WP/21/159

IMF Working Paper

How Have IMF Priorities Evolved? A Text Mining Approach

by Gareth Anderson, Paolo Galang,
Andrea Gamba, Leandro Medina, and Tianxiao Zheng

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Strategy, Policy and Review Department

How Have IMF Priorities Evolved? A Text Mining Approach

**Prepared by Gareth Anderson, Paolo Galang, Andrea Gamba, Leandro Medina, and
Tianxiao Zheng¹**

Authorized for distribution by Ana Corbacho

June 2021

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

This paper assesses how priorities of the IMF's membership have evolved over the past two decades, by using text mining techniques on a unique dataset combining IMFC communiqués and constituency statements. Our results reveal significant variation in priorities across time and constituencies. Statements can be characterized by the weight which they place on three key priorities: (i) growth; (ii) debt and development; and (iii) crisis management and quota reform. Sentiment analysis techniques also show that addressing climate change is a topic which is viewed positively by an increasing number of constituencies.

JEL Classification Numbers: O17, O52, H26

Keywords: sentiment analysis; text mining; text analysis; IMFC, communiqué, statements

Author's E-Mail Address: ganderson2@imf.org; agamba@imf.org; jgalang@imf.org; lmedina@imf.org; tzheng@imf.org

¹ Authors wish to thank Clifford Blair, Chengyu Huang, Deniz Igan, Giulio Lisi, Robert Loo, Jorge Martinez, Natalia Stetsenko, Justin Zake, as well as participants of the SPR seminar, for their useful comments and suggestions. Special thanks to Ana Corbacho and Johan Mathisen for their guidance and support.

Contents

Abstract	2
I. Introduction	4
II. Relation to the Literature.....	4
III. The Data: IMFC Communiqués and Constituency Statements	5
IV. IMFC Communiqués and Member Priorities: Stylized Facts.....	8
V. IMFC Communiqués and Member priorities: A Clustering Analysis Approach	11
Constituency Statements with Topic Orientations.....	11
Constituency Clusters Based on Topic Distributions	15
VI. Emerging Topics.....	17
VII. Conclusion.....	23
References	24
I. APPENDIX	26

I. INTRODUCTION

The International Monetary and Financial Committee (IMFC) is a policy advisory committee to the IMF's Board of Governors, serving as a bridge between the Board of Governors and the IMF's resident Executive Board. It meets twice a year at the time of the IMF-World Bank Spring and Annual Meetings to provide strategic direction to the work of the Fund. After each meeting, it releases a communiqué that sets out the consensus position among members about the current economic outlook, the desired policy response, and the strategic priorities for the work of the IMF. Following the publication of the communiqué, the statements of individual constituencies are released publicly. These provide the viewpoints of the individual constituencies and are often revealing of the differences of opinion amongst the IMF's membership.

In this paper, we combine text mining with clustering techniques to analyze the IMFC communiqués and the IMFC constituency statements over the period 2000–2019. We first determine relevant topics by building a lexicon, and then analyze keyword frequencies over time and across constituencies. Specifically, we calculate the topic distribution of each statement, which captures the relative frequency of a topic normalized by the total number of all topic appearances. We focus on three questions: i) how have economic policy priorities of the IMF's governing body evolved over the past two decades? ii) what are the similarities and differences across individual constituencies and over time? iii) what emerging priorities are likely to be the focus of future policy, and how do members feel about them?

We find that statements of the IMF's membership can be distinguished by the weight which they attach to the priorities of growth, debt and development, and crisis-management. Broadly, the priorities of the membership have evolved through three phases: phase 1 in the early-2000s placed a large emphasis on debt and development issues; phase 2 in the late-2000s focused on the financial crisis and IMF quota reform; and phase 3 in the 2010s highlighted economic growth. We also detected notable differences in the weight attached to key priorities across constituencies, and over time within the same constituency. The United States has placed a larger weight on economic growth, China on reform, and Germany on debt issues. More recently, some constituencies have emphasized the importance of climate change, gender, and fintech, although this shift in focus is not yet universal across the membership.

II. RELATION TO THE LITERATURE

A number of studies have applied text mining techniques to assess the information content of news. One strand of literature focuses on the frequency of certain terms and topics within news articles, for example the occurrence of words relating to economic uncertainty (Baker, Bloom, and Davis (2013)). Other studies focus on the sentiment of news. For example, Fraiberger et al. (2018) construct a simple sentiment score based on the fraction of positive and negative words in each article and, using this, they estimate the response of equity prices

to sentiment shocks. UpLevel (2019) applies both of these approaches to international summits, focusing on news topics and sentiment in the lead-up to and following the 2019 G20 Osaka Summit.

In addition to analyzing news, a growing literature has also utilized text-mining techniques to assess the content of economic statements, although the focus has largely been on central bank communications. Boukus and Rosenberg (2006) apply latent semantic analysis (LSA) – a form of unsupervised machine learning – to characterize FOMC minutes according to their prevailing theme and find that this theme is correlated to current macroeconomic and financial conditions. Acosta (2015) applies LSA to study the effect of public disclosure of FOMC minutes on members’ comments and finds increased conformity after the publication of transcripts started under Alan Greenspan’s chairmanship. Relatedly, Hansen et al. (2018) use a latent Dirichlet allocation (LDA) topic modelling approach to assess how transparency reforms have affected the content of FOMC deliberations.

