Communication, Monetary Policy, and Financial Markets in Mexico*

Ana Aguilar
BIS

Fernando Pérez-Cervantes
ITAM†

Abstract

We determine if the communication of private banks to their clients with financial interests in Mexico changes or not after Mexico's Central Bank communicates its monetary policy decision (MPD) and also two weeks later, with the publication of the minutes of Mexico's Central Bank monetary policy decision (MMPD) between 2011 and 2019. We use unsupervised Natural Language Processing (NLP) techniques to turn the text that private banks send to their clients about the Mexican economy into vectors of topics. We find that every time, private banks cover a large diversity of topics and words before the MMPD with no evident consensus of topics, and that almost always the quantities of terms and topics are reduced and repeated by almost every bank after the MMPD indicating some surprise (notable exception: the *liftoff* in December 2015), and that the topics vary depending on the date of the MMPD. The fact that private banks discuss the same topics and write to their clients with sentences that contain the exact same words indicates that the private banks react to the MMPD, independent of their opinion about the Central Bank's statements. We also found weak evidence that a measure of the size of the changes in the private bank's communication with their clients is positively correlated to changes in the long-term yields but negatively correlated to the size of exchange rate movements.

JEL Classification: C6, E5, E6

Keywords: Natural Language Processing, Unsupervised Sentence Embedding, Central Bank Communication, Mexico

*We thank Oscar Cuéllar-Nevares for superb research assistance. Pérez-Cervantes gratefully acknowledges financial support from Asociación Mexicana de Cultura A.C.

†Department of Business Administration. Río Hondo 1, Ciudad de México 01080, Mexico

1 Introduction

There has been a long and vast debate regarding the benefits of greater Central Banks' transparency and communication. Up to some extent, all modern Central Banks have moved to a strategy that includes better communication tools, more transparency and higher accountability.[1] Associated to this trend, in the last few decades the literature related to this subject has expanded tremendously. Add the recent development of intensive computational tools that transform text into data, and the long and vast debate regarding the benefits of greater Central Banks' transparency and communication has become itself a source of information for any analysis of Central Banks' communication. In a survey paper, Blinder, Ehrmann, Fratzscher, De Haan, and Jansen (2008) state that the literature in this topic has two main strands. The first is regarding the Central Banks' communication effect on financial markets. The second is dedicated to analyze and compare the different strategies of communication across Central Banks or changes in communication through time for a given Central Bank. And any of the two strands can be quantitatively measured and not just qualitatively using the aforementioned computational tools for text analysis.

In this paper, we will focus on a particular aspect of the first strand: we will quantitatively track the changes in what analysts in private banks write to their clients the day before and the day after Mexico's Central Bank (Banxico) publishes *something*. In particular, we focus on the monetary policy decision (MPD) and its minutes of the monetary policy decision (MMPD) which for the case of Mexico, includes an overall assessment of the economy. The monetary policy decision is known two weeks in advance of the publication of the MMPD, but not Banxico's assessment of the economy. That is, analysts can speculate and at most predict the reasons for Banxico's decision, but they won't know for sure until the MMPD is

[1] See Filardo and Guinigundo (2008), Bholat, Hansen, Santos, and Schonhardt-Bailey (2015), and Kahveci and Odabaş (2016) for relatively recent surveys on international evidence for communication by Central Banks.

published. In this paper, we check if the changes in messages sent by private banks to their clients before and after the MMPD, maybe due to the MMPD, affect financial markets.

Monetary policy decisions affect financial markets and the expected monetary policy stance not only through the decision itself (increase, decrease, or no change in the reference rate) but also because the decision is an outcome variable of the Central Bank's overall assessment of the economy. This assessment is extensively covered in the minutes of the decision in Mexico, and it is subject to a lot of educated guesses by analysts.

If the true message from the Central Bank to the economy (published in the MMPD) is transmitted (or amplified) via the private banks, we can use Natural Language Processing (NLP) techniques to measure the effectiveness of this transmission mechanism. If communication is ineffective (i.e. what Banxico writes does not affect financial markets), we can check if Banxico's message was not transmitted by the private banks, or if the message by private banks about what Banxico wrote did not influence financial markets, either because Banxico's assessment of the economy is in line with the educated guesses, meaning that the message was transmitted the day of the publication of the monetary policy decision, or because what private banks write does not affect financial markets.

We first check the effectiveness of the MMPD by checking whether the analysts of private banks (usually the Director of Research, or the Chief Economists) that send memorandums to their clients with financial interest in Mexico (which we denote as *analysts* or simply *banks*) discuss similar topics in the days prior and following the publication of the minutes. Mexico's Central Bank is a unique case to do such an analysis because its MMPD is very structured in 5 sections, namely (1) external outlook, (2) domestic financial markets and macroeconomic stance, (3) economic activity, (4) inflation outlook, and (5) monetary policy forward-looking final paragraph. Also, Mexico is the only country in the region that has increased in 4.25 pp the reference interest rate after the Global Financial Crises due to several external shocks that affected inflation. This means that all monetary policy, the Bank's communication, and the Bank's outreach have been systematically active since late 2015 giving interesting variation to study. To the best of our knowledge, this is the first paper that has ever studied divergence or convergence of *topics* prior and after the MMPD in Mexico.

In addition, a Central Bank can move financial markets directly, affecting the short-term rate and other financial variables. It can also steer expectations of those variables. In this regard, there is "short-run" Central Bank communication, which refers to the economic outlook and forecasts that influence the financial markets. There is also "long-run" communication, which defines the Central Bank goals and overall strategy that provide an anchor to long-run inflation expectations. In regard the short-term influence, there are some clear advantages of a higher quality –not necessarily more- communication: first, Central Bank's communication increases the predictability and understanding of monetary policy authorities' actions. It increases the "signal-to-noise" ratio, enhancing the quality of the information set available for the markets, firms and households. Therefore, anyone could be able to make decisions of investment and consumption for longer horizons and with more certainty. Second, good communication can enhance the transmission from short-term rates (controlled by the Central Bank) to longer ones, increasing the effectiveness of the monetary policy transmission mechanism. A greater proportion of the credit and financial transactions are referred to long-term rates, instead of short ones.

The related literature has taken advantage of the computational capacity that nowadays allows researchers to use text analysis to study the Central Bank messages and their impact on the markets. In the financial markets, the reports that economists from many financial institutions make regarding the Central Banks (the treasury department and other public policies) have become a very important piece of information. In those reports, well trained economist scrutinize all Central Bank communication carefully. They have reports prior and posterior to important Central Bank communications. This paper relies on the assumption that the readers of these reports pay attention to what is written there, are informed *because* of these reports, and respond in the financial markets to the best of their abilities given the information.

Therefore, it is possible to study the impact of the Central Bank communication on the topics they discuss. Specifically, we try to assess if after a Central Bank monetary policy communication there is more congruence or commonality among economic analysts regarding the most important concerns for the Central Bank at the time of the decision. This is an indirect channel of the impact of Central Bank's messages to the financial markets since analysts' reports are in general a product for their clients. Based on their analysis and the interpretation of the Central Bank communications they make recommendations of portfolio positions in the fixed income market, the FX market, and equity market to their clients, investor, and fund managers, among others. For example, if in any given monetary policy decision, the Central Banks assessment of the economic activity reflects concern due to a larger slack in the economy and therefore the Board states that they do not foresee inflation pressures coming from the demand side in the relevant policy horizon, this is a very important piece of information for the market. It is very likely that after the monetary policy statement, most analysts will highlight this in their reports, and even they may show their own economic slack estimations. This could translate into portfolio recommendations for their clients, based on a more *dovish* stance from the Central Bank.

Also, there is full circle. Central Banks follow carefully what analysts and specialised journalists write about its communications to assess if the intended message was transmitted adequately to the specialists and, in general, to the public. Central Banks can look at financial variables to confirm if they were successful with the message. However, analysts provide an additional source of information, since they explain directly if they understand the Central

Bank's decision, if they agree with the Central Bank reading of the economic outlook and its forecasts, if they interpreted correctly the signals about future policy intentions, among other aspects. In addition to that, financial analysts maintain a close and fruitful relationship with authorities. There is a close and frequent communication among them. This is through the visits from analysis and their customers (investors, foreign mainly) with authorities to explain (besides formal communication) the Central Bank view of the economy and to answer direct questions from investors and fund managers. This communication channel is very important for emerging economies, recipients of large capital inflows and where less analysis and forecasts are made in comparison with advanced economies. Thus, studying the interpretation of financial analysts from Central Bank communications is a novel and complementary source of information for Central Banks and in general for the monetary policy communication literature.

There is another line of research, focused in transforming the Central Bank communication into sentiment indicators and studying directly their effect on financial variables. To do so, one must impose certain interpretation of the communication (e.g. labels of what is a *hawk* or a *dove* sentence or statement). We avoid that by keeping the computational text analysis as general and free as possible, and we focus on the impact of communication on the economic analysts.

