

Predicting the impact of central bank communications on financial market investors' interest rate expectations

Andy Moniz¹ and Franciska de Jong^{2,3}

¹ Erasmus Studio, Erasmus University, Rotterdam, The Netherlands
{moniz}@rsm.nl

² Erasmus Studio, Erasmus University, Rotterdam, The Netherlands
{fdejong}@eshcc.eur.nl

³ Human Media Interaction, University of Twente, Enschede, The Netherlands
{f.m.g.dejong}@utwente.nl

Abstract. In this paper, we design an automated system that predicts the impact of central bank communications on investors' interest rate expectations. Our corpus is the Bank of England's *'Monetary Policy Committee Minutes'*. Prior studies suggest that effective communications can mitigate a financial crisis; ineffective communications may exacerbate one. The system described here works in four phases. First, the system detects salient aspects associated with economic growth, prices, interest rates and bank lending using information from Wikipedia. These *economic aspects* are detected using the *TextRank* link analysis algorithm. A multinomial Naive Bayesian model then classifies document sentences to these aspects. The second phase measures sentiment using a count of terms from the *General Inquirer* dictionary. The third phase employs Latent Dirichlet Allocation (LDA) to infer topic clusters that may act as intensifiers/diminishers for the economic aspects. Finally, an ensemble tree combines the phases to predict the impact of the communications on financial market interest rates.

Keywords: sentiment analysis-text mining-link analysis-financial markets

1 Introduction

Post the global financial crisis, there has been a dramatic change in the use of central bank communications as a central bank policy instrument [1,2]. Central banks communicate qualitative information to the financial market through statements, minutes, speeches, and published reports [3]. Communication is an important tool that a central bank can use to avert a crisis, by providing investors with its assessment of the risks and the measures it views as necessary to reduce those risks within the economy [1], [4]. Previous studies suggest that effective central bank communications can mitigate and potentially prevent a financial crisis; ineffective communications may exacerbate one [1], [5]. In [4], the Swedish central bank, the Riksbank, is criticized because its communications were "not clear or strong enough" leading up to the global financial crisis, such that the bank's information went "unnoticed" [1]. In this paper, we design

an automated system that predicts the impact of central bank communications on interest rate expectations, as derived via financial market patterns. For the purposes of this study, we analyze economic sentiment, as expressed in the '*Monetary Policy Committee Minutes*' [2] published by the Bank of England, that details its monthly interest rate decisions.

Financial markets scrutinize central bank communications for "*clues and shades of meaning about its assessment of the economy and the direction of where economic policy may be heading*" [1]. As a prediction task, the measurement and evaluation of sentiment is challenging due to the complexities and subtleties of interpreting bank communications [1]. The formation of economic policy is a balancing act between achieving high economic growth and financial stability, while targeting low inflation [2]. The relative importance of these objectives is dynamic, and varies depending on the prevailing economic conditions [2]. For example under benign economic conditions, high inflation may be construed by financial market investors as a negative signal for the direction of future interest rates. During the financial crisis of 2007-2009, high inflation was considered to be a positive signal by effectively lowering interest rates¹ [6]. This motivates a need for fine-grained sentiment analysis, to automatically detect economic aspects and predict the central bank sentiment expressed towards these aspects [7]. Such a model would provide investors with an automated system to decipher the complexities and interactions of *economic aspects*, to interpret the consequences of these interactions for the future path of interest rates, and to incorporate the information into their investment decisions. For a central bank, such a model would provide it with the ability to predict the impact of its economic policies on the financial markets. The resulting 'price discovery' process [2] may promote a more efficient functioning of financial markets.

Our approach consists of four phases. First, the system detects salient references to economic aspects associated with economic growth, prices, interest rates and bank lending and employs a multinomial Naive Bayesian model to classify sentences within documents. Economic aspects are identified in a pre-processing step, that employs a link analysis using the TextRank algorithm [8,9]. The second phase measures sentiment expressed for the economic aspects, using a count of terms from the *General Inquirer* dictionary [17]. The third phase employs Latent Dirichlet Allocation (LDA) to infer intensifiers/diminishers that may change the meaning of the economic aspects and economic sentiment [10],[7]. Specifically, the model categorizes whether the magnitude of the economic aspects has 'intensified' or 'diminished' over time [11,12]. We refer to the resulting topic clusters as *directional topic clusters*. Finally, an ensemble tree combines the model components to predict the impact of the communications on financial market interest rates over the following day.