Text mining of international organizations’ documents, however, is a relatively recent development. Mihalyi and Mate (2018) analyze IMF Article IV reports to assess the nature of IMF surveillance and, similarly, IMF (2019) consider how the IMF’s social spending surveillance has evolved through a text mining analysis. Amaglobeli et al (2018) construct a database of tax policy measures building on both OECD surveys and tax-related news.

This paper is the first to explore the content of IMFC communiqués and constituency statements through various text mining techniques. Using documents spanning the period 2000-2019, we can exploit both the time and cross-sectional aspects to provide insights on the evolution of the IMF membership’s priorities. Moreover, at any point in time, we are able to interrogate the key differences in priorities across the membership. Finally, by comparing the statements of individual constituencies to those of the IMFC, we explore how the position of each constituency differs from the overall consensus position of the IMFC.

III. THE DATA: IMFC COMMUNIQUE’S AND CONSTITUENCY STATEMENTS

Our analysis draws on two important strands of documents: the IMFC communiqués and IMFC constituency statements. These documents are released semi-annually following the IMFC Meetings, which take place in April and October, at the World Bank and IMF Spring and Annual Meetings.

The communiqué is a document which is agreed upon by the IMFC’s 24 constituencies, and therefore reflects the consensus position about key macroeconomic themes and the strategic direction of the IMF’s work and policies. Following the release of communiqué, the 24 constituencies and key meeting observers (for example, the OECD and the ECB) publish their own statements. These statements represent the views and priorities of the individual constituencies at a given point in time and often reveal the true differences of opinion amongst the IMF’s membership.

We construct a panel dataset that contains the set of IMFC communiqués and constituency statements, hereafter referred to as the “corpus”, between the Spring Meetings 2000 and the Annual Meetings 2019, covering 24 constituencies (representing 189 members) and 13 observers.^{2,3} Our panel includes IMFC communiqué statements since 2000 and constituency statements since 2003. For constituencies with more than one-member country, the nationality of the Executive Director representing the constituency often changes over time. For this reason, our analysis focuses on statements at the constituency level, rather than at the member country level. Overall, we have compiled 1,147 documents, among which 40 are communiqués, 805 are constituency statements, and 342 are statements from observers.⁴

The descriptive statistics for the corpus are presented in Table 2 in the Appendix. The length of statements and communiqués varies across time but they comove closely, with lengthier documents before the 2008-2009 financial crisis compared with afterwards. The communiqué and constituency statements from the 2009 Annual Meetings were particularly long, reflecting increased discussion of the global crisis and the associated policy response (Figures 1 and 2).

Figure 1. Number of sentences per constituency statement.

Figure 2. Number of sentences per IMFC communiqué.

Building a lexicon

We use a dictionary-based approach to extract the main topics from the collection of statements. Beginning with the corpus of IMFC communiqués and constituency statements, we start by eliminating stop words, such as “and”, “the” and other words common in written

² The meeting observers include the European Commission, the European Central Bank, the European Union, the Financial Stability Board (Financial Stability Forum prior to 2008), the International Labour Organization, the IMF’s Managing Director, the Organisation for Economic Co-operation and Development, the United Nations, the United Nations Conference on Trade and Development, the World Bank, and the World Trade Organization.

³ Constituency membership is generally stable, although there are some small changes over time in their composition. For example, in 2012 Belgium, Luxembourg, and the countries of the constituency previously chaired by the Netherlands agreed to create a new constituency.

⁴ We note that there are 35 observations missing, which represents a small portion of the sample. Table 2 in Appendix highlights the number of observations by constituency.

English which do not identify any specific term of relevance for our analysis. Second, we collect bigrams and trigrams- collections of two or three words which are commonly used together, for example “carbon tax”. Following this procedure, we are left with over 11,000 unique words, bigrams, or trigrams. We observe that 1,679 of these have a direct reference to IMF strategic work. Given our corpus is relatively small, we manually map⁵ these 1,679 terms into 133 broader topics. We then build a Document Term Matrix (DTM), whose rows represent each term appearing in the corpus, and whose column represent each document analyzed. In practice, the DTM is a matrix of term frequencies in each document comprising the corpus.

The lexicon-building process, as widely accepted in the literature, involved some discretion. To minimize idiosyncrasies, we developed two independent mapping proposals and validated a reconciled correspondence table with experienced communiqué drafters⁶. An example of the mapping is shown in Figure 3 for the topic of climate change. In the Table 3 of the Appendix we present the mapping for three other prominent topics (growth, debt, and reform). We also present an alternative topic modelling approach in the Appendix that uses latent Dirichlet allocation (LDA), a method which does not rely on the manual assignment of terms to topics and instead uses a probabilistic approach.⁷ The results presented in this alternative approach are consistent with our main findings.

⁵ In the literature, text is often further pre-processed by applying (automated) lemmatization and stemming to reduce variation (HaCohen-Kerner et al. (2020)). Given our corpus is relatively small, we preferred to embed these procedures in the manual construction of the lexicon.

