already adopted by the Central Bank of New Zealand (from 1997) and Japan (from 1999). That is, over the last thirty years many Central Banks drifted apart from the long-term tradition of secrecy and not only became noticeably more transparent but also started to rely more heavily on communication for monetary policy-making. This revolution in thinking was motivated by different reasons. First of all, Central Banks act as public servants and so have the duty to be more open to the public about their decision-making process (Bernanke, 2004). Second, independent Central Banks should be accountable for their actions to their respective Government and, more generally, to the public. Third, as the notion of the importance of managing market expectations became increasingly widespread, many Central Banks started to use communication as a key tool to make monetary policy more effective and improve economic results, especially under exceptional conditions. As illustrated by many, such as Bernanke (2004) and Woodford (2005), although the public tends to focus on short-term rates, such as the fed funds rate, long-term ones are much more important to economic activities as most investment and borrowing decisions depend on their value. However, there is a link between the two: long-term rates reflect market expectations of the future evolution of overnight rates (such as fed funds rate) and other related short-term rates. This is because long-term rates can be decomposed into two components: the expected future short-term rates and a term premium[45]. Thus, as many economists argue, better communication about Central Bank's future short-rate setting contributes to the management of long-term ones and thus in the effectiveness of monetary policy.

ii. Text Mining Techniques to Central Banks' Documents for Content Analysis: Review of Literature

As a result of such global revolution, Central Bank communication nowadays can boast a vast literature, the majority of which focuses on the impact of communication on financial markets or on the link between different communication strategies and the economic performance (Blinder et al., 2008). Most of such focuses on quantitative data and on the application of data mining, techniques that enable to discover patterns in large data sets. More recently there is also a small but growing branch that focuses on the application of text mining on Central Bank's documents, techniques that are able to quantify text data, to infer their tone (for instance, positive or negative, dovish or hawkish[46]) and/or their content. For my analysis I make use of a popular text mining technique, called LDA, to infer the proportion of FOMC minutes devoted to the topic of monetary policy over time. Thus, this thesis contributes to a small but growing literature devoted to the use of text mining techniques to Central Banks' (in particular, Fed's) documents for content analysis. In the following paragraph I provide an overview of major works within such literature.

[45] Such premium reflects the remuneration investors ask for holding assets with longer maturities.
[46] See for instance, Lucca and Trebbi (2009), Bligh and Hess (2013) and Kahveci and Odabas (2016).

One of the very first applications of text mining techniques to Central Banks' documents for content analysis can be found in Boukus and Rosenberg (2006). The authors applied another popular technique, similar to LDA, called Latent Semantic Analysis (LSA) to extract the different topics of the minutes from 1987 to 2005. Their aim was to analyze if the Treasury yields reaction to the minutes release depends also on the specific theme expressed. This work was the very first application of LSA to the Federal Reserve's documents and set the stage for its future use. Hendry and Madeley (2010) performed a similar analysis to the statements of the Bank of Canada from 2002 to 2008, trying to identify the most valuable topics by financial markets. One recent application of LSA is a paper by Acosta (2015). The author applied such technique to both the FOMC minutes and the transcripts, from 1976 to 2008, to analyze how accurately the minutes convey the information discussed during the meetings (and so, the transcripts) and if such accuracy changed as the transcripts became public starting in 1994, that is, as the members knew their discussion would have become public. The most recent paper which uses LSA for topic modeling is the one of Mazis and Tsekrekos (2017) which analyses the contents of the FOMC post-meeting statements (from 2003 to 2014) and their impact on the U.S. Treasury market. The papers described so far all apply the Latent Semantic Analysis (LSA) technique to discover the latent topics. Very recently, also Latent Dirichlet Allocation (LDA), the method I chose for my analysis, has been applied to Central Banks' documents for content analysis. The very first application of such methodology in economics is the one of Hansen et al. (2017), who performed a similar study to the one of Acosta (2015), trying to infer how policy deliberations have evolved from 1987 to 2009 following FOMC members' knowledge of their discussion being public. Another relevant work is the one of Jagadeesh and Wu (2017) who studied the informativeness of the contents of the FOMC minutes from 1991 to 2015 for the stocks and bonds market. For my analysis, as Jagadeesh and Wu (2017), I applied LDA to the FOMC minutes to extract the different topics and compute their proportions over time. However, my work is different from theirs as I addressed a different research question and a different period. Moreover, I made different choices in the application of LDA. The most recent application of LDA to Central Banks' documents for content analysis is the one of Edison and Carcel (2020). The authors applied such technique to the FOMC transcripts from 2003 to 2013 to simply analyze the evolution of their contents over time. Although the minutes can be considered as a summary of the transcripts, and that both my analysis and Edison and Carcel's one address the same period, the results are difficult to compare because of the different choice of the number of topics and of the respective labeling. Another work worth mentioning within such literature is the one of Vallès and Schonhardt-Bailey (2015), who used the paid software Alceste, which, as LSA and LDA, is able to detect the latent topics of a set of documents. They studied how the contents discussed by the Bank of England's Monetary Policy Committee (MPC) during their monthly meetings and by their members in their speeches were shaped by the adoption of forward guidance in 2013. My analysis is similar to theirs as I also focus on how the adoption of forward guidance shaped the contents of the FOMC discussions. In conclusion, the literature here reported shows the benefits from the application of text mining techniques to Central Banks' documents, as they offer valuable tools for the extraction of the contents from text documents and for their transformation into quantitative data that can be used to extract valuable insights.

