From Headlines to Hard Data: Nowcasting the Bahamian Economy with News Sentiment

Crystal Gooding[1]

October 2024

Abstract

Monitoring the economy in real time has become increasingly critical for economists and policymakers. This study utilizes economic and financial news to 'nowcast' economic conditions in The Bahamas. Using a time varying version of the Latent Dirichlet Allocation (LDA) model, Bahamian news articles published between 2015Q1 to 2024Q1, are organized into their main themes. Sentiment scores for these themes are calculated using a sentiment dictionary, scaled by word importance, and aggregated to create a quarterly sentiment index. The index's predictive power is evaluated through the Autoregressive Distributed Lag (ARDL) model and Impulse Response Functions (IRFs). The main findings indicate that an increase in positive news sentiment is linked to immediate GDP growth that persists for 14 months. These results suggest that news sentiment can serve as a valuable leading indicator for real-time economic monitoring, providing timely insights for policymakers.

JEL Classification: C11; C53; E00; E27

Keywords: Nowcasting; Sentiment Analysis; Latent Dirichlet Allocation (LDA)

[1] Crystal Gooding: CLGooding@centralbankbahamas.com. The views expressed in this paper are those of the author and do not necessarily represent The Central Bank of The Bahamas.

Section 1: Introduction

Detecting real-time shifts in economic activity would allow policymakers to assess the immediate impact of events and implement timely interventions. This is why 'nowcasting', which is making forecasts about the current economic conditions, has gained popularity in economic literature (Banbura et al., 2012). However, many traditional economic indicators are often released with a delay, which limits their usefulness in real-time decision-making. In recent years, text data, such as financial/economic news articles have been identified as a potential source of timely economic data (Gentzkow et al. 2019). News articles are available daily and are typically not revised. In addition, there is evidence of a link between news sentiment, consumption and output in the literature, with consumer sentiment serving as the transmission channel (Doms and Morin, 2004; Barsky and Sims, 2012; Shapiro et al., 2022). For example, negative news reports may lead to a deterioration in consumer sentiment, which could eventually result in reduced consumption and output.

The present study evaluates the predictive power of news sentiment by constructing a quarterly news sentiment index over the period 2015Q1-2024Q1 using four Bahamian news sources to nowcast Bahamian economic activity. The main contribution of this study is a variation of the time-varying Latent Dirichlet Allocation (LDA) model that captures relative sentiment across time by identifying the main topics in Bahamian news in each quarter. Then, a sentiment dictionary scores the tone of each topic to create the news sentiment index. Given the size of the dataset, the nowcast is conducted using the short run autoregressive distributed lag model (ARDL), and impulse response functions (IRFs) generated using smooth local projection methodology, which are suitable for smaller sample sizes.

The results suggest that the sentiment index could serve as a leading indicator to predict Bahamian economic activity in the short term. The ARDL results indicate that news sentiment and GDP have a positive and statistically significant relationship in the short run. In particular, a one-unit increase in news sentiment is associated with a 16.6% rise in real GDP *ceteris paribus*. The IRF provides further evidence to support this theory as the response of real GDP to a boost in optimistic news sentiment is an immediate increase, with the effect persisting for around 14 months.

The structure of the paper is as follows. Section 2 reviews the relevant literature. Section 3 introduces the time varying Latent Dirichlet Allocation model. Section 4 presents the newspaper data, then it analyses the topics that the LDA identifies within the newspaper data. Section 5 outlines the process of compiling the sentiment index. Section 6 discusses the nowcasting data and models, respectively. Section 7 presents the outcomes of the 'nowcasting' exercises before the paper concludes in Section 8.

Section 2: Literature Review

Since this study uses data from the Bahamas, it is important to select methodologies that can create a sentiment index and produce reliable nowcasting exercises despite the smaller dataset. However, many sentiment analysis methodologies only utilize subsets of text data. For example, Baker et al. (2016) and Park et al (2024) calculated an index based on the number of news articles that signalled economic uncertainty, while Aguilar et al. (2020) counted articles that indicated uncertainty or a boom. Barbaglia et al. (2024) constructed daily news sentiment indexes by using the fine grained, aspect based sentiment (FiGAS) approach (as developed by Consoli et al. (2022)) to assign sentiment scores to predefined aspects of economic news.

Reducing the current dataset would negatively impact this analysis, so it is preferable to employ a technique that utilizes all of the observable text data. Such techniques include the dictionary-based approach formulated by Nguyen and La Cava (2020), who used a sentiment dictionary to calculate an index based on the difference between the count of positive and negative word. However, this straightforward method does not fully account for the contextual importance of words because it does not analyse the relationships between the words (Ash and Hansen 2023). For example, while the word “growth” might be classified as positive, it can carry a negative connotation in a phrase like “slow growth”. Without considering context, the sentiment index might not reflect the true underlying sentiment of the text.