⁶ The iterations with drafters, combined with the manual selection of relevant bigrams and trigrams, were essential to properly classify words that could be used across different policy contexts into specific topics (e.g., “debt growth” was classified into “debt” while “employment growth” was classified into “employment”).

⁷ A further possible extension would be to consider dynamic topic models, whereby the nature of topics evolves over time (see, for example, Blei and Lafferty (2006) and Wang, Blei, and Heckerman (2008)).

Figure 3. Building a lexicon: an example for the topic of climate change

Term	Topic
bio-fuel	climate_change
biofuel_production	climate_change
carbon	climate_change
carbon_pricing	climate_change
carbon_tax	climate_change
climate_change	climate_change
climate_change_mitigation	climate_change
climate-related	climate_change
climatic	climate_change
conserve_energy	climate_change
cop	climate_change
drought	climate_change
emission	climate_change
energy_conservation	climate_change
energy_intensity	climate_change
energy-efficient	climate_change
energy-saving	climate_change
extreme_weather	climate_change
global_warming	climate_change
low-carbon	climate_change
niño	climate_change
tackling_climate_change	climate_change

Using the 133 topics in our lexicon, we count the number of topic appearances in each text comprising our corpus. Let n_{it}^j denote the number of appearances of topic j from the statement of constituency i at date t , then we can write the total number of all topic appearances from the statement of constituency i at date t as $N_{it} = \sum_j n_{it}^j$. We then define the topic distribution as $K_{it} \equiv [k_{it}^1, k_{it}^2, \dots, k_{it}^j \dots]'$, where $k_{it}^j = n_{it}^j / N_{it}$ is the fraction of attention devoted to topic j in the statement of constituency i at date t , with $\sum_j k_{it} = 1$.

The topic distribution K_{it} provides a useful way to represent the topic patterns present across the collection of statements. Each element of K_{it} represents the weight of a topic. In other words, it captures the relative frequency of a topic normalized by the total number of all topic appearances. Therefore, the topic distribution is useful in summarizing the frequency of topics and is the main object we use in our subsequent analysis.

IV. IMFC COMMUNIQUÉS AND MEMBER PRIORITIES: STYLIZED FACTS

Descriptive Statistics

The topic distributions of IMFC communiqués reveal that these documents tend to focus on a few selected topics. On average, each topic gets a weight of around one percent in each constituency statement—but a few select topics get a much higher weight, while several other topics receive a very low weight. Figure 4 shows the most popular ten topics based on

IMFC communiqués across the sample period. Growth is the most common topic, on average accounting for around ten percent of the topics occurring in constituency statements. This reflects the fact that maintaining sustainable and vigorous economic growth is consistently a top priority among members. The top ten list also includes reform, debt, risk, fiscal issues, and development, among others. The topic content of IMFC communiqués shows a significant overlap with those of the constituency statements. Figure 5 shows the most popular ten topics in constituency statements across the sample period. Notably, growth and reform are the most common topics in both IMFC communiqués and constituency statements.

Figure 4. Top-10 topics of IMFC communiqués (weighted frequency)

Figure 5. Top-10 topic of constituency statements (weighted frequency)

Looking at how topics have evolved over time in constituency statements, it is clear that growth has persistently been a top priority (Figure 6). However, in 2009, the focus shifted away from growth towards crisis management measures, following the onset of the global financial crisis. Debt, which was frequently mentioned in the context of debt relief initiatives in the early 2000s, has been less common following the financial crisis. Instead, the topic of reform has been prominent, in part reflecting discussions on IMF quota and governance reforms. The topic of climate change has only begun to gain prominence very recently.

Figure 6. Topic evolution over time (weighted frequency), averaged over all constituencies

Although a small number of topics tend to get a very high weight across many constituencies, there is often substantial differences across constituencies in the weights they place on these topics. Figure 7 plots the standard deviation across constituencies of the weight placed on key topics over time. This summarizes relative differences across all constituencies on a certain topic over time, with a lower standard deviation reflecting greater similarity across constituencies at that time. Constituencies consistently differ the most on the weight they attach to the topic of growth, with other topics having much lower variation. The standard deviation for growth drops sharply during the global financial crisis, likely reflecting a unified shift across the membership towards focusing on crisis policies. Following the crisis, the standard deviation for growth increases, potentially reflecting idiosyncratic domestic economic conditions and differing priorities across the membership.

Figure 7. Cross constituency variation of top 4 topics. (Standard Deviation)

Focusing on the five largest constituencies by quota share, we also find significant variation in the frequency of key topics (Figure 8). Statements from United States (US) and China (CC) on average place more emphasis on growth and reform than the other large constituencies. Germany (GR) places a higher weight on debt, while the NE constituency, which includes the Netherlands, places a relatively high weight on risk.

Figure 8. Frequency of top 4 topics for the five largest shareholders

V. IMFC COMMUNIQUÉS AND MEMBER PRIORITIES: A CLUSTERING ANALYSIS APPROACH

In this section, we use principal components and clustering analysis techniques to more formally evaluate the priorities of the IMFC membership and how they have shifted over time.