III. Methods

Below a chart that summarizes all steps implemented for my analysis that I describe along this section.

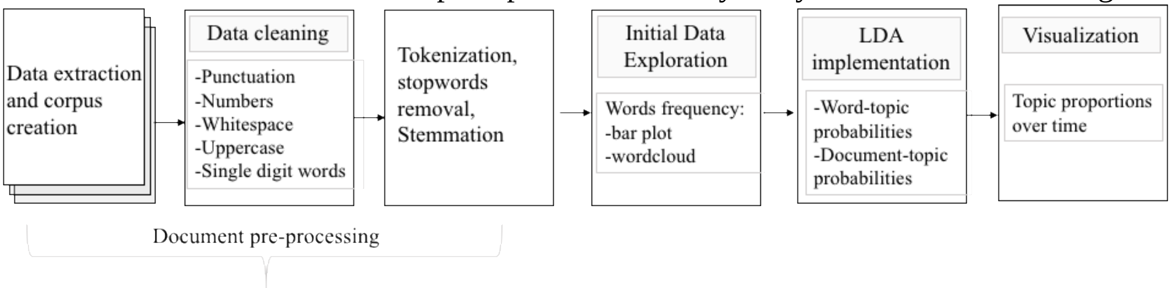

i. Description of the Data

I analyzed the minutes of all the scheduled FOMC meetings from 2002 to 2013, for a total sample of 96 meeting minutes. The minutes of the Federal Open Market Committee are an accurate summary of the contents discussed by the participants during the so-called "FOMC meetings". The FOMC started to draft the minutes since its earlier days, in 1936, and to publish them in 1967. The minutes are characterized by a common language and by a clear and recurring structure made up of four major parts: administrative matters; staff overview and outlook of the economic and financial situation; participants' view on economic conditions and outlook; and Governors' discussion on policy decision, outlook for future monetary policy and the vote as well as dissenting views on the policy action and on the post-meeting statement. I have decided to focus my analysis on the FOMC minutes instead of the post-meeting statements or transcripts because of the minutes' relatively better balance between information content and timing. Although the minutes are released after a three-week lag, whereas the statements immediately after the FOMC meeting, the minutes provide a much richer information content about the meeting than the statements. The transcripts, on the other hand, although provide the most detailed record of the meetings, are released after a five-year lag and thus lose power in terms of communication relative to the minutes. I chose to investigate the period of 2002-2013 as it was marked by greater policy transparency via the use of the forward guidance tool as well as by the leadership of two Fed Chairmen with a different perspective on transparency. In particular, I included also year 2002 to analyze the difference in contents between 2002, the year still characterized by the "balance of risks" language and 2003, the starting year of the forward guidance.

ii. Document Pre-processing and Initial Data Exploration

I freely downloaded each minutes between 2002 and 2013 from the Federal Reserve Board of Governors' official website and then transformed each file into a txt format, each saved with the corresponding date. Before starting to process the data, I removed manually from all the minutes the first section and the voting part, and directive to the Fed of NY of the last section as I did not deem them to be useful for my analysis. I removed also the second section since, like Edison and Carcel (2020), I was only interested to analyze the contents discussed by the Governors (and so, not the data and reports provided by the staff). I then broke each minutes into different paragraphs keeping record of the specific minutes where a paragraph belongs[47]. A total sample of 818 documents (paragraphs) was created. The first steps consisted of the download of the required packages, load of the paragraph documents, and creation of the corpus, that is, a collection of the text documents under mining. Then, as suggested by the text mining guide provided by Bholat et al. (2015, p. 5-8), I performed some pre-processing passages very common in any text mining analysis to prepare the documents in a format suitable for analysis. That is, I removed the numbers, punctuation, excess whitespace, single-digit words, and converted uppercase to lowercase. I proceeded by breaking the corpus into individual tokens or words, a process known as "tokenization", so that one token-per-document-per-row table was created. The latter format, known as "tidy text" format, is particularly useful for further text processing as it allows to use some common text mining packages. Afterward, I removed from the documents the so-called "stopwords", very frequent words that do not provide any insight about the contents of documents, such as *"the", "and"* (for a total of 1,149 stopwords), as well as an additional list of stopwords ("mystopwords") with meaningless words commonly seen in the FOMC language, such as *"committee" and "participants"*. I then stemmed the words, that is, removed the suffix or prefixed of words and maintained just their root forms so that the algorithm can recognize them as the same word. At the end of such pre-processing steps the number of total words was more than halved, from an initial of 185,249 to 84,404. To initially explore the data and so investigate the documents' most frequent words I created a Term Document Matrix (TDM), a matrix where each row represents one unique word, each column one document and each entry the frequency of a word in a document. The matrix created is characterized by 2,437 unique words and 818 documents. The next steps consist of the creation of the Document Term Matrix and of the application of LDA.