This paper is organized as follows. Section 2 describes the Natural Language Processing (NLP) techniques used to read millions of paragraphs in Spanish, which then guide the computer to classify both Banxico's MMPD and the analysts sentences into vectors of information. It also describes the data used both for training the computer and to check the analysts' response to the MMPD. Sections 3, 4, 5, and 6 describe the results both of the techniques used in section 2 and with the financial data, and finally section 7 concludes and describes possible further research with the same techniques.

2 Data and Methodology

Our data consists on raw text of the analysts' reviews of Banxico's policy decisions and minutes between 2011 and 2019 and on the text of Banxico's minutes. There are 24 unique banks that review Banxico's minutes, and they issue four reports for every monetary policy decision: 1. Before the announcement of the monetary policy decision, 2. After the communique of the monetary policy decision, 3. Before the release of the minutes of the meeting for the policy decision (MMPD) and 4. After the release of the MMPD. These reports have in average 13 and 28 sentences respectively. Not every bank sends their clients text about every minute, and there are 4380 pre-release and 16860 post-release sentences.

We now describe the Natural Language Processing (NLP) methods used to assess the effectiveness of Banco de México's communication between 2011 and 2019. We use unsupervised sentence embeddings to simplify common NLP tasks such as classification and semantic similarity. The lack of supervision is important since there is no subjectivity implied in the model. While we humans can possibly classify a statement in a low dimensional space for its "hawkishness" and "dovishness", a computer can do this in many dimensions. This technique can therefore extract more information, which, as will be seen, is useful to explore and interpret what the main components and tones of a text corpus are. "Tone" in this setting ends up being what the computer found to be the way the analysts write about, but also the emphasis, repetition and composition of words that are used accompanying the economic concepts.

During the last decade, NLP saw a rise in popularity and many algorithms were developed, many tasks that were previously considered very difficult are now easily solved by simple baseline methods. Many of them relied on computing a set of high quality vector representations of words that reduce the dimensionality of a word dictionary. Simple methods such as Latent Semantic Analysis existed long ago, but Neural Networks and other recent algorithms outperformed them. In a Neural Network, large quantities of raw text are fed into a language model in which words are represented by a continuous vector learned by the algorithm. Using this vector representation, the Neural Network attempts some task such as predicting words given their context (see Mikolov, Sutskever, Chen, Corrado, and Dean (2013)). Once the model is fitted, the inferred word vector encodes information of the *words* themselves. In general, the representation of words as continuous vectors is called the Vector Space Model and word vectors are called word embeddings.

Learned word vectors have been found to accurately capture semantic and syntactic structures. This gives them great representation power which is useful in many NLP tasks and thus they have been studied thoroughly. In this vector space, "near" words share similar meaning and satisfy some algebraic regularities. Common examples are relationships like the ones below:

$$\text{"king"} - \text{"queen"} \approx \text{"man"} - \text{"woman"} \tag{1}$$

$$\text{"brother"} - \text{"man"} + \text{"woman"} \approx \text{"sister"} \tag{2}$$

A strong advantage of vector space models is that no explicit supervision is needed. These models are fitted using only the relative position of words in a text corpus. This makes their bias dependent only on the selected text corpus and not on human tagging or classification. That is, we are not forcing the computer to find something for us (such as a "hawkish sentence") but let it decide whether or not to cluster those types of sentences as to be in the same group of "tone". Common algorithms, such as GloVe (see Pennington, Socher, and Manning (2014)) and Word2Vec (see Goldberg and Levy (2014)), generate vector representations of words through their context, mapping words that have similar meanings close to each other in a vector space that resembles their actual semantic relationships which are not chosen by a human. Composition of topics in a document can be modeled using Latent Dirichlet Allocation (see Blei, Ng, and Jordan (2003)), or using human classifications in vectors such a vector counting how many times the words "inflation" or "GDP" appear in a text. Primitive approaches to measure semantic similarity between documents usually represent them as a set of words and measure their similarity.

For the Spanish language, we computed word embedding using the popular algorithm GloVe, which is considered good at capturing semantic information of words and is fairly simple. We fitted a custom GloVe model using a Spanish text corpus with approximately 800 million words, containing the Spanish version of Wikipedia, a sample of some years of around 200 Mexican news sites with a focus on economic, political and financial news, and diverse content regarding monetary policy, central banking and financial economics. Fitting our word embeddings with a suitable corpus such as this one is essential for getting good quality vectors in this domain.

Although some research has been conducted, there is no generally accepted rule for choosing the dimensionality of word embeddings other than considering that the broader the domain, the higher the dimension should be.[2] For regular models, vectors of dimension between 50 and 500 are common, although for very big models sometimes up to 1000 dimensions are used. For choosing the best embedding dimension, it has been found that, depending on the training corpus, dimensions that are lower than some threshold may restrict performance, while by increasing the upper bound the performance tends not to decrease much and the restrictions come in higher computational requirements. We thus use the common upper bound of 300 for our word vectors, considering that the training corpus is broad, and we do not want to restrict word vector expressiveness, while our model needs are not those of the

[2] See Patel and Bhattacharyya (2017) and Das, Ghosh, Bhattacharya, Varma, and Bhandari (2019) for some examples.

biggest models reported.

The rest of the paper does three things, based on this methodology. We first make a qualitative assessment of how well does the algorithm work in texts about the Mexican economy, given that they most likely have a different structure than any random text in Spanish. Second, we describe how does the frequency of the topics detected by the algorithm change as information by the Bank of Mexico is revealed (for example, before vs. after the decision, before vs. after the minutes, and after the decision but before the minutes vs. after the decision and after the minutes). Last, and following standard literature, we will define the magnitude of the changes in the frequency of topics as the magnitude of surprises in the information revealed over time to study how does the magnitude of the surprises translate to changes in financial variables.

3 Results: Qualitative Assessment of the Financial Topics

Our algorithm is trained with 800 millions words in Spanish. Only a small fraction of them belong to sentences written in financial news articles, the minutes of the Central Bank's decision, or the analyst's texts to their clients. The majority of the text used to train our algorithm comes from Wikipedia articles discussing just about every topic in the Spanish language. So, a natural question to evaluate this algorithm is how well does it separate the financial topics, since we want to study financial texts. If every word in the analysts' texts is lumped in a single topic, e.g. a unique "financial topic" then there is no variation to study in our analysis. But if words are successfully and systematically separated into different topics with a natural interpretation and a natural separation, then we can use this variation to draw some conclusions about the words that Banxico uses, or the words that the analysts use. So, in this section, we study word embeddings, that is, the 300-dimension vector that represent each word in Spanish. For simplicity, we illustrate the results with an example.

Take the following 11 words in Spanish: (1) Banxico, (2) inflación, (3) monetaria, (4) minuta, (5) economía, (6) consumo, (7) ingresos, (8) fiscal, (9) crecimiento, (10) Hacienda, and (11) haciendas. They, respectively, mean (1) Central Bank of Mexico, (2) inflation, (3) monetary, (4) minutes (and in Spanish, this word is not an homonym of the unit of time), (5) economy/economics (in Spanish, these words are homonyms), (6) consumption, (7) revenues/number of people entering somewhere (in Spanish, these words are homonyms, but in singular then there are at least 5 concepts for the word *ingreso*), (8) fiscal/prosecutor (in Spanish, these words are homonyms), (9) growth, (10) Ministry of Finance in Mexico/large colonial house (in Spanish, these words are homonyms), and (11) large colonial houses.

An educated reader will immediately notice that the first 10 words are "financial" words in the sense that if they appear in a sentence, then the sentence is most likely discussing something financial or about the economy, and, conversely, if it is known that a sentence is "financial" then the probability of seeing that word spikes (and the probability of many other words plummets). An even more educated reader will also notice that the first four words are not only financial, but could be easily classified as "Central Banking" words. Then as the list goes on, they slightly become more neutral (but still "financial") and then eventually become "Ministry of Finance" words. The 11th word is not a financial word, but it is the plural of the homonym of a financial word. The general idea of the Latent Dirichlet Allocation methodology (explained in section 2) is the following: is it possible to feed a computer with hundreds of millions of sentences in Spanish and transform this computer into an *educated reader*, without any human intervention? How about an *even more educated reader*? Can it find that *Hacienda* has two meanings from the words around it? Can it find that *haciendas* does not have two meanings although it is maybe written in the same sentence as *Hacienda* with very high probability? Can it find that discussing the economy is not an exclusive subject of

Pairs of Words	Banxico	inflación	monetaria	minuta	economía	consumo	ingresos	fiscal	crecimiento	Hacienda	haciendas
Banxico	1	0.582	0.617	0.529	0.22	0.13	0.17	0.251	0.243	0.203	0.009
inflación	0.582	1	0.669	0.344	0.356	0.278	0.322	0.3	0.432	0.113	-0.007
monetaria	0.617	0.669	1	0.394	0.415	0.156	0.227	0.411	0.291	0.164	0.031
minuta	0.529	0.344	0.394	1	0.066	0.089	0.112	0.184	0.039	0.17	-0.011
economía	0.22	0.356	0.415	0.066	1	0.353	0.412	0.32	0.476	0.364	0.152
consumo	0.13	0.278	0.156	0.089	0.353	1	0.37	0.177	0.376	0.135	0.021
ingresos	0.17	0.322	0.227	0.112	0.412	0.37	1	0.41	0.376	0.37	0.167
fiscal	0.251	0.3	0.411	0.184	0.32	0.177	0.41	1	0.19	0.488	0.014
crecimiento	0.243	0.432	0.291	0.039	0.476	0.376	0.376	0.19	1	0.093	0.063
Hacienda	0.203	0.113	0.164	0.17	0.364	0.135	0.37	0.488	0.093	1	0.481
haciendas	0.009	-0.007	0.031	-0.011	0.152	0.021	0.167	0.014	0.063	0.481	1

Figure 1: Cosine of the angle between 11 selected "financial" words

the Central Bank or the Ministry of Finance? Can it find some pattern that we are not even aware of?