¹ The real interest rate is the rate of interest a borrower expects to pay on debt after allowing for inflation and is equal to the nominal interest rate (set by the central bank) minus the rate of inflation [2]

The rest of this paper is structured as follows. Section 2 draws on literature from the fields of economics and discusses the implications for sentiment analysis and keyword detection. Section 3 models the individual components of the system. Section 4 outlines the corpus of central bank communications, provides an evaluation of the model components and then discusses the results. Section 5 concludes and suggests avenues for future research.

2 Related Work

2.1 Background: central bank research

Post the financial crisis, several central banks have identified communications, particularly ‘enhanced forward guidance’, as an important policy instrument within their economic toolkit [1,2]. Effective communications enhance a central bank’s public transparency, accountability and credibility [13], which in turn aids its ability to implement economic policies [14]. To date, there has been little research into text mining of central bank communications. In [14], the impact of different types of communications (press releases, speeches, interviews, and news conferences) are analyzed to determine which media sources impact interest rate expectations. The analysis does not, however, classify the language used in the documents. In [3], a term counting approach is adopted to analyze the sentiment contained within the meeting minutes of the US central bank (the Federal Reserve). In [3] and [15] Latent Semantic Analysis is employed to analyze the sentiment contained within the Bank of Canada’s minutes. The intention of this study is to design a fine-grained sentiment analysis approach to analyze the impact of central bank communications on financial market investors. To our knowledge, this remains an unexplored avenue of research.

2.2 Background: sentiment analysis

Traditionally, fine-grained sentiment analysis has been researched for the classification of online user reviews of products and movies [16]. Readers are often not only interested in the general sentiment towards an aspect but also a detailed opinion analysis for each of these aspects [7]. Evaluation is conducted by comparing model classifications versus ratings provided by users. The evaluation of economic sentiment is arguably a harder task, due to the lack of a clearly defined outcome to assess model performance. For example, which economic variable should a model’s predictions be evaluated against? The relative importance of the aspects (e.g. economic growth/inflation/interest rates) is subjective, may vary over time, and the measurement of the aspects is only known with significant time delay.

The traditional approach to text-mining within the field of finance is to count terms using the General Inquirer dictionary [17,18]. The dictionary classifies words according to multiple categories, including 1,915 positive words and 2,291 negative words. The *General Inquirer* was developed for psychology and sociology research and

while it is used for text mining within the field of finance, little research has been conducted as to its suitability within finance [19]. Aspects that are frequently mentioned in central bank communications, such as the terms ‘employment’, ‘unemployment’ and ‘growth’, are not classified by the *General Inquirer* dictionary. Adjectives are often needed before investors can interpret the patterns in the economy to form their interest rate expectations [3]. Furthermore, the terms ‘inflation’ and ‘low’ are classified as negative by the dictionary, yet ‘low inflation’ is a positive characteristic and indeed achieving this is a central bank’s core objective [2]. The terms ‘fall’ and ‘decline’ are classified as negative terms in the *General Inquirer* dictionary, yet the opposite terms ‘rise’ and ‘increase’ are not classified at all.

2.3 Background: keyword detection

Graph-based algorithms have received much attention [8] as an approach to keyphrase extraction and are considered to be state-of-the-art unsupervised methods [20]. In a graph representation of a document, nodes are words or phrases, and edges represent co-occurrence or semantic relations. The underlying assumption is that all words in the text have some relationship to all other words in the text. Such an approach is statistical, because it links all co-occurring terms without considering their meaning or function in text. Centrality is often used to estimate the importance of a word in a document [22], and is a way of deciding on the importance of a vertex within a graph that takes into account global information recursively computed from the entire graph, rather than relying only on local vertex-specific information [23]. The main advantage of such a representation is that selected terms are independent of their language [21].