iii. Latent Dirichlet Allocation (LDA): a Text Mining Technique

Before explaining the remaining passages of my analysis, I will first give a brief overview of text mining, LDA, and Gibbs sampling (the method used to implement LDA) and my choice of the number of topics.

a) An Overview of Text Mining and its Application to Central Banks Documents

Text mining refers to the wide set of computational tools and statistical techniques that are able to quantify text data, that is, to treat it quantitatively to extract valuable insights (Bholat et al., 2015, p.1). They have been widely applied in many social science literatures, such as politics and marketing, but have been less frequently applied in economics, especially within Central Bank studies (Bholat et al., 2015, p.1). Indeed, most of the empirical works in economics tend to focus on quantitative data and to apply data mining techniques. However, there are many gains also from the application of text mining ones. First of all, they

[47] Splitting the documents into different paragraphs improves the performance of the LDA in the topic identification as paragraphs (although not all of them) are devoted to a lesser number of topics.

enable to analyze large sets of text data in a relatively shorter time and also enable to extract valuable information that may be disregarded by a human reader (Bholat et al., 2015, p.1). Second, they can be used for a wide range of purposes, the most common being information retrieval, content comparison, topic classification and modeling, among others. In such context, observations, instead of being numbers, are text data, also known as documents. Candidate text data for central banking analysis are Chairman and Governor speeches, meeting materials, reports, presentations, etc. Text mining techniques can be distinguished between supervised (require a user-classified set of documents) and unsupervised and between abductive (without generalization to other cases) and deductive. LDA the technique I applied for my analysis, belongs to the unsupervised and abductive techniques.

b) The Latent Dirichlet Allocation (LDA) Algorithm and Gibbs sampling

First developed by Blei et al. (2003), LDA is an unsupervised generative probabilistic topic model algorithm that is able to detect underlying topics (also known as latent variables) in a collection of text documents (that is, a corpus). The word "Latent" stands for hidden, in this case the topics; "Dirichlet" refers to the distribution of distribution, in this case each document can be described as a probabilistic distribution over latent topics, each of which is a probabilistic distribution over words; and "Allocation" refers to the action of allocating topics to the documents and words to the topics. It requires the number of topics T to be defined a priori by the user and to be fixed for the entire process. As described by Schwarz (2018), LDA relies on two baseline assumptions. The first is that documents with similar topics will use a similar group of words. The second is that each document (in our case paragraph) can be described as a probabilistic distribution over latent topics and each topic as a probabilistic distribution over words. In other words, each document is a mixture of topics and each topic is a mixture of words. Words identified by the model with a high probability of belonging to a certain topic cluster usually give a good insight into the content of such topic. Therefore, if a word has a high probability of belonging to a certain topic, the documents with such word will more probably belong to such topics as well. LDA is a generative process that works backward as if it were to construct the documents by itself. Its utmost goals are the estimation of the document topic mixture (θ) and the word distribution of each topic (ϕ). Without going deep into the math (you can find it in the extended version of this thesis), LDA maximizes a likelihood function which however is computationally unfeasible. Thus, some other methods for the LDA such as the Gibbs sampling have been developed. Gibbs sampling is a Monte Carlo Markov-chain algorithm that iteratively changes the topic assignment of one word, conditional of the topic assignment of all the other words. The method will first randomly assign a topic to each word in the corpus of documents (as said before, the number of topics must be chosen beforehand). As a consequence, each document will be randomly assigned to some topics. Afterward, the Gibbs sampling will assume that all the word assignments are correct but the current one in document d. The Gibbs sampling will assign a new topic to such word following a probabilistic model: it multiplies the proportion of words in document d that are currently assigned to topic t ($P\ (topic\ t \mid document\ d)$) to the proportion of assignments to topic t over all documents coming from that word (P $(word\ w \mid topic\ t)$). That is, in assigning the new topic to the current word, the Gibbs sampling will take into account both the proportion of words of document *d* assigned to topic *t* (excluding the current word) and the proportion of assignments to topic *t* over all documents that came from this word *w*. This process will be repeated for each word *w* of each document *d* of corpus D until a steady-state is reached. In summary, the power of LDA consists in its capability of clustering words belonging to the same topics without using any user pre-defined clusters of words. Like Latent Semantic Analysis, another popular technique for topic modeling illustrated before in the literature review, LDA does so by assuming words are not independent of one another but rather linked together by latent topics based on their appearance alongside each other. LDA only requires the number of topics *T* to be defined a priori. This advantage is what makes popular such method with respect to Boolean and dictionary-based techniques[48] (Bholat et al., 2015, p.11).