Constituency Statements with Topic Orientations

Principal components analysis provides a useful way of reducing the dimensionality of our large dataset and assessing the importance of each topic. Using our dataset consisting of the topic distributions of all IMFC communiqués and constituency statements over time, we project the data onto the first two principal components. The loading plot, presented in Figure 9, shows how strongly each topic influences the first two principal components. The projected values of the vectors show the weight each topic has on the principal components and the angles between the vectors provide an indication of how the topics correlate with each other. For each principal component, we present the five topics that have the largest absolute weight.

The topic of growth is particularly important in driving the variation across statements. Figure 9 shows that the topic “growth” has the strongest influence on the first principal component and “crisis” and “surveillance” have a moderate influence. For the second principal component, “debt”, “development”, and “reform” have a moderate influence.

Figure 9. Principal component loadings

Clustering analysis groups together statements which have similar topic distributions, allowing us to broadly characterize the nature of different statements. The K-means clustering algorithm is one popular clustering technique and proceeds as follows. First, the number of clusters, k , is chosen and for each of the k clusters, a centroid is also defined. Second, each point in the dataset is assigned to the cluster associated with its nearest centroid, using the Euclidian distance metric. Third, once all the data points have been assigned, the centroids are recalculated as the centers of the clusters which result from the second step. The second and third steps are then repeated until no data point is reassigned. Guided by an “elbow plot”, we choose to use the K-means algorithm to divide the data into three clusters.⁸

To visualize the clusters in two-dimensional space, we can combine the clustering analysis with principal components analysis over all statements. Each point in Figure 10 represents a statement, with the x-axis showing a statement’s score for the first principal component and the y-axis showing the score for the second principal component. The color of each point shows the cluster which the statement belongs to. To show how each topic influences the principal components, we also overlay the factor loadings, which were previously shown in Figure 9. Figure 10 shows that statements can be broadly characterized into three groups: a cluster of statements which tend to emphasize “growth”; another cluster around “debt” and

⁸ We also considered hierarchical clustering, which produced similar results. Alternative clustering methods could also be used, in particular to assess the robustness of the results to outliers, for example K-medians.

“development”; and a third cluster which focuses on “crisis” management and “quota” issues.⁹

Figure 10. Principal component biplot with clustering

To assess how the content of statements has evolved over time, in Figure 11, we plot the statements relating to three different periods: a pre-crisis period from 2003-2006, a crisis-period from 2007-2011 that encompasses both the global financial crisis and the euro-area crisis, and a post-crisis period from 2012-2019.

Prior to the financial crisis, constituency statements placed a large focus on “debt” and “development” issues, with relatively few statements falling into the “growth” or “crisis” and “quota” issues clusters. Unsurprisingly, statements from 2007-2011 are largely grouped into the “crisis” and “quota” issues cluster. From 2012-2019, there has been a much greater number of statements within the “growth” cluster.

⁹ For robustness, we also implemented the K-means algorithm with the number of clusters, k , set to four. We find there is still a strong distinction between “growth” statements and “debt” statements. The “crisis” cluster is split into two clusters: one which has a high focus on “quota” and one which has a high focus on “crisis”.

Figure 11. Principal component biplots by period
2003-2006

2007-2011

2012-2019

Constituency Clusters Based on Topic Distributions

We can use topic distributions to analyze which constituencies produce similar statements to each other and which constituencies tend to differ in the topics they place emphasis on.

First, we group constituency statements with similar distributions at a given point in time, by computing the Bhattacharyya distance between each pair constituency statements and Communiqué biannually. The Bhattacharyya distance metric is designed to measure the similarity of probability distributions and provides an alternative to the Euclidean distance metric used in the K-means analysis above. Recall that in our dataset, the topic distribution of country i is depicted by topic distribution K_{it} . The Bhattacharyya distance between country i and country z is defined as

$$D_B(K_{it}, K_{zt}) = -\ln\left(\sum_j \sqrt{k_{it}^j k_{zt}^j}\right)$$

where $0 \leq D_B \leq \infty$. If the topic distributions K_{it} and K_{zt} are identical, the Bhattacharyya distance is 0. Whereas if there is no similarity between the topic distributions, the Bhattacharyya distance approaches infinity.

Using the measure of Bhattacharyya distance, we are able to calculate a distance matrix for each period in time, which highlights how similar constituency statements are to each other and to the overall IMFC communiqué. Based on our sample from April 2003 to October 2019, we produce 34 distance matrices. Figure 12 shows the distance matrix for October 2019, with each square illustrating the distance between the row constituency and the column constituency. Lighter shading indicates increased similarity between two statements. For example, the statements of CO have darker shading than most others, indicating it differed substantially in its topic distribution compared to other constituencies.

Figure 12. Bhattacharyya Distance Matrix, October 2019

To formally consider which groups of statements are similar in their topic content at a given point in time, we use a hierarchical clustering technique. In this technique, initially each constituency is considered as an individual cluster. At each iteration, the two clusters which are most similar, based on the Bhattacharyya distance measure, are merged. The process continues, until clusters are formed.

The result is a cluster tree for each biannual matrix, as visualized by the cluster dendrogram below in Figure 13 for October 2019 statements. The vertical axis of the dendrogram represents the distance or dissimilarity between clusters. The horizontal axis represents the objects and clusters. Each fusion of two clusters is represented by the splitting of a vertical line into two vertical lines. The vertical position of the split, shown by the short horizontal bar, gives the distance (dissimilarity) between the two clusters. In our analysis, we divide the samples into four clusters. The four clusters identified by the algorithm are highlighted by colored rectangles.