Figure 1 contains the cosine of the angle between the 300-dimension vectors of the 11 words. Take the formula of the cosine of the angle between two vectors in equation 3

$$\cos\left(\theta_{x,y}\right) = \frac{x \cdot y}{\|x\| \, \|y\|} \tag{3}$$

where x is a 300-dimension vector with a word, y is a 300-dimension vector with another word, $\cdot$ is the dot-product operator, $\theta_{x,y}$ is the angle between x and y, and $\|x\|$, $\|y\|$ are the norms of vectors x and y. The larger the cosine, then the "more parallel" the vectors are (meaning the angle is smaller). As the cosine gets closer to the zero value, then the angle is larger and approaching 90 degrees (meaning that the vectors are closer to orthogonal, or linearly independent). Two words with a large cosine of the angle between them, in this context, imply they have very similar relative intensities in the 300 dimensions of the embedding. Two words with a small cosine of the angle between them imply that in whatever dimension some word is strong, then the other words is not, and vice versa. Two words with negative cosine imply that in whatever dimension some word is strong, then the other words is strong but with the opposite sign.

From this figure, it is noticeable that the "Central Bank" word pairs are more parallel than any other type of pair, and that the cosine is small between these words and the "Ministry of Finance" words. Notice that the "Ministry of Finance" words have small angles between them, and that *haciendas* is basically orthogonal to every words except ingresos (revenues) and Hacienda (which is the singular of haciendas). Also notice that words such as economy/economics and growth are not particularly linked to an entity such as the Central Bank or the Ministry of Finance, but they are both strongly linked.

We did hundreds of such comparisons, with very similar conclusions: the algorithm is capable of understanding the general meaning of words, and the general use and context of words. We can say that, for the purposes of this paper, the computer became an *even more educated reader*. Now in the next section we check what does it find when we make it read all the minutes, and all the analysis by private banks over a period of 9 years, but instead of analyzing individual words, it analyzes phrases by "adding" all the words (vectors) in the same sentence.

4 Results: Words into Sentences, Sentences into Topics

Now that we know the computer can appropriately represent each word in a natural way, we need to look at the word interactions by adding up the embeddings in the words of each sentence in a text. In order to capture the compound meaning of multiple words in sentences, which in turn can translate into representing semantic content like sentiment features or topics, an option is to compute analogous sentence vector representations. Many methods exist for obtaining in-domain and general-purpose sentence vectors. Surprisingly,

Por su parte el documento senala que se estima un crecimiento inercial de los ingresos tributarios en linea con la actividad economica del pais
Ademas se espera que el endurecimiento de las condiciones fiscales se espera que el deficit fiscal caiga un digito digito del pib este ano afecte la actividad economica
El igae de julio de hoy proxy del pib mensual sorprendio en el lado negativo con una tasa de crecimiento interanual del digito podria estar senalando la tan esperada desaceleracion del consumo del sector privado ademas de la debil inversion privada y publica que compromete el pib presupuestario del proximo ano...
Creemos que el objetivo fiscal implicara un superavit primario de una magnitud importante probablemente un superavit primario del digito digito del pib
El dato del crecimiento del pib en el primer trimestre digito a a la menor tasa de crecimiento anual desde digito confirma las senales de debilidad de la actividad economica sobre todo en el sector industrial
Hemos comentado que una baja del pib petrolero de digito sumado al efecto negativo del recorte de gasto publico restaria aproximadamente digito puntos al crecimiento del pib lo cual consideramos podria compensarse por un crecimiento mas robusto de la economia de Estados Unidos
Los datos del pib del segundo trimestre mostraron que la economia habia acelerado su crecimiento
Sin embargo el bajo desempeno de la actividad economica desde el cierre del digito que se prolongo hasta los primeros meses del digito determina que la prevision apunte a un menor crecimiento del pib
Del lado positivo una actividad mas dinamica en los ee uu aunado al estimulo fiscal por parte del gobierno beneficiaran a la actividad economica este ano
Se preve que el crecimiento del pib se acelere a digito el proximo ano
Un miembro dijo que esto es coherente con un ritmo mas lento de crecimiento economico el deficit de la cuenta corriente probablemente disminuyo a fines de digito y la desaceleracion de los prestamos del sector privado
El estancamiento de la actividad economica
Los ingresos registran un crecimiento de digito anual real en donde se incluye el remanente de operacion trasferido por el banco de mexico mientras que el gasto neto presupuestal crecio solo digito
El digito restante es esencialmente solo para compensar los menores ingresos fiscales este ano como resultado de un crecimiento mas dÃ©bil

Note: the word "digito" represents any number including a year, a percentage, etc.
Figure 2: The nearest neighbors of "Se estima un crecimiento inercial de los ingresos fiscales con la actividad economica del país"

even simple models like word vector averaging usually obtain promising results, sometimes even outperforming complex models as Recurrent Neural Networks (see Wieting, Bansal, Gimpel, and Livescu (2015)). Depending on the method, sentence vectors have been shown to capture features such as sentiment and topic. For classifying a sentence into a sentiment or topic, a very simple model can yield very good results if used on top of these sentence vectors, however, we use a model that, in this context, misses sentiment features *on purpose*. We do not use this feature because we do not want to know if the banks agree or not with Banco de México's overall assessment of the economy. We just want to know what do the banks write about and use that information to make assessments about commuication strategies.

For computing sentence vectors we use a simple method. We average word vectors in a sentence written either by Banxico or by analysts, and, using a text corpus, remove the projection on the first k components of a principal components analysis (PCA) decomposition of its sentences.[3] This method can outperform complex models in producing expressive sentence vectors for some tasks and is suitable for in-domain adaptation, in this case, the domain of bank analysis of Central Bank minutes. Also, it is simple and inexpensive to compute. We use $k = 5$ as we found it to obtain expressive sentence vectors, but our results are not very sensitive to the value of k. These sentence vectors capture well the semantics present in a sentence in the Central Banking domain, grouping together sentences that address the same topics. For example the nearest neighbors of "Se estima un crecimiento inercial de los ingresos fiscales con la actividad economica del pais", a sentence that was written by one of the banks in an undisclosed date, are in figure 2.[4] Readers of the Spanish language can immediately notice that ignoring on purpose word ordering information also misses the difference between good and bad outcomes regarding, in this case, economic growth and fiscal revenue, but can also immediately notice that, no matter the bank's sentiment about a topic, they are all the same topics. This is what we were looking for and this is what we found. We found topics, independent of their sentiment, and without human intervention. Notice that coincidentally, the "farthest" sentences to the original, at the end of the table, are opinions by the banks on what authorities in Mexico (namely Banxico and the Ministry of Finance) are expected to do given the information provided. By fitting word embeddings with text relative to central banking and the Mexican economy, and fitting sentence embeddings using bank's reviews and Banxico's minutes, we obtained meaningful semantic representations that accurately capture all subjects ever mentioned in any review or minute. Sentences are well distributed along their main variation dimensions, making easy to characterize the components of the narrative of both Banxico and analysts. As will be shown below, these components seem to group not only topics in an economic or monetary sense, but also the

[3] See Arora, Liang, and Ma (2016) for details on the creation of the vectors.

[4] The translation for this sentence is "It is estimated an inertial growth of fiscal revenue with the economic activity of the country"

ways analysts and the Central Bank refer to them.