c) The choice of the number of topics

One of the main challenges of LDA is the ex-ante choice of the number of topics *T* as there is a trade-off between choosing few topics, with the risk of mixing different topics and loose interpretability, and choosing too many, with the risk of creating too specific topics and losing the general picture. Generally, the number of topics is chosen by the researcher based on the context of the data under analysis, although there are some statistical methods (such as k-means clustering) that are able to guide the choice. As Edison and Carcel (2020), I chose the number of topics arbitrarily, after a careful analysis of the content of the minutes over time. In particular, I performed the LDA process using a different number of topics, from eight to twelve. The aim of this step was to find the number of topics such that the top words from each topic cluster were as distinct as possible so that I was clearly able to identify a topic labeling from such words. I identified nine as the most suitable number of topics for my analysis.

d) Application of Latent Dirichlet Allocation (LDA) on FOMC Minutes

After the pre-processing steps described before (needed to prepare the documents in a format suitable for analysis), I transformed my text data into a Document Term Matrix (DTM) as this is the object of the function `LDA()` from the *topicsmodel package*, the one I decided to use for LDA implementation. A DTM is just the transpose of the TDM illustrated before. I used Term Frequency (TF) weighting for its creation as I judged it to be the most appropriate for my analysis[49]. The dimension of the Document Term Matrix obtained is 818 paragraphs times 2,437 unique words. Afterward, I implemented the LDA algorithm using as object the obtained matrix and as method the Gibbs sampling, the one I illustrated before, to create a 9 topic LDA model. Then, I used the function `tidy()` from the *tidytext package* to calculate what the function calls "betas", the expected word-topic-probabilities, that is the expected probabilities of words of belonging

[48] They require the user to pre-define a set of words of interest to identify the clusters of topics.
[49] This is not the only possible method for term-weighting. Please, see the extended version of this thesis for further details.

to a certain topic (what I previously denoted as ∅). For each topic *t*, the model will compute the probability of each word being generated from that topic, for a total of 21,933 probabilities (2,437 unique words times 9 topics). Of course, as previously illustrated, the higher the probability for a given word *v* the more likely such word characterizes topic *t*. Table 1 in Topics Labeling section reports the top 20 words for each topic.

e) The Estimation of the Proportions of Topics Over Time

Besides modeling the "betas", the `tidy()` function also estimates what it calls the "gammas", which represent the expected document-topic-probabilities, that is the expected probabilities of documents (in this case, paragraphs) of belonging to a certain topic (what I previously denoted as θ). For each paragraph *d*, the model will compute the probability of belonging to each of the topic *t*, for a total of 7,362 probabilities (818 paragraphs times 9 topics). I aggregated such gammas per year, normalizing by the different number of paragraphs per year in order to identify the topic proportions per document over the years 2002-2013. Then, I obtained a second figure which illustrates only the topic of Monetary Policy & Communication, the one of interest. The two figures are shown in the next section.

IV. Results

i. Exploratory Results: Words Frequency

As illustrated before, I computed the Term Document Matrix (TDM) with the aim of initially exploring the data. Therefore, I used the matrix to generate a frequency table and a word cloud, a popular text mining representation of words, where the frequency of each word is shown with font size and color.

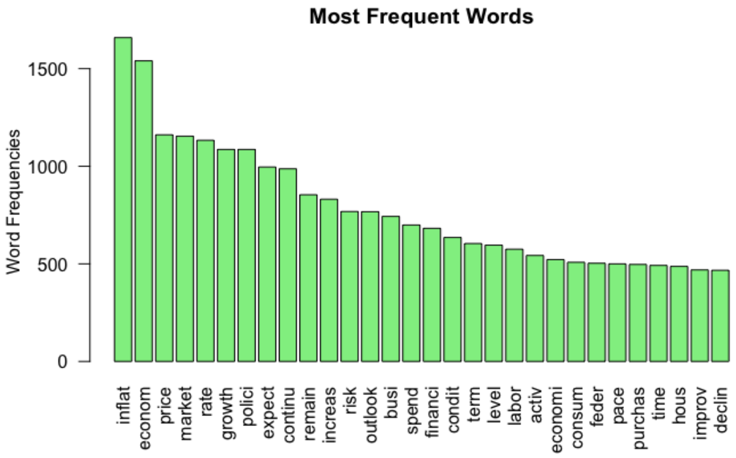

Figure 1: the 30 most frequent words

Figure 2: Word Cloud
The size and color of each word indicate the frequency

As can be seen, the seven most common words in the FOMC minutes are "inflation", "economy", "price", "market", "rate", "growth" and "policy", words closely related to some of the most discussed topics of Central Banks, that is, inflation, economy, and monetary policy. Other very frequent words are "business", "spending", "financial" and "labor", the first two related to key sectors, the third to financial markets, and the last to labor market topics. It is not surprising that the most frequent word in the minutes in the considered period is "inflation" (and the third one "price"). Indeed, price stability, generally interpreted as low and stable inflation, is one of Federal Reserve's two mandates (together with maximum employment). Another reason why it is not surprising is that Alan Greenspan and Ben Bernanke, the two Chairmen during the period analyzed, were inflation fighters first and foremost as they both accredited the benefits of low and stable inflation.

ii. Topics Labeling

I implemented the LDA algorithm with the aim of splitting the corpus into nine distinguishing topics. For each word cluster, table 1 reports the top twenty words with the highest "betas", the probabilities of words of belonging to a certain topic: the higher the probability for a given word *v* the more likely such word characterizes topic *t*.