3 Model to predict changes in investors’ expectations

In this section we describe the four phases of the system. First, the system detects salient references to economic aspects and employs a multinomial Naive Bayesian model to classify sentences within documents. The second phase measures sentiment expressed for the economic aspects, using a count of terms from the *General Inquirer* dictionary. The third phase employs a LDA model and categorizes whether the magnitude of the economic aspects has ‘intensified’ or ‘diminished’ [11,12]. Finally, an ensemble tree combines the model components to predict the impact of the communications on financial market interest rates over the following day.

3.1 Aspect detection

In [3] it is shown that *tf-idf* weighting selects infrequent terms that relate to major news events or economic shocks. Our approach is intended to detect the common economic themes that are discussed in central bank communications, that are more likely to influence investors’ interest rate expectations on a day-to-day basis [2]. To determine salient references, we employ a link analysis approach that detects the most frequently mentioned terms within two Wikipedia pages on Central Banking and In-

flation. The model employs TextRank [8], a ranking algorithm based on the concept of eigenvector centrality, to compute the importance of the nodes in the graph. Each vertex corresponds to a word. A weight, w_{ij} , is assigned to the edge connecting the two vertices, v_i and v_j . The goal is to compute the score of each vertex, which reflects its importance, and use the word types that correspond to the highest scored vertices to form keywords for the text [23]. The score for v_i , $S(v_i)$, is initialized with a default value and is computed in an iterative manner until convergence using recursive formula shown in Equation (1).

$$S(v_i) = (1 - d) + d \times \sum_{v_j \in \text{Adj}(v_i)} \frac{w_{ji}}{\sum_{v_k \in \text{Adj}(v_j)} w_{jk}} S(v_j) \quad (1)$$

where $\text{Adj}(v_i)$ denotes v_i 's neighbors and d is the damping factor set to 0.85 [8]. Figure 1 displays the resulting TextRank clustering of terms. The size of each node is directly proportional to the TextRank score of the respective economic aspect.




Fig. 1. Link analysis of frequently occurring terms. Different nodes colors reflect different communities identified using the Clauset-Newman-Moore algorithm.

We define economic aspects by employing a greedy algorithm to detect communities of terms within the network [31]. The *economic growth aspect* detects the frequency of the terms: ‘demand’, ‘goods’, ‘services’, ‘investment’. The *prices aspect* detects the terms: ‘inflation’, ‘prices’, ‘money’, ‘markets’, ‘currency’. The *interest rate aspect* detects the occurrence of: ‘interest’, ‘rates’, ‘policy’ and a *bank lending aspect* detects the terms: ‘banks’, ‘lending’ and ‘assets’. It is not surprising to see these terms appear in the link analysis, given a central bank’s remit is to maintain price and financial stability. The choice of terms is consistent with the text mining

research of [3] which identifies 'growth', 'price', 'rate', and 'econom' as the most frequently occurring terms for the US economy. Using the four economic aspects, the system employs a multinomial Naive Bayesian model [24] to categorize sentences within each document. The resulting classification labels form the basis upon which sentiment analysis is applied.

3.2 Polarity detection

In the second phase, the model computes a measure of economic sentiment associated with each of the economic aspects. We measure polarity by counting the number of positive (P) versus negative (N) terms, $(P - N)/(P + N)$ identified using the *General Inquirer* dictionary [17]. In line with [16], our goal is not to show that a term counting method can perform as well as a Machine Learning method, but to provide a baseline methodology to measure central bank sentiment and to draw attention to the limitations of the approach that is widely adopted by text mining studies in the field of finance as indicated in Section 2.2. The sentiment metrics that are associated with the economic aspects: economic growth, prices, interest rate and bank lending are labelled $Tone_{growth}$, $Tone_{prices}$, $Tone_{interest_rates}$ and $Tone_{bank_lending}$ respectively. A fifth sentiment metric, $Tone_{overall}$, is computed to measure the polarity associated with the overall document, without conditioning upon the *economic aspects*. The five sentiment metrics are included as separate components within the ensemble tree.