Now we dedicate a couple of paragraphs discussing technical considerations regarding the removal of the Principal Components. A common and natural question for any reader who is not familiar with these techniques is how is it possible that the algorithm does such a great job grouping topics that are covered by a Central Bank by being trained with text that very infrequently uses these words by just removing the PC of the "financial" texts. On the one hand, we have the issue that the writer of the sentence is expecting that a human with some knowledge of the topic is reading the sentence. The raw vectors would have made, for example, the concept "inflation", or "price index", to weigh a lot more in the similarity measures, and thus would have made more difficult to discriminate among the secondary topics that appear together with it, much in the same way one could expect that describing a function using its derivative can complicate what "positive" and "increasing" mean. On the other hand, we have the frequency as a deterrent. Sentences in a random Wikipedia article that discuss inflation and growth would have been lumped closer with sentences that discuss inflation and Europe. By removing the Principal Components we address the issues in both hands. On the one hand, we force the computer to magnify the differences in context implied by small differences in the frequency of word pairs, and on the other, we force the computer to put those sentences in different classifications (one being related to a monetary policy stance for the future and the other one being related to a description of past or current events), which is exactly what we want.

Once the sentence vectors have been obtained, it is easy to explore the relevant topics and analyst's ways to address them. We use the t-SNE algorithm for visualizing the sentences.[5] t-SNE is a nonlinear algorithm that maps high dimensional data into a low dimension space (usually 2 or 3 dimensions), such that near points in the original space are mapped to near

[5]See Maaten and Hinton (2008) for details.

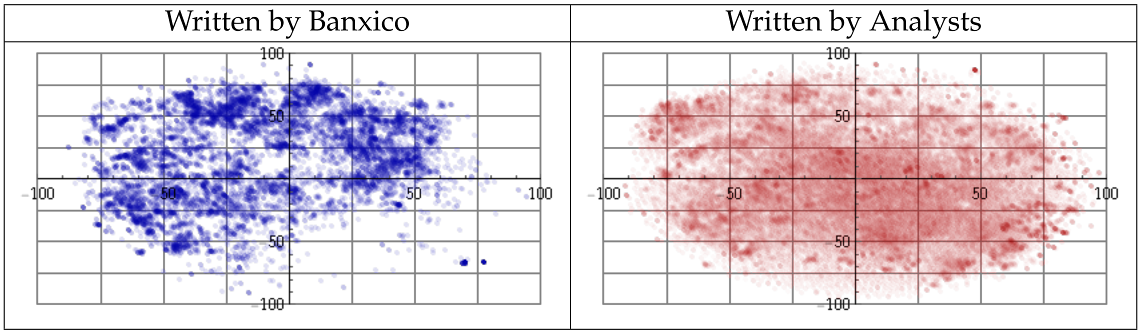

Figure 3: Two-dimensional representation of the relative frequency of all the subjects of all the sentences written by Banco de México and analysts that write to clients with interests in Mexico, 2011-2019

points in the lower dimensional space with high probability (see figure 3 for a 2-dimensional representation of all the sentences written by Banxico or analysts during 9 years, between 2011 and 2019).

This low dimensional representation is useful to interpret the main topics mentioned by analysts and to gain insight into how they address them, possibly enabling us to later make informed choices for supervision among a lower set of dimensions, not just "hawkishness" or "dovishness". For example, it would be useful that a computer learns to classify a sentence as whether the analysts confirmed their previous expectation on a certain topic or if their beliefs were changed, and also, to quickly build a training set aiming for the desired supervision. Later on in the paper we plot the output of the t-SNE for the sentences written by Banxico and by analysts in many different ways, and one such example is in figure 3.

In figure 3 it is possible to see a 2-dimensional projection of all sentences written by Banxico and by analysts between 2011 and 2019. Recall that each dot is a sentence in a 300-dimension vector space but projected in 2 dimensions using the t-SNE method: a sentence is the average of its word vectors, then subtracted its 5 principal components of "financial" sentences.

If the entire 2-dimensional plane was representative of all text in Spanish subtracted its first 5 principal components, then it is visible that analysts discuss just about every subject contained in reviews or minutes, and that Banxico does not discuss certain subjects. More importantly, both for Banxico and the analysts, the topics seem to be clustered, that is, the frequency of some topics is much larger than others. Interestingly enough, Banxico has its own clusters, and the analysts have their own clusters.

In this lower dimensional space, it is easier to group together sentences according to their meaning. We do that by grouping all the sentences of all the banks from 2011 to 2019 into 100 clusters using K-means. Here, a cluster does not necessarily mean a well defined topic but only a grouping of similar meaning sentences; nearby clusters could be attributed to the same topic, while sentences in a cluster could be grouped in different topics, depending on the reviewer. Once we have computed a cluster membership for every sentence, we use the probability distribution of sentences across topics to model the effects of the Central Bank communication by means of the persistence and emphasis on topics and tones in the general narrative of banks and in the Central Bank's minutes.

5 Results: Changes over time in Banxico's and in the Analysts' text

In this section, we do two things. First, we study how does the frequency of each topic (cluster) change over time. Second, we study how does the change depend on what was the nature of the text.

5.1 Change in the Frequency of Topics

We decompose the elements of all the communication done by Banxico and analysts, by date and by bank. That is, we will date each dot of figure 3, attach it to the bank that wrote it, and see if we can find some patterns in the changes of the location of the dots. As it can be seen in figure 4 on the following page, Banxico is very consistent with the sentences they *don't* write. On the contrary, it seems that analysts still write just about everything, but the emphasis, or

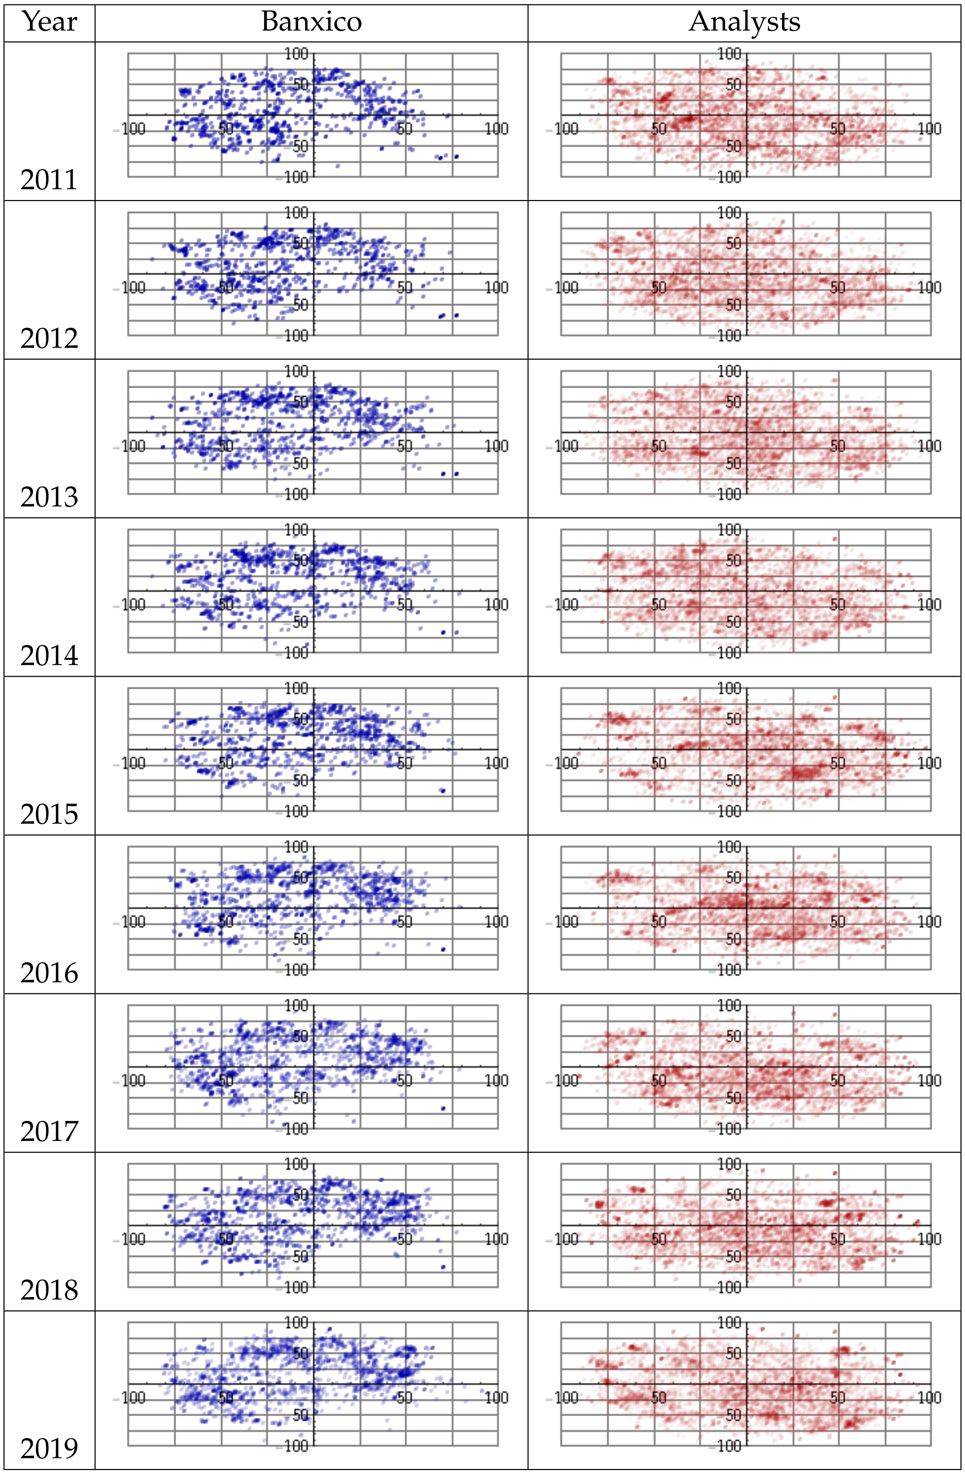

Figure 4: Two-dimensional representation of the relative frequency of all the subjects of all the sentences written by Banco de México and analysts that write to clients with interests in Mexico.

the clusters, vary by year. For example, take 2011, 2015, and 2018 as well as clusters 7, 9, 51, 68, and 96, pictured in figure 5.