Topic 1: Asset Purchase Program		Topic 2: Inflation		Topic 3: Labor Market		Topic 4: Fiscal Policy & International Economy		Topic 5: Current and Prospective Economic Developments & Economic Projections	
Word	**Beta**	**Word**	**Beta**	**Word**	**Beta**	**Word**	**Beta**	**Word**	**Beta**
purchas	0.0564	inflat	0.1451	labor	0.0585	growth	0.0798	econom	0.0700
addit	0.0353	price	0.0906	market	0.0379	economi	0.0452	rate	0.0437
asset	0.0319	expect	0.0558	employ	0.0348	fiscal	0.0346	expect	0.0392
term	0.0274	increas	0.0337	rate	0.0306	econom	0.0341	outlook	0.0382
secur	0.0230	pressur	0.0327	declin	0.0271	activ	0.0315	anticip	0.0287
agre	0.0223	energi	0.0247	unemploy	0.0264	recoveri	0.0202	level	0.0278
program	0.0218	cost	0.0245	continu	0.0261	demand	0.0197	run	0.0270
feder	0.0204	measur	0.0227	level	0.0261	support	0.0197	percent	0.0257
reserv	0.0199	remain	0.0216	suggest	0.0217	export	0.0185	continu	0.0253
outlook	0.0187	core	0.0212	low	0.0202	data	0.0158	pace	0.0250
pace	0.0186	resourc	0.0194	increas	0.0201	effect	0.0157	growth	0.0239
time	0.0165	risk	0.0171	remain	0.0176	contribut	0.0153	moder	0.0222
accommod	0.0159	rise	0.0155	product	0.0173	foreign	0.0149	gradual	0.0207
polici	0.0153	commod	0.0151	indic	0.0167	incom	0.0143	inform	0.0194
effect	0.0153	concern	0.0129	growth	0.0158	govern	0.0143	project	0.0191
monetari	0.0143	moder	0.0126	slow	0.0148	restrain	0.0134	consist	0.0191
stabil	0.0134	term	0.0116	improv	0.0137	uncertainti	0.0124	term	0.0179
potenti	0.0119	slack	0.0105	output	0.0128	stimulu	0.0119	remain	0.0158
hold	0.0117	read	0.0103	factor	0.0123	reduc	0.0117	judg	0.0154
condit	0.0114	util	0.0100	gain	0.0120	balanc	0.0110	inflat	0.0143
Topic 6: Financial Markets		**Topic 7: Economic Outlook**		**Topic 8:Key Sectors**		**Topic 9: Monetary Policy & Communication**			
Word	**Beta**	**Word**	**Beta**	**Word**	**Beta**	**Word**	**Beta**		
market	0.0859	outlook	0.0251	spend	0.0536	polici	0.0917		
financi	0.0771	econom	0.0248	busi	0.0446	econom	0.0401		
condit	0.0431	expans	0.0247	consum	0.0375	rate	0.0393		
credit	0.0372	uncertainti	0.0187	hous	0.0374	feder	0.0300		
bank	0.0254	economi	0.0182	sector	0.0307	fund	0,0299		
risk	0.0212	prospect	0.0171	household	0.0258	monetari	0.0256		
increas	0.0186	includ	0.0171	invest	0.0246	risk	0.0239		
remain	0.0171	signific	0.0171	report	0.0235	statement	0.0223		
improv	0.0169	factor	0.0168	sale	0.0206	agre	0.0220		
loan	0.0152	evid	0.0166	continu	0.0187	futur	0.0164		
eas	0.0136	strengthen	0.0159	incom	0.0160	stanc	0.0156		
declin	0.0127	persist	0.0151	price	0.0158	accommod	0.0150		
tighten	0.0101	anticip	0.0146	inventori	0.0148	forward	0.0150		
concern	0.0099	develop	0.0145	home	0.0147	target	0.0142		
real	0.0086	comment	0.0142	capit	0.0142	assess	0.0139		
strain	0.0086	product	0.0137	contact	0.0140	commun	0.0136		
continu	0.0084	effect	0.0135	activ	0.0117	action	0.0114		
deterior	0.0084	posit	0.0128	construct	0.0111	guidanc	0.0109		
liquid	0.0083	continu	0.0125	mortgag	0.0109	decis	0.0106		
investor	0.0083	provid	0.0125	firm	0.0104	view	0.0104		