Figure 13. Cluster dendrogram, October 2019

The figure shows that the constituency CO is a cluster by itself, indicating that its statement was somewhat unique in the topics it focused on. Constituencies of JA and FF form a second cluster, US, CC, RU and several other constituencies form the third cluster, UK, GR and a number of other constituencies form the fourth cluster. Note that a cluster indicates that constituencies inside it tend to have similar topic distributions in their statements. It does not necessarily mean that they share agreed views on these topics, as clarified by the sentiment analysis performance in the next section.

Importantly, constituency clusters evolve over time. Our next step is to identify constituencies which frequently have similar topic distributions. To achieve this, we calculate the percentage of time that a constituency falls into the same cluster with another.

Table 1 shows for each of the five largest shareholders, which other constituencies cluster with them most often. We also consider IMFC communiqués in this analysis. Over the full sample period, the constituency CC, is in the same cluster as the US almost 50 percent of the time, indicating that these two large constituencies have frequently had similar topic priorities. The two large European constituencies, GR and NE, are also frequently in the

same cluster (58 percent of the time). Out of the five largest constituencies, the IMFC communiqué is most commonly in the same cluster as constituency GR (52 percent of the time).

Table 1. Constituencies which Cluster with the Largest Shareholders

Cluster 1		Cluster 2		Cluster 3		Cluster 4		Cluster 5	
US centered		CC centered		JA centered		GR centered		NE centered	
Top 5	% of time	Top 5	% of time	Top 5	% of time	Top 5	% of time	Top 5	% of time
CC	48%	IN	55%	CO	29%	NE	58%	GR	58%
AG	39%	AG	52%	US	29%	IMFC	52%	SZ	58%
BR	36%	IT	48%	AF	26%	CC	45%	IMFC	42%
CO	32%	US	48%	AP	26%	NO	45%	AP	39%
EC	32%	EC	45%	MD	26%	SA	45%	MI	39%

VI. EMERGING TOPICS

As the previous analysis shows, members' priorities evolve over time and new topics of interest frequently emerge. Indeed, the Fund has changed considerably in the past 20 years, and while the topics of growth, debt, and reform are persistently important across the membership, many other topics are attracting increasing attention. In this section we look at emerging topics – topics whose prominence increased significantly in recent statements and communiqués, even if they may not have made it to the top priorities yet. In what follows, we consider the topics of climate change, gender, fintech, and inequality; for those, we use sentiment analysis to gauge what is the likely attitude each constituency may have towards these emerging topics.

Panel A of Figure 14 highlights the weight placed on climate change over time by each of the constituencies, with darker shades indicating a greater topic frequency.¹⁰ It shows that climate change gained traction among some constituencies in 2015 and is slowly becoming more prominent as a topic of interest across the membership, with a significant jump in interest in 2019. The constituencies of FF, CO, and NO were early adopters of the climate change topic and on average the topic frequency of climate change is higher for these constituencies than for others. In contrast, the topic frequency of climate change for the US is less than for other constituencies. Not surprisingly, the IMFC communiqué ranks in between, as they are drafted to reflect the aggregate view of the membership.

¹⁰ The interpretation of the figures in Panel A and B are similar. The size of the square represents the absolute value of topic frequency (Panel A) or sentiment score (Panel B). For instance, in Panel B, the bigger the size, the more polarized the sentiment score is (sentiment can be positive or negative). A very small square represents a low frequency (Panel A) or a sentiment score near zero (Panel B). To differentiate a positive sentiment score from a negative sentiment score in Panel B, a color is also attached to the square box with dark blue representing positive score and dark red representing negative score.

Figure 14. Climate Change

To assess the sentiment attached to the topic of climate change, we use the popular “Valence Aware Dictionary and Sentiment Reasoner” (VADER) lexicon and rules-based sentiment tool (see Hutto and Gilbert 2015). VADER maps a dictionary of words to a sentiment score, based on whether a word conveys positive or negative sentiment.¹¹

Figure 14, Panel B shows the sentiment attached to climate change across membership and over time. The blue color in Panel B represents positive sentiment while the red color indicates a negative sentiment. Constituencies are ranked on the vertical axis according to their mean sentiment score across time. Consistent with the message of Panel A, Panel B shows that more members express their sentiment on climate change in recent years. Constituencies of NO, IT, and CC have showed the most positive sentiment on climate change over time. Other constituencies, like AF, a constituency representing a number of African countries, and MD, a constituency representing a number of countries in the Middle

¹¹ Studies have shown that VADER has certain limitations in capturing the subtle nuances when analyzing the sentiment of expressions that are domain and context dependent (see, for example, Rizun & Revina (2019)).

East, express mostly negative sentiment towards climate change. This could, however, reflect expressions of concern or worry on the implications of climate change by these constituencies.