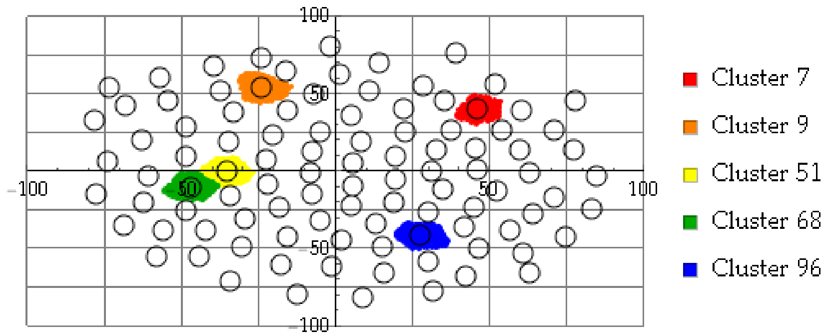

Figure 5: Location of all the sentences in clusters 7, 9, 51, 68, and 96. The disks represent the centroid of each of the 100 clusters found in the data.

In 2011, both Banxico and the analysts used a lot of sentences from cluster 51 (see Appendix to find the location of all the clusters). This cluster contains sentences that discuss gasoline prices, oil prices, and in some cases even food prices, always in the context of the exchange rate and the Mexican inflation. The reader will remember that 2011 was the year of the Middle East crisis, which spiked oil prices and reduced transportation routes all over the world. In 2015, the analysts used a lot of sentences from cluster 96. This cluster has explicit opinions about what the Bank of Mexico will do with respect to the Fed ("Taking into consideration the last communications by Banxico, we reiterate our opinion that Banxico will try to wait for the Fed and raise at the same time", "If the break-evens remain high, then Banxico can raise the rate in September, having also more information about the Fed", and many other examples, also listed in the Appendix).[6] Our algorithm found that Banxico did not write any sentence in this cluster in the whole 9-year period. This was, of course, the year of the *liftoff* by the Fed, which happened in December. In 2018, again Banxico and the analysts discuss heavily on one cluster, in this case, cluster 7. This cluster includes the sentences about negotiations, exchange rate, trade, tariffs, Canada, United States, and North America. This was

[6]Free translation done by the authors to "teniendo en cuenta las últimas comunicaciones de Banxico, reiteramos nuestra opinión de que Banxico tratará de esperar a la Fed y subir al mismo tiempo" and "si los break-evens se mantienen altos, Banxico podrÃa subir tasas en septiembre, teniendo también más información sobre la Fed", respectively

the year of the final rounds of the new North American Free Trade Agreement.

On the other hand, every year, Banxico writes a lot of sentences from clusters 9 and 68. The former contains phrases such as "upward revision", "sustained growth", and similar concepts, while the latter contains phrases such as "increasing prices", "rising inflation", and similar. Meanwhile, the analysts do write sentences from that cluster, they do it much less frequently. Notice how cluster 51 and cluster 68 are similar in the sense that they discuss prices and inflation, but cluster 51 discusses the specific goods (price of oil, mostly), and cluster 68 talks about prices in general. It is exactly this type of classification we were looking for when doing the clusters, and the algorithm seems to do it very well. We allowed for 100 clusters to be chosen arbitrarily, but we believe that with some human intervention, the results can even become stronger, which will be a new dimension added to our further research. For example, it is interesting to review the clusters found by the model as they segment topics and the ways of addressing them along the directions of most variance in the minutes and reviews. Most of the topics are well defined but sometimes in unexpected ways, for example, there is a topic about punctual forecasts about inflation and one about forecasts on inflation but with a strong argumentative narrative; there is a topic about the components of economic growth and one that reviews the Central Bank's perspective on these same components of growth. The granularity at which topics are defined depends on the number of clusters selected and on the number of principal components removed from sentences vectors.

5.2 Change in Analysts' topics after Banxico communicated something

In this subsection, after seeing that the algorithm did a good job classyfying sentences given the overall context of a specific year, we study if it does a good job by observing the change in the frequency in short periods of time. According to standard literature, we will name the

Cluster (topic)	Frequency in 2 & 4	Frequency in 2, 3 & 4	% change in frequency	Topic Interpretation
19	24	24	0.00%	Inflation expectations
25	21	21	0.00%	Foreign outlook: developed countries
31	24	24	0.00%	Real economy: private consumption, industrial, investment
35	30	30	0.00%	Output gap, growth and slackness
42	22	22	0.00%	Drivers of financial flows into Mexico
49	29	28	-3.40%	Risk balance on inflation: references to CB communication
20	26	25	-3.80%	Interpretation about inflation expectations
10	24	23	-4.20%	Punctual CB's forecasts for inflation
21	36	34	-5.60%	Labor market and its slackness
32	23	21	-8.70%	Growth forecasts: national and international
27	22	20	-9.10%	References to MPC and CB
0	33	29	-12.10%	Inflation forecasts: core inflation, gasoline and agricultural
23	31	27	-12.90%	CB views on growth risks
51	31	27	-12.90%	Effect of FX and inputs on inflation
48	21	18	-14.30%	Discussion about inflation drivers
40	25	21	-16.00%	inflation and rate forecasts: shocks and impact of expectations
12	24	20	-16.70%	Risks, forecasts and levels of indices: inflation
30	32	25	-21.90%	Risks, possible outcomes: growth, inflation, FX
46	27	21	-22.20%	On the votes for the MPD
47	25	19	-24.00%	Drivers of voting in the MPD
6	41	31	-24.40%	MPC and its actions on inflation
14	22	16	-27.30%	Interpretations of future CB actions
37	25	18	-28.00%	Components of growth
53	37	26	-29.70%	Arguments on forecasts
15	33	23	-30.30%	References to CB communication
28	37	24	-35.10%	Interpretations of forecasts on MPD
5	24	15	-37.50%	FX, inflation and volatility: references to CB communication
24	31	17	-45.20%	CB actions on FX and financial markets
13	27	14	-48.10%	Forecasts, opinions and interpretations
3	29	13	-55.20%	Risk balances on inflation: short term
2	25	9	-64.00%	Future decisions of MPC; tone of its communication
26	41	14	-65.90%	Analysts opinions
17	31	8	-74.20%	FX levels
45	24	3	-87.50%	Forecasts, opinions: rate changes
33	24	2	-91.70%	About the tone of CB communication
41	27	1	-96.30%	Analysts opinions on international outlook
8	32	0	-100.00%	Announcements about MMPD
39	30	0	-100.00%	Rate forecasts

Table 1: Frequency of each topic (cluster) in the text written by analysts each time there was a monetary policy decision in bank's reviews (moments 2 & 4), and in both bank's reviews and the MMPD (moments 2, 3 & 4).

measure of the changes in the distribution of the topics as a surprise, in the sense that it was unnecessary to write about that subject before, but then something happened and now it is necesary.

We measure the surprise between distributions of sentences in 4 different moments in time. For a particular MPD we consider four moments, which in temporal order are:

1. Analyst reviews before the MPD
2. Analyst reviews after the MPD
3. Minutes of the monetary policy decision meeting (MMPD)
4. Analyst reviews after the release of the MMPD

In this setting, we can isolate the effect of the MPD announcement in the analysts narrative by looking at the differences between reviews in moments 1 and 2. Similarly, we isolate the effect of the MMPD by the differences between moments 3 and 4.

In table 1 we show our interpretation of some topics, together with their mean change of probability (in percentage) before and after the MMPD (column 2), and before and after the MPD (column 3). Column 4 shows the number of times that a topic was mentioned by at least two banks both after the monetary policy decision and after the MMPD, and column 5 shows the number of times Banxico *also* wrote a sentence in this topic.