Table 1: Distribution of the Top 20 Words per Topic

While on one hand LDA manages to cluster words belonging to the same topics, on the other, it does not provide topic labels, which must be defined by the researcher in accordance with the subject being studied. After a careful analysis, I identified the following topics associate to each cluster: Asset Purchase Program, Inflation, Labor Market, Fiscal Policy & International Economy, Current and Prospective Economic Developments & Economic Projections, Financial Markets, Economic Outlook, Key Sectors (which comprise Consumer Spending, Business Investments and Housing Market) and Monetary Policy & Communication, the one of interest. Another aspect worth noticing is that the probability characteristic of LDA is important, especially within this analysis, as it lets words to be associated with different topics. For example, the word "price" has a high probability of being associated with the topic of Inflation but also has a modest probability of being associated with the topic of Key sectors. Indeed, the level of prices is also relevant to the discussion of business and consumer spending.

iii. Proportion of Topics Over Time

As previously described, for each paragraph *d* the model computed the probability of belonging to each of the topic *t*. The sum of such probabilities for each paragraph is one and so they can be interpreted as the proportion devoted to a certain topic. Thus, a higher probability means a higher proportion devoted to a certain topic. I used the obtained gammas to identify the topic proportions per document over the years 2002-2013, as Figure 3 shows. As I expected, there is not a significant variation in the proportion of the topics throughout time. This is because the meetings (and so the minutes) follow a particular structure such that most of the topics are covered in each of them. Despite this aspect, we can still infer some relevant information. Inflation is amply discussed throughout the entire period, but especially between 2004-2006 when inflation pressures were a major concern for the FOMC.

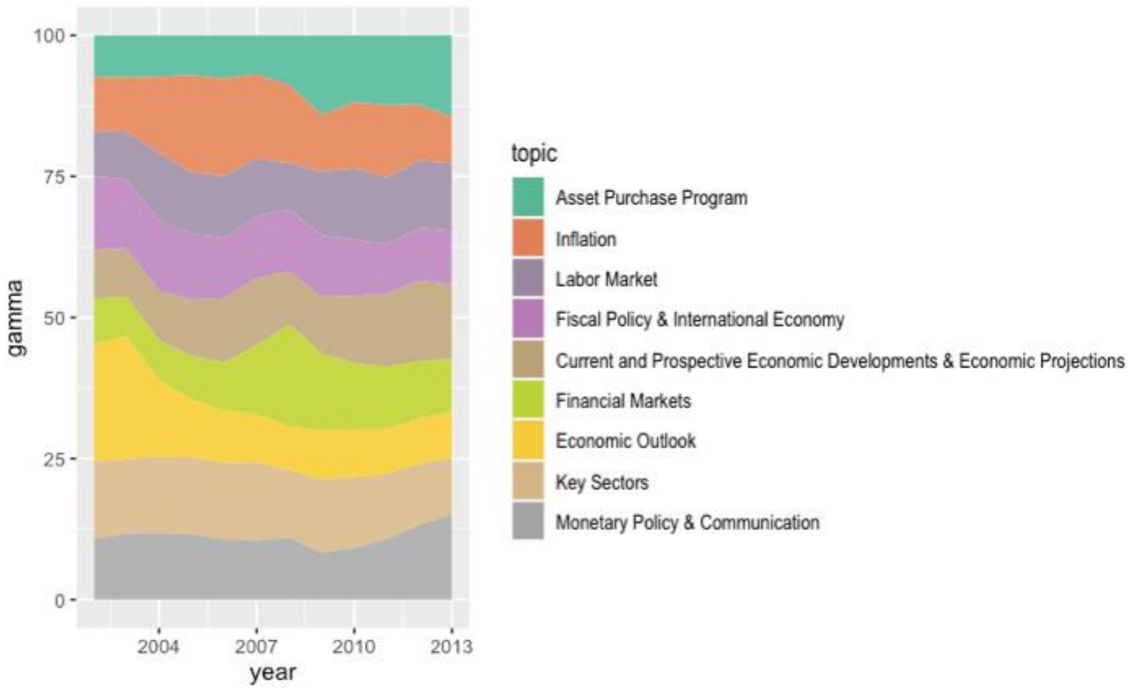

Figure 3: Proportion of FOMC minutes topics over time

Also the topic of Key Sectors is significant throughout the entire period as it encloses consumer spending, business investments, and housing market, information crucially important for the assessment of the economic conditions. As expected, due to the serious deterioration of financial markets, the proportion of the minutes devoted to such topic increased substantially during the years of the Global Financial Crisis. Following the same reasoning, also the topic of the Asset Purchase Program increased in the last part of the figure due to the large-scale program initiated in late 2008 (ended in 2014).