Panel C summarizes the findings from Panel A and B. It decomposes constituencies into four quadrants according to their rankings on topic frequency and sentiment, with the 50th percentile as cutoffs. The four quadrants are (High, High), (High, Low), (Low, High), and (Low, Low). For example, constituencies like NO, IT, and FF fall into the upper right quadrant (High, High), as they are above the median on both topic frequency of climate change and their sentiment towards it.

Moving beyond climate change, inequality has been a popular topic for a number of years (Figure 15, Panel A), but has gained increased prominence in the past three years. Gender is another topic that has received increased attention from a number of constituencies, though not all, with generally positive sentiment (Figure 16, Panels A and B). The positive sentiment largely reflects constituencies “supporting” or “welcoming” gender-based initiatives. In recent years, constituencies began giving attention to fintech (Figure 17, Panel A), reflecting formal recognition of the importance of this rapidly growing industry. Sentiment on fintech tends to be positive across the membership (Figure 17, Panel B).

Figure 15. Inequality **Panel C: Quadrant**

Figure 16. Gender **Panel C: Quadrant**

Figure 17. Fintech **Panel C: Quadrant**

VII. CONCLUSION

Using a unique dataset which combines IMFC communiqués and constituency statements, we use text mining techniques to assess the evolving priorities of the IMF's membership. We find that statements show significant variation in their topic content across time and constituencies. We broadly find that there are three types of statement: those that emphasize growth, those that emphasize debt and development, and those focused on crisis management. While there are notable differences in which topics constituencies choose to focus on, we find that several constituencies persistently produce similar statements.

The dataset and the text mining techniques presented in this paper provide a useful basis for further research into what influences the priorities of the IMF membership. We provide a glimpse of potential applications by identifying key topics across constituencies and deploying sentiment analysis to uncover emerging topics. Tracking the evolving priorities and understanding their nature is particularly important to ensure the IMF continues to meet the needs of its membership.

REFERENCES

- Acosta, M., 2015, FOMC Responses to Calls for Transparency. Finance and Economics Discussion Series 2015-060. Washington: Board of Governors of the Federal Reserve
- Amaglobeli, D. and Crispolti, V. and Dabla-Norris, E. and Karnane, P. & Misch, F., 2018, Tax Policy Measures in Advanced and Emerging Economies: A Novel Database. IMF Working Papers 18/110, International Monetary Fund.
- Antweiler, W. and Frank, M., 2004, Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards. *Journal of Finance*, v59(3):1259–1295.
- Apel, M. and Grimaldi, M., 2012, The Information Content of Central Bank Minutes. Sveriges Riksbank Working Paper Series, No. 261.
- Baker, S. R., Bloom, N. and Davis, S. J., 2013, Measuring Economic Policy Uncertainty. Chicago Booth Research Paper 13-02.
- Bholat, D., Hansen, S., Santos, P. and Schonhardt-Bailey, C., 2015, Text mining for Central Banks: Handbook. Centre for Central Banking Studies (33). pp. 1-19. ISSN 1756-7270
- Blei, D.M. and Lafferty, J.D., 2006, June. Dynamic Topic Models. In *Proceedings of the 23rd International Conference on Machine Learning* (pp. 113-120).
- Bonacich, P., 1972, Technique for Analyzing Overlapping Memberships. *Sociological Methodology*, 4:176–185.
- Boukus, E. and Rosenberg, J., 2006, The Information Content of FOMC Minutes. Mimeo, Federal Reserve Bank of New York.
- Bruno, G., 2017, Central Bank Communications: Information Extraction and Semantic Analysis, Presented at IFC-Bank Indonesia Satellite Seminar on “Big Data” at the ISI Regional Statistics Conference 2017 Bali, Indonesia, 21 March 2017.
- Coleman, M. and Liao, T. L., 1975, A Computer Readability Formula Designed for Machine Scoring. *Journal of Applied Psychology*, 60:283–284.
- Fraiberger, S., Lee D., Puy D. and Ranciere R., 2018, Media Sentiment and International Asset Prices, IMF Working Paper IMF/18/274
- Gentzkow, M. and Shapiro, J. M., 2010, What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica*, Vol. 78, No. 1, pages 35-71.
- Greene, D., and Cross J. P., 2017, Exploring the Political Agenda of the European Parliament Using a Dynamic Topic Modeling Approach. *Political Analysis* 25 (1): 77–94.
- Grimmer, J., and Stewart B. M., 2013, Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political analysis* 21 (3): 267–297.
- Gunning, R., 1952, *The Technique of Clear Writing*. McGraw-Hill.