Note that the CB minutes don't make forecasts about future rate changes (topic 39), neither they issue opinions on these forecasts (topic 45). On the other hand, changes in comments about the tone of Banxico's communication (topic 33) are more likely to be induced by the MMPD than by the MPD. Analogously, comments about the drivers of MPD voting (topic 47) are induced by the MMPD rather than by the MPD.

But we can go much further. We want to obtain correlations. In each report by each analyst at any date, cluster membership for each sentence is distributed $Categorical(p_i)$ where p is a vector with the cluster appearance probabilities across the 100 topics and i indexes every date and each of the four report types. Sentences in each date-report type are distributed $Multinomial(n, p_i)$. We use the Hellinger distance to quantify changes in these distributions and we call it a measure of surprise. If the observed distribution of sentences in a given date-report is very different from the expected distribution it should have a higher surprise. Hellinger distance between two discrete distributions P and Q is given by

$$H(P,Q) = \sqrt{\frac{1}{2}\sum_{i}(\sqrt{p_i} - \sqrt{q_i})^2} \tag{4}$$

Here, p_i and q_i are probabilities for cluster i of each of the two distributions P and Q. Hellinger distance is a metric bounded by 1 and straightforward to understand; it is bigger if the probability of each of the topics changes more. Even though there are many ways for measuring similarity between probability distributions, we chose this one for its simplicity, and because it is symmetric and satisfies the triangle inequality. Also, it has ideal qualities that we are looking for when we discuss a change of subject on a text. For example, if a subject that is largely discussed in the past (measure P) is no longer largely discussed in the present (measure Q), the Hellinger distance will show a larger change if only one subject that was not discussed in the past picks up the relative frequency than a subject that already was largely discussed. Transitively, changes are smaller if several topics each receive a fraction of the relative frequency, but a change is considered larger the more subjects that were not discussed at all are discussed now.

Given what topics the analysts address and the characteristics of how they address them, there are topics that don't appear in some of the reports or in the Central Banks' minutes. For example, topics about the analyst's opinion on the Central Bank's views on growth, or analyst's forecasts for Mexican Government bonds will never appear in the minutes. When comparing distribution of topics covered in more than one type of reviews we simply don't consider the topics that usually don't appear in any of the implied reports.

We found that the surprise induced by the MMPD on analyst reviews is (almost always) greatly reduced if the initial analyst distribution is updated with the sentences in the minute, this suggests that the emphasis on topics covered by the minutes is absorbed by the analysts narrative. As a reference, we compute the surprises between analyst reviews and the uniform distribution to test whether the concentration of topics increases after the MPD and after the MPDD. Proximity to the uniform distribution (in which entropy is maximal) is interpreted as a low topic concentration, which suggests whether analysts don't agree in which topics are a priority for the CB and might have driven the MPD or that they write just about every topic. We found the concentration of topics to increase after the MMPD 58% of times, whereas the concentration of topics decreases 96% of the times after the MPD.

Another natural dimension would be to study topics as random variables, and, given that the mean does not represent anything, but the relative frequencies do, measure the variances of the distributions over time. In figure 6 it is possible to see that that 83 percent of times, the variance in coverage among analysts decreases after the release of the MMPD. It is noticeable that the level of the variance is correlated with the level of the monetary policy rate, that the drop in the variance is correlated with the fact that the monetary policy rate changed (most noticeably when it went down in 2014), and that the slope of the yield curve (measured by the difference between the 10-year bond rate and the monetary policy rate) is correlated with the level of the variance of the analysts' statements before the MMPD but not so much after.

Notice, for example, the large drop in the variance of the sentences written by the banks in February 2015. The entire set of sentences is pictured in figure 7. Notice that before the MMPD, the banks basically wrote about every topic across the board, with no emphasis on any particular subject. But after the publication of the MMPD, there is clear dominance of three subjects: The United States, The Federal Reserve Board, and the possibility of large

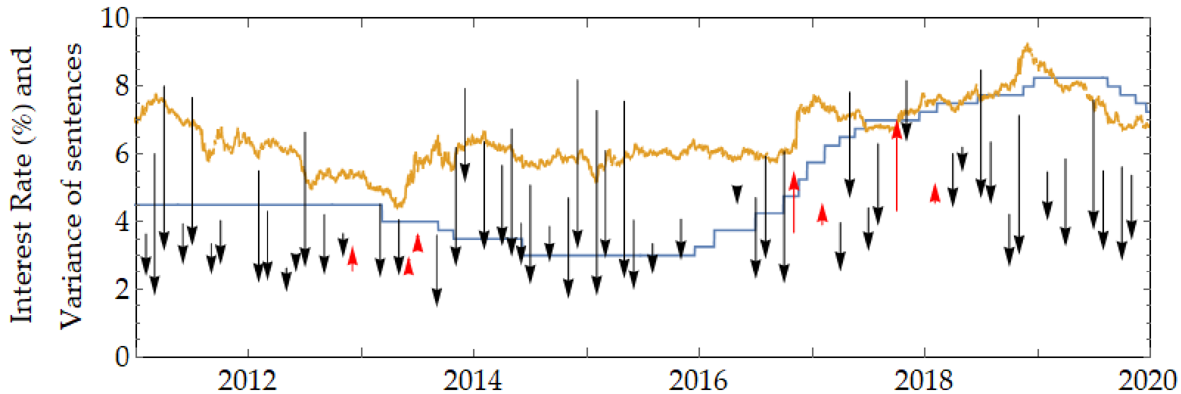

Figure 6: Change in variance of the analysts' statements before and after the MMPD, 2011-2019. Blue Line: target monetary policy rate, Yellow Line: 10-year bond yield. The red arrows are the only occassions in 9 years that variance increased after the MMPD.

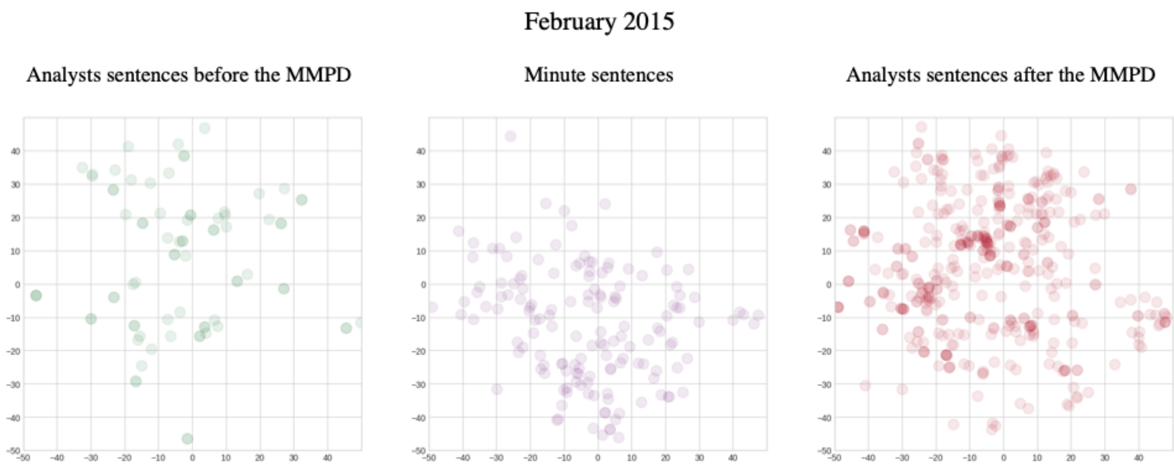

Figure 7: Sentences covered by the analysts before and after the publication of the Minutes of the Monetary Policy Decision in February, 2015

fluctuations in the exchange rate. This pattern was repeated during all of 2015. The U.S. dollar kept appreciating everywhere, the banks discussing just about every topic before the MMPD and then discussing the United States economy and the Federal Reserve Board as basically the only topics the day after the MMPD, as can be clearly seen in figure 8. Notice that this trend ended in December, when the *liftoff* happened. Notice that, comparing to 2014 and 2016, it is a similar effect, but in other topics (see figure 9 and figure 10).

	20-year Bond		10-year Bond		USDMXN Abs Chg		USDMXN Chg	
	(1)†	(2)	(3)†	(4)	(5)†	(6)	(7)†	(8)
MMPD Surprise	0.321^{**}	0.265^{**}	0.248^{*}	0.190^{*}				
	(0.100)	(0.089)	(0.097)	(0.084)				
US Yield Curve Chg	0.063	0.098	0.219^{*}	0.1968^{*}	0.009	0.017	0.000	0.000
	(0.112)	(0.102)	(0.108	(0.096)	(0.012)	(0.012)	(0.021)	(0.005)
MPD Surprise					-0.014^{*}	-0.012^{*}	0.023^{*}	0.020^{+}
					(0.006)	(0.006)	(0.011)	(0.010)
Rate Spread Before					−0.003	−0.001	0.004	0.000
					(0.004)	(0.003)	(0.006)	(0.005)
Rate Spread After					0.000	0.000	−0.003	−0.001
					(0.003)	(0.003)	(0.006)	(0.004)
Intercept	−0.108	−0.074	−0.068	−0.042	0.022	0.015	−0.018	−0.007
Observations	63	63	63	63	64	64	64	64
R^2	0.329	0.147	0.297	0.143	0.312	0.180	0.235	0.080

Standard errors in parenthesis.
† Marked models include year fixed effects.
$^{+}$ $p < 0.1$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 2: Regression models linking surprise in analysts narrative induced by the MMPD with absolute change in government bond yields and surprise in analysts narrative induced by the MPD with percent change in USDMXN.