iv. A Focus on the Topic of Monetary Policy & Communication

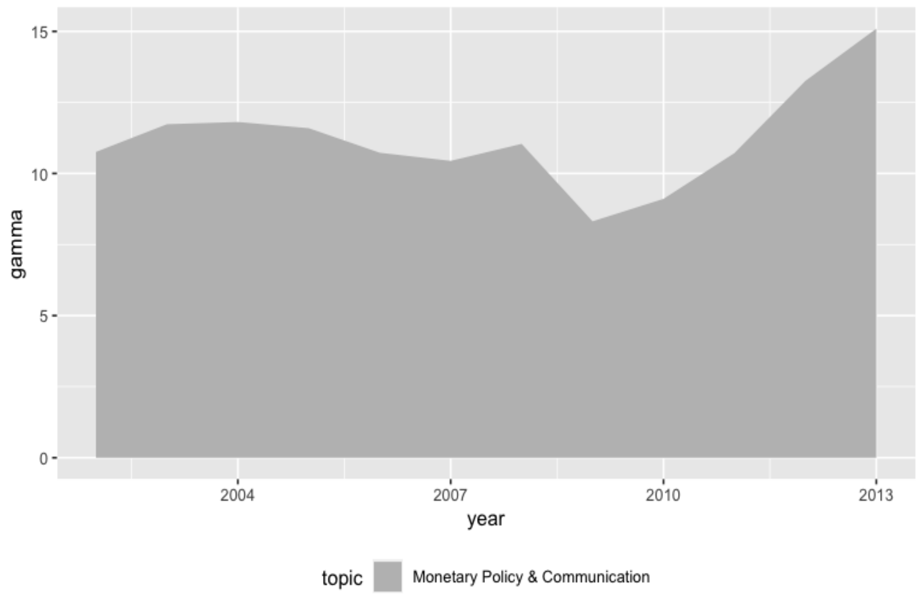

Figure 4: Proportion of FOMC minutes topic Monetary Policy & Communication over time

The goal of this thesis is to analyze if the increased policy transparency via greater use of forward guidance in the FOMC post-meeting statements has transformed the contents of their minutes between 2002-2013. In particular, I wanted to observe if this greater transparency in the post-meeting statement meant also a greater proportion of contents devoted to monetary policy (defined as references to current and future setting of the fed funds rate) in the FOMC minutes and if so, if this increase is more visible during and after the Global Financial Crisis, under Bernanke's chairmanship, when the FOMC used a firmer and more explicit form of forward guidance language. The findings support my initial hypothesis as there is not a significant variation in the proportion of topics over time (see figure 3), including the topic devoted to monetary policy, but still there is a visible increase in the proportion of the topic under analysis starting from the end of 2008 (see figure 4), when the FOMC, under Bernanke's guidance, made greater use of its forward guidance tool. In the following paragraph I will present and comment on the results in more detail.

The period under analysis is marked by the leadership of two Fed Chairmen, Alan Greenspan, at the helm at the beginning of the sample until January 2006 and Ben Bernanke, Chairman from February 2006 until the end of the sample. Alan Greenspan and Ben Bernanke share a lot of points of interest. However, they were different in many other ways. One of these was that while Greenspan was very well known for his “Constructive Ambiguity” language and for his initial reluctance in moving toward greater policy transparency, Bernanke, from the early days as Governor in 2002, was a strong supporter of transparency and policy plans communications. Although Greenspan set the stage, Bernanke steered the Federal Reserve in the direction of greater policy transparency, especially thought greater use of the forward guidance language during and after the crisis. The sample period starts in 2002 when the FOMC didn’t yet employ a direct form of forward language in its post-meeting statements. In August 2003, worried about a possible emergence of deflation, after lowering the fed funds rate almost to the zero-lower bound, the FOMC tried to manage expectations with a firmer and time-dependent form of forward guidance: *“policy accommodation can be maintained for a considerable period”*. Although such language started only at half year, we can see from the figure a slight increase in the proportion of the topic in the minutes between 2002 and 2003. The *“considerable period”* language was used for the subsequent meetings until the statement of January 2004 when *“can be patient in removing its policy accommodation”* was introduced, which was then substituted in May by *“accommodation can be removed at a pace that is likely to be measured”*, signaling that as the economic conditions were changing, the fed funds rate would have likely been raised soon, as it was the case. In June 2004 the FOMC repeated the same language of May and at every subsequent meeting the Fed increased its rate by 25 basis points until the last meeting of Chairman Greenspan, January 2006. Such forward language was then changed in December 2005 with some looser statements indicating some likely policy firming. The direct interest rate guidance that started in August 2003 ended in June 2006. Such a drop in the language is also reflected in the minutes (although not from the beginning) as we can see a decrease in the topic proportion until the end of 2008-beginning of 2009. The forward guidance language was then resurrected in December 2008, once the FOMC had clearer ideas regarding the economic outlook, with a purely qualitative language indicating that policy accommodation would have continued for *“some time”*. The December 2008 statement was slightly modified in March 2009, when *"for some time"* was replaced with *"for an extended period"*. Such language was maintained until June 2011, after which the Committee significantly changed the statement with a less qualitative and more explicit language introducing a specific date. In August 2011 the statement indicated that *“likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013”*. This type of language is also referred to as date-based forward guidance. Such date was then pushed out twice, *"at least through late 2014"* (January 2012) and *"at least through mid-2015"* (September 2012). Afterward, starting from December 2013 until the end of the sample, the language underwent a second important change as the future path of the policy rate was tied to the Fed’s economic objectives. As Bernanke (2013) stated, the FOMC believed the date-based statement to be limited and opted for a language that explained how future policy would be affected by future economic conditions, the so-called “state-contingent language”. In conclusion, during and after the Global Financial Crisis, the FOMC under Bernanke used a firmer and more explicit form of forward guidance language. If we consider the last part of the figure, we can infer that this increase in policy transparency via greater use of forward guidance in the post-meeting statement is also reflected in the minutes, as we observe an increase in the proportion of topic devoted to monetary policy starting from the end of 2008-beginning of 2009 until the end of the sample.