- HaCohen-Kerner, Y., Miller, D., Yigal, Y. (2020). The influence of preprocessing on text classification using a bag-of-words representation. *PLoS ONE* 15(5): e0232525
- Hansen, S., McMahon, M., and Prat, A., 2018. Transparency and Deliberation within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*, 133(2), 801-870.
- Hutto, C.J. and Gilbert, E., 2015. VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text.
- IMF, 2019, A Strategy of IMF Engagement on Social Spending, Background Paper V. Trends and Patterns in Fund Engagement on Social Spending: A Text Mining Analysis.
- Mihalyi, D. and Mate, A., 2018, Text-Mining IMF Country Reports – An Original Dataset, Unpublished Manuscript.
- Mimno, D., Wallach, H., Talley, E., Leenders, M. & McCallum, A. (2011). Optimizing semantic coherence in topic models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '11)*. Association for Computational Linguistics, USA, 262–272
- Rizun, N., and Revina, A. 2019. Business Sentiment Analysis. Concept and Method for Perceived Anticipated Effort Identification. In A. Siarheyeva, C. Barry, M. Lang, H. Linger, & C. Schneider (Eds.), *Information Systems Development: Information Systems Beyond 2020 (ISD2019 Proceedings)*
- UpLevel, 2019, Feature G20 project, available at <https://www.uplevel.work/blog/feature-examining-a-political-summit-using-text-analytics>
- Vapnik, V., 1995, *The Nature of Statistical Learning Theory*. Springer-Verlag, New York.
- Vapnik, V., and Lerner A., 1963, Pattern Recognition using Generalized Portrait Method. *Automation and Remote Control*, v24.
- Wang, C., Blei D., and Heckerman D., 2008. Continuous Time Dynamic Topic Models. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI'08)*, pp 579–586

I. APPENDIX

Table 1. Members of the IMF's 24 Constituencies

Constituency	Constituency members
AE	Angola, Botswana, Burundi, Eritrea, Ethiopia, The Gambia, Kenya, Lesotho, Liberia, Malawi, Mozambique, Namibia, Nigeria, Sierra Leone, Somalia, South Africa, South Sudan, Sudan, Eswatini, Tanzania, Uganda, Zambia, and Zimbabwe
AF	Benin, Burkina Faso, Cameroon, C.A.R., Chad, Comoros, D.R. Congo, Rep. Congo, Côte d'Ivoire, Djibouti, Equatorial Guinea, Gabon, Guinea, Guinea Bissau, Madagascar, Mali, Mauritania, Mauritius, Niger, Rwanda, São Tomé & Príncipe, Senegal, Togo
AG	Argentina, Bolivia, Chile, Paraguay, Peru, and Uruguay
AP	Australia, Kiribati, Korea, Marshall Islands, Federated States of Micronesia, Mongolia, Nauru, New Zealand, Palau, Papua New Guinea, Samoa, Seychelles, Solomon Islands, Tuvalu, and Vanuatu
BR	Brazil, Cabo Verde, Dominican Republic, Ecuador, Guyana, Haiti, Nicaragua, Panama, Suriname, Timor-Leste, and Trinidad and Tobago
CC	China
CE	Colombia, Costa Rica, El Salvador, Guatemala, Honduras, Mexico, Spain, and República Bolivariana de Venezuela
CO	Antigua and Barbuda, The Bahamas, Barbados, Belize, Canada, Dominica, Grenada, Ireland, Jamaica, St. Kitts and Nevis, St. Lucia, and St. Vincent and the Grenadines
EC	Austria, Belarus, Czech Republic, Hungary, Kosovo, Slovak Republic, Slovenia, and Turkey
FF	France
GR	Germany
IN	Bangladesh, Bhutan, India, and Sri Lanka
IT	Albania, Greece, Italy, Malta, Portugal, and San Marino
JA	Japan
MD	Afghanistan, Algeria, Ghana, Islamic Republic of Iran, Libya, Morocco, Pakistan, and Tunisia

MI	Bahrain, Egypt, Iraq, Jordan, Kuwait, Lebanon, Maldives, Oman, Qatar, United Arab Emirates, and Republic of Yemen
NE	Andorra, Armenia, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Georgia, Israel, Luxembourg, Moldova, Montenegro, Netherlands, Republic of North Macedonia, Romania, and Ukraine
NO	Denmark, Estonia, Finland, Iceland, Latvia, Lithuania, Norway, and Sweden
RU	Russian Federation and Syrian Arab Republic
SA	Saudi Arabia
ST	Brunei Darussalam, Cambodia, Fiji, Indonesia, Lao People's Democratic Republic, Malaysia, Myanmar, Nepal, Philippines, Singapore, Thailand, Tonga, and Vietnam
SZ	Azerbaijan, Kazakhstan, Kyrgyz Republic, Poland, Serbia, Switzerland, Tajikistan, Turkmenistan, and Uzbekistan
UK	United Kingdom
US	United States

Table 2. Summary Statistics

Constituency	No. of Documents	No. of Sentences	Word count (max)	Word count (min)	Word count (mean)	Word count (Std.)
AE	33	73	97	10	30	14
AF	33	79	76	8	27	13
AG	34	231	116	7	29	15
AP	34	70	84	9	28	13
BR	34	87	75	9	27	12
CC	34	83	66	8	26	12
CE	33	62	64	9	27	11
CO	34	82	91	8	28	13
EC	14	55	71	10	27	11
FF	33	114	70	7	26	12
GR	34	108	60	7	24	11
IN	34	137	68	7	25	11
IT	34	167	77	7	26	12
JA	34	65	76	9	30	14
MD	33	33	64	12	31	13
MI	34	74	75	9	27	12
NE	35	146	83	7	24	12
NO	34	89	50	7	22	9
RU	33	71	56	9	25	10
SA	32	78	54	8	26	10
ST	33	66	61	11	28	11
SZ	34	89	58	8	23	10
UK	33	77	70	8	27	12
US	33	56	54	8	25	10
IMFC	40	67	85	8	27	14