6 Results: Correlations with financial variables

In this section, we not only find correlations between topics, but find if the changes in the distribution of the topics and the changes in financial variables are correlated. The main results are in table 2. There, we regressed the measure of the surprise in the analysts' texts with changes observed the day of the publication of the relevant information by Banxico. For example, in the first column, we regressed the change in the Mexican 20-year bond yield the day of the publication of the minutes of the monetary policy decision to the measure of the change in the frequency of topics that analysts wrote about between the text associated to the publication of the monetary policy decision and the text associated to the publication of the minutes of the monetary policy decision. In theory, the only difference should be that now the banks *know for sure* the reasons of the decision and before that they could only speculate. So, the idea is to correlate the movement in the yield can be associated to the change in that information, which in this paper is measured with the change in the frequency of topics. In that column, we controlled for the year (which in the last subsection was found to contain a lot of the variation in the topics) and still found that the larger the surprise, the larger the increase in the 20-year bond rate. Column 3 does the same for the 10-year yield, and we found the same pattern. In columns 5-8 we also found the measure of the change in the topics the day before and the day after the monetary policy decision (which was also studied in figure 6), and found that, controlling for foreign and domestic yield curves, a large surprise does not induce a large movement in the exchange rate.

7 Conclusion and further research

After a Central Banks communication, namely monetary policy statement, its corresponding minutes, or the Quarterly Inflation Report, we could expect to see some reaction from financial variables and economic analysts from financial institutions. This would be the case if the analysis and messages displayed by the Central Bank add some information to the markets. After studying the statements of economic analysts before and after the minutes for the monetary policy decisions, we observe more agreement and congruence around the topics concerning the Central Bank. In other words, the explanation of the decision, the outlook of the economy, and the evaluation regarding the forecasts provided by the Central Bank in the minutes have an impact and a great influence in the analysts' own outlook, at least in the topics they consider relevant to write their clients about. This could mean that the Central Bank's messages make sense to the analysts and consider that they are a good source of information regarding the explanation of the current monetary policy decision and future policy actions, or at least they decide to echo through the same topics. It would be unfortunate for the Central Bank if after a communication, analysts do not pay any attention to the authorities' messages and the dispersion of the subjects covered by them remain the same as before the communication. Our study does not say anything about the agreement (or disagreement) of the economists regarding the Central Banks stance or outlook for the economy. However, the content in the analysts' reports constitutes their opinion of the outlook and the expectation of what the Central Bank will do regarding monetary policy in the future. Those reports are in general an important input for their clients. Based on their analysis and the interpretation of the Central Bank communications, analysts make recommendations of portfolio positions in the fixed income market, FX market, and equity market to their clients. Those clients are local and foreign investors, traders, and fund managers, among others. Firms also take these reports as an input for investment and overall planning decisions. Therefore, even a simply echo of the Central Banks outlook means that this information displayed by the Central Bank will have impact in the financial markets through this channel.

In future steps of this research, we plan to analyse the qualitative content of the monetary policy statements and its evolution over time to revise how this information is interpreted by the financial markets. As other research, we will construct a sentiment Index from Banco de México's communiques (hawkish versus dovish) and evaluate if such index have a relationship with the financial markets main variables. In addition, we can estimate if it includes information about the future path of the monetary policy rate in addition to the actual monetary policy decision. We then could study how fixed income market (e.g. level and slope of the yield curve) respond to the monetary policy statements and confirm if those communications instruments help to guide expectations regarding future monetary policy. Finally, we could study if such an index has an impact only on short-term rates or also in long-term ones.

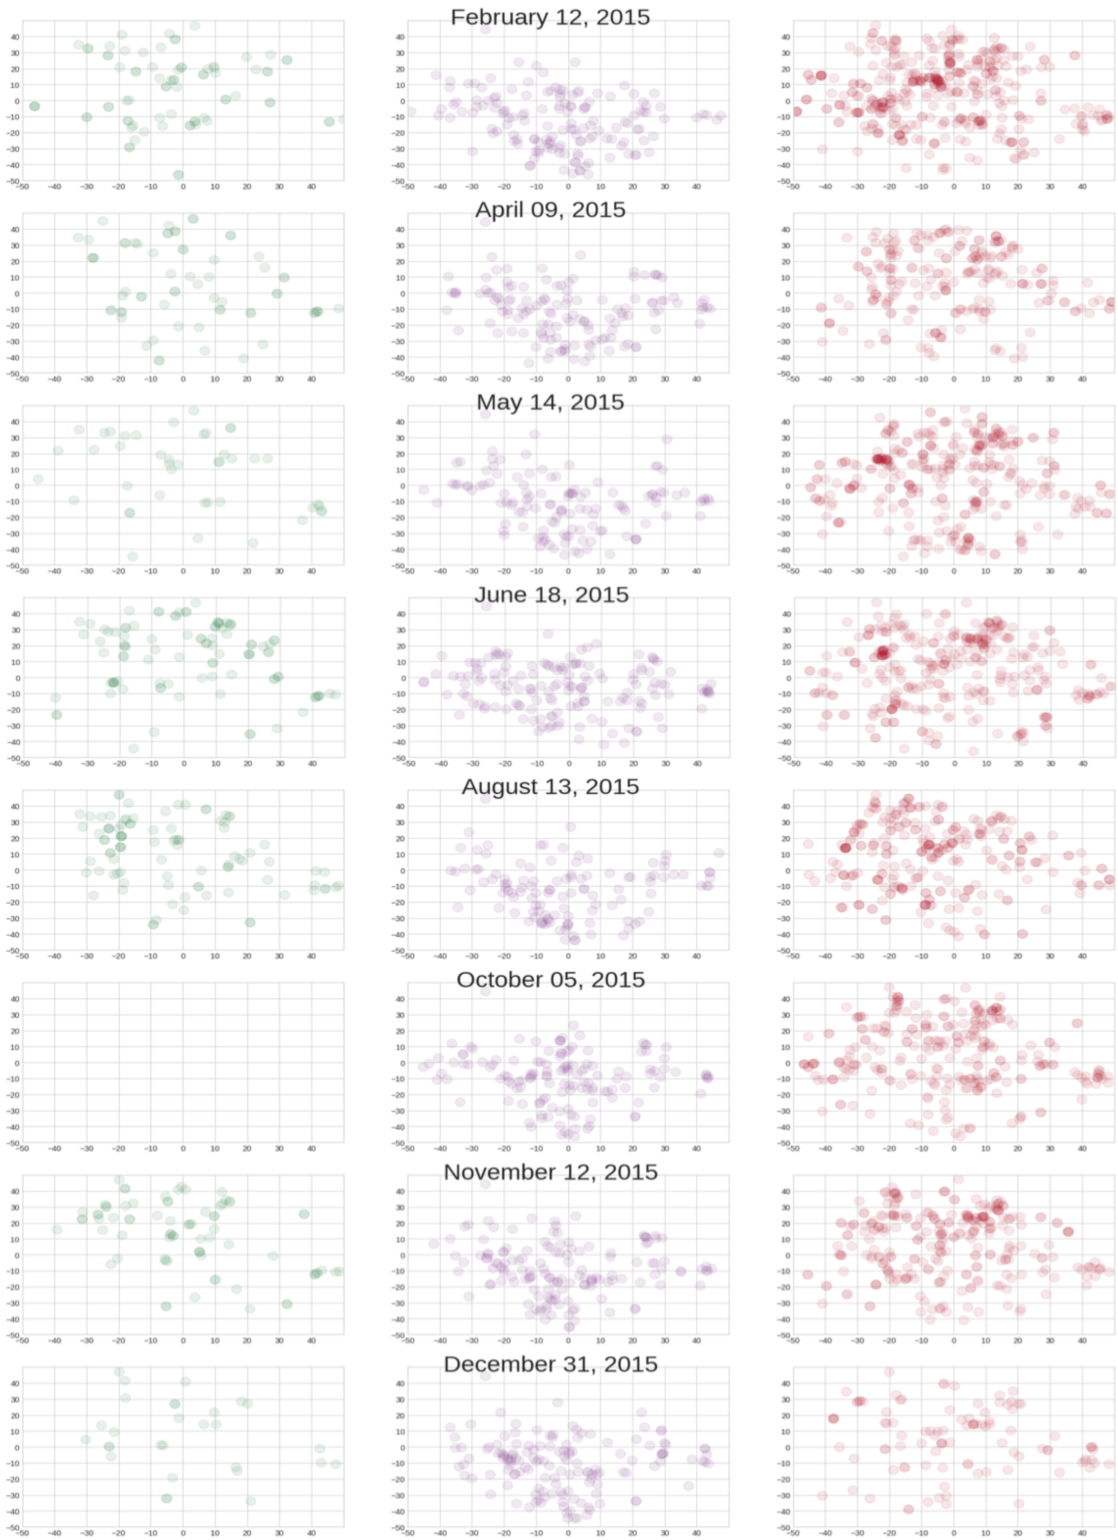

Note: left is before the MMPD, center is the minute sentences, and right is after.
Figure 8: The huge dominance of the Fed and the United States economy during 2015 is only seen after the MMPD, not before.