3.3 Detection of LDA directional topic clusters

Next we extend the baseline term-counting method by taking intensifiers and diminishers into account [11,12]. These are terms that change the degree of the expressed sentiment in a document (see Section 2.2). In the case of central bank communications, the terms describe how economic aspects have changed over time. We employ an implementation of LDA [10], and represent each document as a probability distribution over latent topics, where each topic is modeled by a probability distribution of words. In [7], LDA is found to capture the global topics in documents, to the extent that topics do not represent ratable aspects associated with individual documents, but define clusterings of the documents into specific types. For the purposes of training the LDA model, we consider each sentence within each central bank communication to be a separate document. This increases the sample size of the dataset (see Section 4.1) and is intended to improve the robustness of the LDA model for statistical inference. We implement standard settings for LDA hyper-parameters, $\alpha = 50/K$ and $\beta = .01$, where the number of topics K is set to 20 [25]. We manually annotate two of the topic clusters that capture 'directional' information [1] and appear to act as intensifiers/diminishers of meaning. We label the clusters *directional topic clusters*. Table 1 identifies the top terms associated with the two clusters. Representative words are the highest probability document terms for each topic cluster.

Table 1. Representative document terms associated with the directional topic clusters

'intensifier cluster'		'diminisher cluster'	
word	prob.	word	prob.
increase	0.150	moderated	0.190
strong	0.107	lower	0.161
accelerate	0.081	downwards	0.123
strength	0.063	difficult	0.102
support	0.058	less	0.070

Next for each central bank communication the LDA model infers the probabilities associated with the ‘intensifier’ and ‘diminisher’ clusters within each of the three economic aspects detected by the Naïve Bayesian classifier. The output of the model is a vector of six topic probabilities that proxy the central bank’s assessment that the economic aspects are intensifying/diminishing. We label the model *directional LDA model* and the respective probability vectors: $\text{Topic}_{\text{growth}\uparrow}$, $\text{Topic}_{\text{prices}\uparrow}$, $\text{Topic}_{\text{interest_rates}\uparrow}$ and $\text{Topic}_{\text{bank_lending}\uparrow}$ if the economic aspects are increasing and $\text{Topic}_{\text{growth}\downarrow}$, $\text{Topic}_{\text{prices}\downarrow}$, $\text{Topic}_{\text{interest_rates}\downarrow}$ and $\text{Topic}_{\text{bank_lending}\downarrow}$ if the economic aspects are decreasing. We include the topic probabilities as components within the ensemble tree.

4 Experiments

In this section we discuss the corpus of central bank communications and describe the investor patterns data used to evaluate the impact of the central bank communications on investors’ interest rate expectations. We then outline the evaluation of the ensemble classification tree, present the results and provide a discussion.

4.1 Data

We choose to analyze the interest rate minutes of the *Bank of England*. As cited in [3], central bank minutes are closely watched by investors to gauge the future direction of economic policies. Similar datasets for the US and Canadian central banks’ minutes are examined in [3] and [15]. The Bank of England announces the level of UK interest rates on the first Thursday of every month. The details that underpin this decision are only provided two weeks later and are published in the Bank of England’s ‘*Monetary Policy Committee Minutes*’. The communications are interesting to analyze, because changes in investors’ expectations on the day of the central bank communication may be attributed to the qualitative information contained within the meeting minutes rather than the interest rate decision announced two weeks before. Minutes typically include summaries of committee members’ views on economic conditions and discuss the rationale for their interest rate decisions [26]. The central bank’s minutes are, on average, 12 pages long (including a header page), and contain around 55 bullet points, typically with 5 sentences in each bullet. The documents are available from 1997, the year when Parliament voted to give the Bank of England operational independence from the UK government. We retrieve all meeting minutes

available between July 1997-March 2014² to create a corpus that consists of 199 documents. For the purposes of aspect detection and to train the LDA model, we remove the header page and define a document as an individual sentence within each of the meeting minutes. This expands the corpus to a collection of 53,195 documents.