Table 3. Lexicon for Top 3 IMFC Topics

Growth	Debt	Reform
grown faster	bond issuance	bankruptcy code
growth	bonded_debt	bankruptcy_law
growth-	debt	cut_red_tape
growth—an	debt-	deregulating
growth-and	debt/gdp_ratio	deregulation
growth-boosting	debt_overhang	insolvency_law
growth-critical	debt_overhang	reduce_red_tape
growth-critical	debt_service-to-export_ratio	reform
growth-enhancing	debt_servicing	reform_fatigue
growth-enhancing	debt_servicing	structural_reform_agenda
growth-friendly	debt_sustainability	structural_reform_agenda
growth-friendly	debt_sustainability	structural_reforms
growth-friendly_manner	debt-to-gdp_ratio	unfinished_reform_agenda
growth-inducing	debt-management	
growth-oriented	debt-related	
growth-oriented	debt-service	
growth-promoting	debt-services	
growth-supporting	debt-sustainability	
growth-supporting	debt-to_gdp	
growth-supportive	debt-to-gdp	
private_sector-led_growth	debt-to-gdp_ratio	
pro-growth	debt-to-gdp-ratio	
	debt—was	
	dis-indebtedness	
	enhanced_hipc_initiative	
	heavily_indebted	
	high_debt	
	highly_indebted	
	highly-indebted	
	hipc-to-hipc_debt	
	indebted	
	indebtedness	
	indebtedness	
	nonconcessional_borrowing	
	nonconcessional_borrowing	
	non-concessional_borrowing	
	non-paris_club	
	non-paris_club	
	non-paris_club_creditor	
	non-paris_club_creditor	
	non-paris_club_official	
	paris_club	
	paris_club	
	paris_conference	
	paris_declaration	
	public-debt-to-gdp_ratio	
	sovereign_bond	
	sovereign_bond_issuance	
	sovereign_default	
	syndicated_loan	
	upper-credit_tranche	

ANNEX 1. Modelling with Latent Dirichlet Allocation (LDA)

To check the robustness of our key findings, we consider an alternative topic modelling approach. In our baseline approach, we identify topics using a dictionary-based approach. In the dictionary approach, the initial assignment of words to topics involves a degree of subjectivity. In this section, we check the robustness of our results using the latent Dirichlet allocation (LDA) topic modelling approach, which does not rely on the manual assignment of words to topics, but instead uses a probabilistic approach.

In the LDA approach, documents are treated as collections of words which arise as a result of a probabilistic process. Specifically, suppose that a corpus is made up of N documents with W unique words. The LDA approach assumes a fixed set of X topics, where each topic X is a probability distribution over the W words. Each document has its own distribution over topics. Therefore, the content of any given document depends on both its distribution over topics and the distribution over words for each topic. The LDA approach uses Bayesian inference to find the parameters governing the distribution of topics and the distribution of words within each topic, to maximize the likelihood of the observed words in the documents (see Hansen, McMahon and Prat (2018) for a detailed discussion of LDA and its application in economics).

We apply the LDA approach at the sentence level for constituency statements and IMFC communiqués. Following the parlance above, a sentence is treated as a ‘document’. We use a model with 10 topics ($X=10$). Therefore, for each sentence we obtain a distribution over the 10 topics, which loosely can be considered as the contribution of each topic to the sentence. For each topic we obtain a distribution over all the words within our corpus. Figure A.1 shows the ten words which have the highest probability within each topic, with the size of the word reflecting its probability. The topics cover a range of economic concepts, including fiscal policy (Topic 4), growth (Topic 7), and the quota formula and reform (Topic 10).

Using the distribution over topics for each sentence, we can identify the dominant topic within each sentence. For each statement, we then calculate the distribution over the ten topics by considering the proportion of sentences for which each topic is the dominant sentence. This provides a useful comparison to our baseline dataset, where we use the dictionary approach to calculate the frequency of topics within documents.

Figure A.1. Wordclouds for the topics identified using the LDA approach

Topic 1

Topic 2

Topic 3

Topic 4

Topic 5

Topic 6

Topic 7

Topic 8

Topic 9

Topic 10

We re-run our clustering analysis using the distributions identified by the LDA approach. Specifically, given the distribution for each statement, we use the K-means algorithm to group statements into three clusters.¹

Following the clustering approach in Section V, we combine the clustering analysis with principal components analysis, so that we can visualize the clusters in two-dimensional space. The results are presented in Figure A.2 using a loading plot, which shows the topics that most strongly influence the first two principal components. It shows that one cluster is characterized by a high prevalence of topic 10. Looking at the words which are most common in topic 10 in Figure A.1, we can see that this is a topic in which the quota formula and reform are particularly important. A second cluster is characterized by the prevalence of topics 4 and 7. These topics relate to growth and fiscal issues. In the final cluster, topics 5 and 8 are prevalent, which tend to relate to developing country and monetary policy issues. Overall, the nature of the three clusters identified with the LDA approach is very similar to the three clusters we identify using the dictionary approach in our baseline analysis.

Figure A.2. Principal component biplot with clustering using LDA approach

¹ The elbow plot suggests that three topics is appropriate, as is also the case with the dictionary approach.