Note: left is before the MMPD, center is the minute sentences, and right is after.

Figure 9: Data for 2014

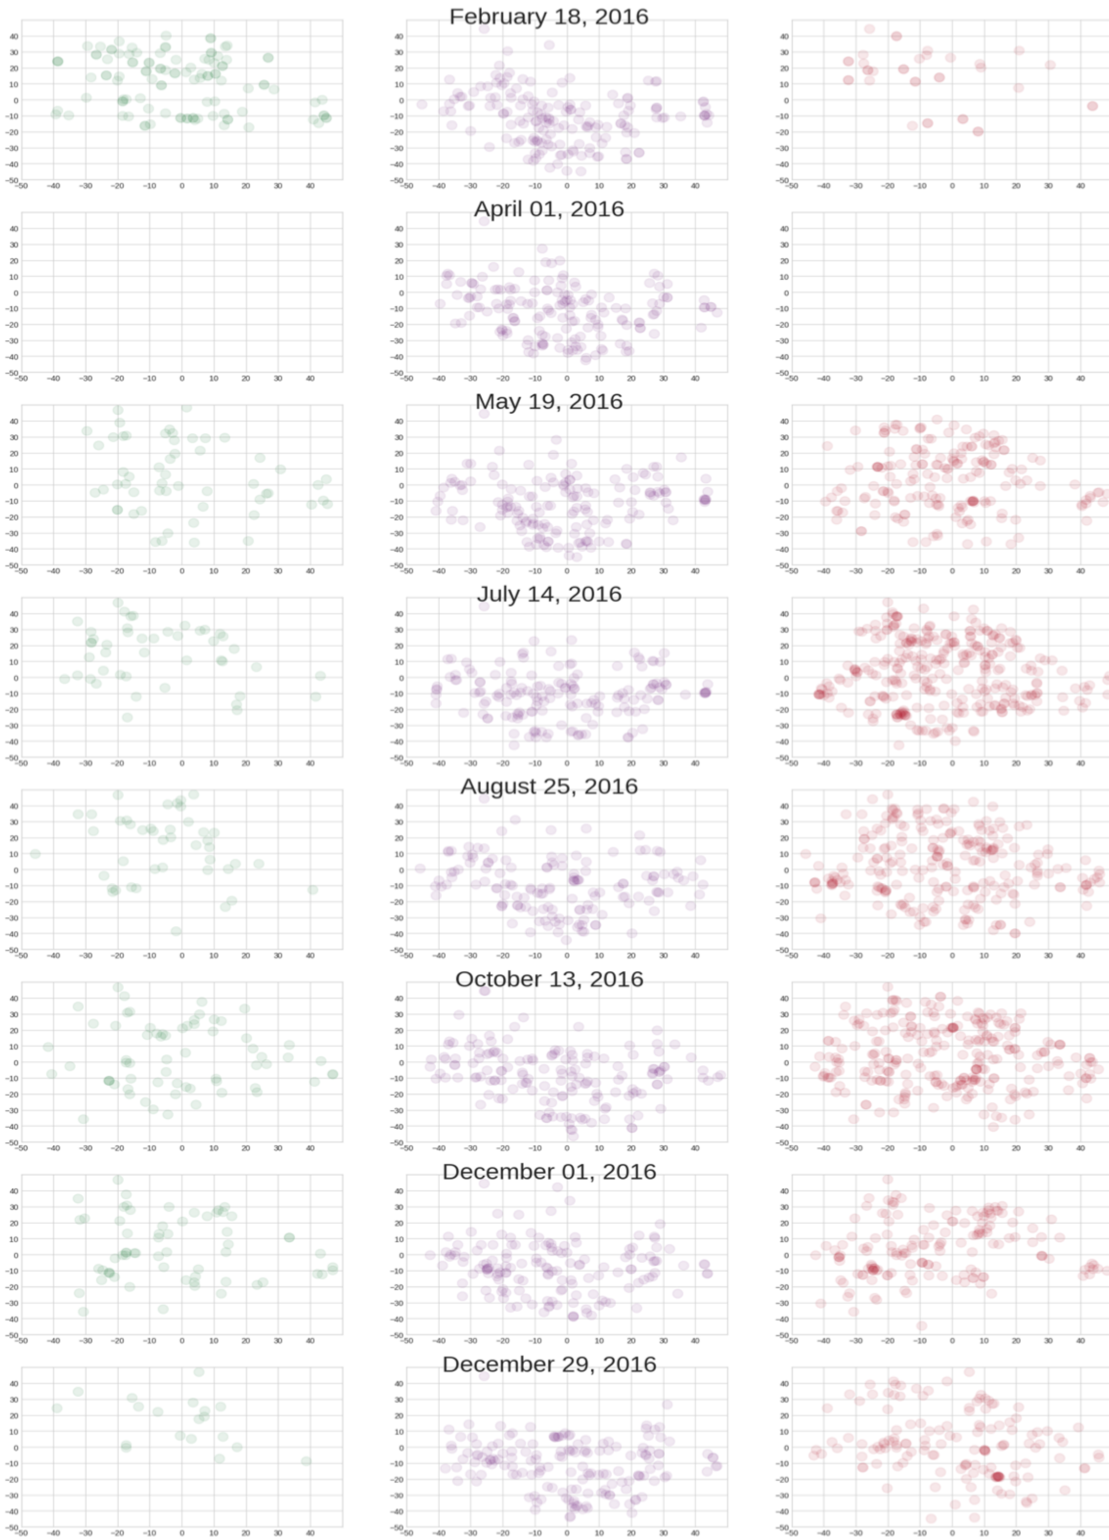

Note: left is before the MMPD, center is the minute sentences, and right is after.

Figure 10: Data for 2016

References

ARORA, S., Y. LIANG, AND T. MA (2016): "A simple but tough-to-beat baseline for sentence embeddings," *Princeton University MIMEO*.

BHOLAT, D., S. HANSEN, P. SANTOS, AND C. SCHONHARDT-BAILEY (2015): "Text mining for central banks," *Available at SSRN 2624811*.

BLEI, D. M., A. Y. NG, AND M. I. JORDAN (2003): "Latent dirichlet allocation," *Journal of machine Learning research*, 3(Jan), 993–1022.

BLINDER, A. S., M. EHRMANN, M. FRATZSCHER, J. DE HAAN, AND D.-J. JANSEN (2008): "Central bank communication and monetary policy: A survey of theory and evidence," *Journal of Economic Literature*, 46(4), 910–45.

DAS, S., S. GHOSH, S. BHATTACHARYA, R. VARMA, AND D. BHANDARI (2019): "Critical Dimension of Word2Vec," in *2019 2nd International Conference on Innovations in Electronics, Signal Processing and Communication (IESC)*, pp. 202–206. IEEE.

FILARDO, A., AND D. GUINIGUNDO (2008): "Transparency and communication in monetary policy: a survey of Asian central banks," in *BSP-BIS High-Level Conference on Transparency and Communication in Monetary Policy, Manila*, vol. 1.

GOLDBERG, Y., AND O. LEVY (2014): "word2vec Explained: deriving Mikolov et al.'s negative-sampling word-embedding method," *arXiv preprint arXiv:1402.3722*.

KAHVECI, E., AND A. ODABAŞ (2016): "Central banks' communication strategy and content analysis of monetary policy statements: The case of Fed, ECB and CBRT," *Procedia-Social and Behavioral Sciences*, 235, 618–629.

MAATEN, L. V. D., AND G. HINTON (2008): "Visualizing data using t-SNE," *Journal of machine learning research*, 9(Nov), 2579–2605.

MIKOLOV, T., I. SUTSKEVER, K. CHEN, G. S. CORRADO, AND J. DEAN (2013): "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, pp. 3111–3119.

PATEL, K., AND P. BHATTACHARYYA (2017): "Towards lower bounds on number of dimensions for word embeddings," in *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 31–36.

PENNINGTON, J., R. SOCHER, AND C. MANNING (2014): "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543.

WIETING, J., M. BANSAL, K. GIMPEL, AND K. LIVESCU (2015): "Towards universal paraphrastic sentence embeddings," *arXiv preprint arXiv:1511.08198*.

Appendix