To evaluate the ensemble tree’s predictions we utilize information obtained from financial market patterns. Interest rate futures contracts are financial instruments that enable investors to insure against or speculate on uncertainty about the future level of interest rates [27]. Changes in the price of the futures contracts therefore reflect changes in investors’ views on the future direction in central bank interest rates. Investors’ interest rate expectations for the following three, six and twelve months are derived and published daily by the Bank of England. We utilize investors’ twelve month ahead forecasts. This data series has the greatest data coverage compared to the three and six month series. Furthermore, the twelve month forecast horizon is consistent with the time horizon over which that the Bank of England conducts its economic policies [2]. To isolate the effect of the central bank communication on investors’ expectations, we compute the percentage change in the interest rate futures contract, as measured from the close of business on the day of the communication announcement until the close of business one day after. This narrow time window helps to minimize the influence on investors’ interest rate expectations from other financial market factors that may occur at the same time [28].

4.2 Experiment setup

We design the evaluation of the prediction model in stages to enhance our understanding of the model components. For a baseline, we evaluate the system’s predictions by using only the tone of the overall document (see Section 3.2). The approach does not take into account individual economic aspects or diminishers/intensifiers [11,12]. We label the model *naïve tone*. This approach is consistent with the methodology adopted by the financial literature [18]. Next we compare the outcomes of an ensemble model that combines the tone associated with each of the economic aspects: economic growth, inflation and interest rates (see Section 3.2). We label this the *economic aspects model*. A third model compares the outcomes from an ensemble model that combines the tone of eight directional economic aspects, that combines the intensifiers/diminishers associated with the four economic aspects (see Section 3.3). We label this the *directional LDA model*. Finally, we combine the components in a single ensemble tree and refer to the system as the *joint aspect-polarity* model.

Learning and prediction is performed using an ensemble tree. The goal of ensemble methods is to combine the predictions of several models built with a given learning algorithm in order to improve generalizability and robustness over a single model. We use the Random Forest algorithm [30], that employs a diverse set of classifiers by introducing randomness into the classifier construction. Experiments were validated

² Central bank communications announced in August 1997 were excluded from the analysis because the communication document was not readily available in a machine readable format.

using five-fold cross validation in which the dataset is broken into five equal sized sets; the classifier is trained on four datasets and tested on the remaining dataset. The process is repeated five times and we calculate the average across folds. For evaluation, we select Mean Absolute Error (MAE), Root Mean Squared Error and Spearman's rho (ρ). We also examine Spearman's rho since prediction may be considered to be a ranking task. The formulae are displayed in Equation (2) below.

$$MAE = \frac{1}{n} \sum_{i=1}^n |O_i - E_i| \quad , \quad RMSE = \left[\frac{1}{n} \sum_{i=1}^n |O_i - E_i|^2 \right]^{1/2} \quad , \quad \rho = 1 - \frac{6 \sum (O_i - E_i)^2}{n(n^2 - 1)} \quad (2)$$

where E_i is the model's predicted value, O_i is the realized value, and n is the number of observations. MAE measures the average magnitude of the forecast errors without considering direction; RMSE penalizes errors and gives a relatively high weight to large errors. A smaller value of MAE or RMSE indicates a more accurate prediction. Spearman's rho is a non-parametric measure of the degree of linear association between the predicted and realized values, and is bound between the range -1 to +1 [29]. A positive Spearman's rho indicates the model's predictive ability; a negative value indicates a poor model fit.

4.3 Experiment results

The evaluation metrics from the model components are shown in Table 2.

Table 2. Evaluation of the model components

Model	MAE	RMSE	ρ
naive tone	0.022	0.016	-0.187 ^{***}
economic aspects	0.018	0.013	-0.044
directional LDA	0.019	0.014	0.041
joint polarity model	0.015	0.011	0.034

The asterisks provide the level of significance where *** indicates that the model's predictions versus forecasts are statistically, negatively significant at the 0.1% level.

The naïve tone model, which is similar to the approach commonly adopted by text mining studies in the field of finance, shows the worst performance. It exhibits the highest MAE and RMSE. The rank correlation of the model's forecasts with realized changes in investors' interest rate expectations is negative and is highly statistically negative, implying that documents that are predicted to have a positive/negative impact on investors' interest rate expectations end up having the reverse effect. The economic aspects and directional LDA models exhibit monotonic decreases in MAE and RMSE, suggesting a slight improvement in the model fit. Spearman's rho, however, is again negative, albeit to a lesser extent. Finally, the joint aspect-polarity model, that includes all model components in the ensemble tree, displays the lowest MAE and RMSE. The mildly positive Spearman's rho is consistent with previous studies within the field of finance. As cited in [19], many factors influence the financial mar-

kets; a low, positive correlation provides sufficient comfort of the model’s predictive power.