Another aspect worth discussing regards the Large Asset Purchase Program (also known as Quantitative Easing). As figure 3 shows, the Asset Purchase Program topic proportion increased starting from late 2008 when the FOMC, besides lowering the fed funds rate and using the forward language tool, provided additional stimulus via the unconventional tool of Large-Scale Asset Purchase Program (LSAPs). As Bernanke (2013) illustrated, both forward guidance and LSAPs affect longer-term interest rates but in different ways: forward guidance affects long-term rates by shaping the expectations of the future short-term rates whereas LSAPs by influencing the term premium. However, the LSAPs can also be seen as a commitment of a Central Bank to maintain the target rate to zero for a certain period[50] and thus it may also affect the short-term market expectations. Therefore, the increase of the proportion of such topic visible in the last part of the figure can be partly considered as a contribution to greater monetary policy transparency in the minutes.

V. Discussion

As illustrated in the previous section, the findings support my initial hypothesis. The same research question of this thesis can also be addressed to the transcripts. If the LDA process is applied in an analogous way the results would probably be the same, as the minutes can be considered as a summary of the transcripts. It would also be interesting to apply the same research question to a different period or to Greenspan's and Bernanke's speeches to observe the difference in the proportion of contents devoted to monetary policy transparency of the two Chairmen. Moreover, through the use of topic modeling techniques such as LDA, future researchers can investigate in more detail the evolution over time of other contents, linking the results, for instance, to the FOMC use of different monetary policy rules over time.

A second but important goal of this thesis was also being able to correctly identify, through the use of LDA, the topic discussed by the FOMC during their meetings. Without an appropriate word clustering each document could have been assigned to the topics with improper proportions, thus invalidating figure 3 and so the entire analysis. As we can see from table 1, the top words from each topic cluster are mostly distinct, enabling us to recognize without doubts the relevant topic labels for each of them. Although its application is not simple, the LDA is a very powerful tool, with clear advantages with respect to other text mining techniques, and so I expect to see in the future a greater application of it in economics.

[50] This is because, by implementing a Large-Scale Asset Purchase Program a Central Bank commits itself to keep reserves well above those needed to maintain the benchmark rate to zero (Bernanke and Reinhart, 2004)

Although topic modeling studies such as mine overcome the subjective scoring of the manual approach, still their process and so the output is influenced by the researcher's judgment and knowledge on the subject. Thus, for instance, studies that make use of the same topic modeling technique to the same or similar documents still may produce different results. As I mentioned before, the technique I applied, LDA, requires the number of topics to be defined ex-ante by the user. Furthermore, LDA does not provide topic labels, which must be instead identified by the user in accordance with the subject being studied. These aspects introduce some form of subjectivity in the analysis.

VI. Conclusion

In the past thirty years we have seen a revolution in the Federal Reserve's communication. Before such openness the Fed provided little information to the public about its policy decisions, fearing the market would overreact or that communication about the future course of fed funds rate would have reduced FOMC discretion to adjust it according to new economic conditions. That is, there was a widespread belief among central bankers that openness could damage the effectiveness of monetary policy. During those years, indeed, the market was accustomed to look how the FOMC systematically responded in the past to certain economic conditions to infer its next policy moves and to a monetary policy that relied on changes to the fed funds rate. The first publication of the post-meeting statement in 1994 marked a turning point. That revolution in thinking was motivated by different reasons but mostly by a considerable amount of evidence of enhancement of the effectiveness of monetary policy via the management of the future fed funds rate expectations. Although with initial reluctance, starting in 1999, the FOMC provided some forms of future rate guidance in the form of "tilt" or "bias" and then "balance of risks". Afterward, in August 2003, fearing a possible deflation, the FOMC issued what is today recognized as its first forward guidance statement. As Bernanke's successor Janett Yellen (2013) commented, *"The FOMC had journeyed from "never explain" to a point where sometimes the explanation is the policy".* The use of forward guidance accelerated during and after the Global Financial Crisis when the FOMC couldn't rely anymore neither on its past behavior for communicating its future moves nor on its conventional tools to stimulate the economy. Bernanke's successors, Janet Yellen and Jerome Powell also embraced the benefits of clear communication with the public, for instance, supplementing post-meeting statements with more frequent press conferences. To conclude, the Federal Reserve's communication policy evolved over time according to the evolving needs of the Committee. As the coronavirus demonstrates, surely new challenges lie ahead of the Fed. Those will inevitably lead to an evolution in the Federal Reserve's conduct of monetary policy as well as communication practices to preserve its maximum employment and stable price goals.

Page intentionally left blank