4.4 Discussion

One interpretation of the experiment results is that multiple aspects are needed to improve the accuracy of the system. The positive Spearman’s rho for the joint model versus the negative Spearman’s rho for the naïve tone, economic aspects and directional LDA models may be indicative of a non-linear relationship between the components that is only evident when the models are combined rather than considered in isolation. One of the strengths of a regression tree is that it does not assume a functional form, allowing it to detect interactions between model components. To aid our understanding of prediction in the joint model, Figure 2 displays the decision tree results for one of the folds. The values in the grey boxes provide the predicted percentage change in investors’ interest rate expectations associated with the sentiment contained within the central bank communication. A positive value indicates that the impact is expected to lead to an increase in investors’ interest rate expectations, while a negative value indicates an expected decrease in interest rate expectations.

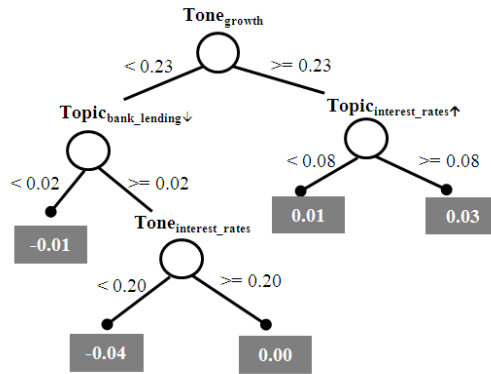


Fig. 2. Example decision tree from one of the folds

The regression tree identifies the interaction between the *directional topic clusters* and *Tone* measures. The primary decision in the decision tree is central bank sentiment towards economic growth. The right hand path indicates that if a central bank communication emphasizes positive economic growth and discusses interest rate increases, the model predicts that investors’ expectations of future interest rates will rise by 3%. The left hand path predicts that if a central bank tone towards economic growth is low, declining bank lending and the tone towards interest rates is negative, investors will reduce their expectations of future interest rates by 4%.

5 Conclusion

The goal of central bank communication is to make messages as clear, simple and understandable as possible to a wide range of audiences [1]. In this study, we focus on one specific audience, namely financial market investors. Investors play a key role for the implementation of a central bank's economic policies [1,2]. The outcome of our study may feed the design of a system that can predict the impact of central bank communication on formation of investors' interest rate expectations. The results of the joint aspect-polarity model suggest that investors may benefit by incorporating a measure of central bank sentiment to forecast interest rates.

In this study we evaluate model performance using prices from financial market instruments. The market price of an interest rate contract implicitly measures the average investor's interest rate expectations [27]. It is also possible to compute an 'implied probability distribution' of those expectations [27]. In future work we plan to evaluate a range of metrics, including the dispersion of the expectations as a proxy of investor uncertainty. Post the 2007–09 financial crisis, central banks have broadened the range of their communication, including the use of social media, live broadcasts, podcasts and blogs, to deliver their messages quickly and efficiently [1]. In future research, a wider range of central bank communications, including those expressed via social media, will be integrated into our study. We also intend to examine alternative approaches to select economic aspects, including dynamic approaches that reflect the usage of terms as central bank communications change over time.

Acknowledgement

The research leading to these results has partially been supported by the Dutch national program COMMIT.

References

1. Vayid, I. Central Bank Communications Before, During and After the Crisis: From Open-Market Operations to Open-Mouth Policy, Bank of Canada Working Paper (2013)
2. Bank of England. Monetary policy trade-offs and forward guidance (2013)
3. Boukus, E., Rosenberg, J., V. The information content of FOMC minutes (2006)
4. Meyersson, P., Karlberg, P.P, A Journey in Communication: the Case of the Sveriges Riksbank SNS Förlag (2012)
5. Viñals, J. Lessons from the Crisis for Central Banks. IMF Speech (2010)
6. Danthine, J, P. Causes and consequences of low interest rates. Speech by Mr Jean-Pierre Danthine, Vice Chairman of the Governing Board of the Swiss National Bank, at the Swisscanto Market Outlook 2014, Lausanne (2013)
7. Titov, I, McDonald, R., T. Modeling online reviews with multi-grain topic models. Proceeding of the 17th WWW (2008)
8. Mihalcea, R., and Tarau, P. TextRank: Bringing order into texts. In Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (2004)

9. Brin, S., Page, L. The anatomy of a large-scale hypertextual web search engine. *Comput. Networks ISDN Systems*, 33:107–117 (1998)
10. Blei, D., M., Ng, A., Jordan, M., I. Latent Dirichlet Allocation, *Journal of Machine Learning Research* 3, 993-1022 (2003)
11. Kennedy, A., Inkpen, D. Sentiment Classification of Movie Reviews using Contextual Valence Shifters. *Computational Intelligence*, vol.22(2), pp.110-125 (2006)
12. Polanyi, L., Zaenen, A. Contextual Valence Shifters, *AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications* (2004)
13. Carney, M. Panel discussion comments to the BIS Conference on The Future of Central Banking under Post-Crisis Mandates (2010)
14. Fay, C., Gravelle, T. “Has the Inclusion of Forward-Looking Statements in Monetary Policy Communications Made the Bank of Canada More Transparent?” *Bank of Canada Discussion Paper* (2010)
15. Hendry, S., Madeley, A. Text Mining and the Information Content of Bank of Canada Communications, *Bank of Canada Working Paper* (2010)
16. Pang, B., Lee, L., Vaithyanathan, S. Thumbs up? Sentiment classification using machine learning techniques. *Proceedings of EMNLP-02* (2002)
17. Stone, P., Dumphy, D. C., Smith, M. S., and Ogilvie, D. M. *The General Inquirer: A Computer Approach to Content Analysis*. The MIT Press. (1966)
18. Tetlock, P., Saar-Tsechansky, M. Macskassy, S. More Than Words: Quantifying Language to Measure Firms’ Fundamentals. *Journal of Finance*, Vol. LXIII, No. 3, 1437-1467 (2008)
19. Loughran, T., McDonald, B. When is a liability not a liability? Textual analysis, dictionaries and 10Ks. *Journal of Finance* 66, 35–65 (2010)
20. Liu, F., Pennell, D., Liu, F., Liu, Y. Unsupervised approaches for automatic keyword extraction using meeting transcripts. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American* (2009)
21. Litvak, M., Last, M. Graph-Based Keyword Extraction for Single-Document Summarization. *Proceedings of the 2nd Workshop on Multi-source, Multilingual Information Extraction and Summarization, Coling 2008*, pp. 17-24. Association for Computational Linguistics (2008)
22. Opsahl, T., Agneessens, F., Skvoretz, J. Node centrality in weighted networks: Generalizing degree and shortest paths (2010)
23. Boudin, F. A Comparison of Centrality Measures for Graph-Based Keyphrase Extraction (2013)
24. McCallum, A., and Nigam, K. A comparison of event models for naive Bayes text classification. *AAAI-98 Workshop on Learning for Text Categorization* (1998)
25. Griffiths, T. L., Steyvers, M. Finding scientific topics. *Proceedings of the National Academy of Science*, 101, 5228-5235 (2004)
26. Danker, D., J., and Luecke, M., M. Background on FOMC Meeting Minutes, *Federal Reserve Bulletin* 175-179 (2005)
27. Clews, R., N. Panigirtzoglou, Proudman, J. Recent developments in extracting information from options markets, *Bank of England Quarterly Bulletin* (2000)
28. Mackinlay, A., C. *Event Studies in Economics and Finance*. *Journal of Economic Literature* (1997)
29. Maritz, J., S. *Distribution-free statistical methods*. Chapman and Hall, (1984)
30. Breiman, L. *Random Forests*. *Machine Learning* (2001)
31. Clauset, A., Newman, M. E. J., and Moore, C. Finding community structure in very large networks. *Physical Review* (2004)