As a result, this study employs a time varying version of the Latent Dirichlet Allocation (LDA) model, which is an unsupervised machine-learning model that was originally introduced by Blei et al. (2003). It infers the broader context of text data by capturing the relationships between words when it groups related terms into topics. It assumes that text data is a mixture of hidden topics, and that each topic is a mixture of words. Although this algorithm requires less human intervention than the supervised algorithms presented by Shapiro et al. (2022), the model’s output remains easy for humans to interpret, especially compared to other models (Ash and Hansen 2023). In addition, LDA does not need any information about the subject matter beforehand unlike approaches such as FiGAS that requires predefined aspects.

The time-varying extension of LDA tracks the evolution of topics over time, or in other words, the relationships between words over time. Blei and Lafferty (2006) developed an early version of this model, while newer adaptations include the rolling LDA (Bittermann and Rieger, 2022), which uses a moving window to capture gradual shifts in topics. Refining this version of LDA, Van Dijk and de Winter (2023) incorporated hierarchical Bayesian estimation to derive layered topics, and then applied a sentiment dictionary to these topics for their economic activity nowcast. Other notable studies include Larsen and Thorsrud (2018), who employed another version of the time-varying LDA to examine the influence of media narratives on business cycles across different countries, and Huang et al., (2018) who modelled topics to create absolute and relative news sentiment indexes.

The results of these studies contended that news sentiment improved the accuracy of short-term economic forecasts. However, some used methods that were again better suited to large datasets, such as vector autoregressions (Baker et al, 2016 and Aguilar et al., 2020), dynamic factor models (Van Dijk and de Winter, 2023 and Larsen and Thorsrud, 2018) and Mixed Data Sampling regressions (Barbaglia et al., 2024). In contrast, models such as the ARDL (Park et al., 2024), IRFs generated using local projection methodology, (Nguyen and La Cava, 2020; Shapiro et al., 2022; Park et al., 2024) and probit models (Aguilar et al., 2020 and Huang et al., 2018) are more suitable for small sample sizes, aligning with the approach in this study.

Section 3: Time Varying Latent Dirichlet Allocation Model

The first stage in developing the news sentiment index is to organize the news articles into quarterly time slices, spanning from January 1, 2015, to March 31, 2024. This allows the sentiment index to be directly comparable to other quarterly macroeconomic indicators. This approach also ensures that each quarter is a snapshot of the news content during that specific time. This approach also reduces computational load and complexity because the LDA model analyses smaller chunks of the data sequentially as opposed to the entire dataset at once. Next, several pre-processing steps transform the raw text into a format suitable for analysis (discussed in Section 4.2). Once pre-processed, each quarterly time slice is analysed individually by the LDA algorithm. As is common with models trained using Bayesian inference techniques, LDA begins with preliminary predictions (or priors) about the underlying themes in the data (Papanikolaou et al., 2017). The predictions are the number of topics present in the text data, the generative process of the text, and the distribution of topics and words. The model uses these priors to estimate the true mixture of topics and words in the documents and across time. The steps of this process are discussed in the following sections.

Section 3.1: Choosing the Number of Topics

The researcher must estimate number of topics that exist in the text data. This decision significantly affects the quality of the results. Underestimating the number of topics results in vague topics, while overestimating leads to redundant or overly similar topics (Gan and Qi, 2021). To determine the optimal number of topics, a series of diagnostic plots are generated following several methodologies (Arun, 2010; CaoJuan, 2009; Deveaud, 2014; Griffiths, 2004). These plots measure the coherence, similarity and redundancy of the model's topics at different topic counts (as seen in Figures A1-A4 in the appendix). These diagnostics indicate that 19 topics is a reasonable approximation for each quarter.

Section 3.2: The Generative Process

LDA assumes that the text data arises from a generative process that defines a joint probability distribution over hidden variables (topics) and observed variables (words) in the text, which can be used to compute the conditional (or posterior) distribution of the hidden variables given the observed variables (Blei et al. 2012). The conditional distribution is expected to differ for each quarter, since news media coverage evolves over time. By creating independent time slices, the model captures the relevant topics in each quarter.

Section 3.3: Initial Topic and Word Assignments

To approximate the posterior distribution of topics and words, the model begins by using a Dirichlet multivariate probability distribution as its prior to randomly assign topics to documents and words.

Section 3.4: Refining the Posterior Distribution Estimates

After initial predictions about the conditional distribution are generated, the model iteratively updates these topic assignments. For each word, it observes how frequently it appears alongside other words in similar contexts. Based on these co-occurrences, the model creates probability distributions for each topic. Each distribution consists of a list of words that represent a topic, with probabilities indicating the likelihood of each word belonging to a topic (Valtonen et al., 2024). A word can belong to multiple topics but with varying probabilities, which means that it is more likely to be associated with certain topics.

With each iteration, the model improves its estimates of which words belong to which topics, and the topic proportions of each document in each quarter. To simplify this process, the model incorporates the collapsed Gibbs sampling algorithm[2], which iteratively processes each word by temporarily excluding its current topic assignment, and then reassigns it conditioned on the topic assignments of all other words. This process continues until the model converges to a steady state, at which point topic assignments are used to estimate topic proportions for the documents and word-topic distributions (Steyvers and Griffiths 2007). Given that the early iterations of this algorithm are typically poor estimates of the conditional distribution, it is common to disregard the initial sampled topic assignments, a phase known as the burn-in period (Papanikolaou et. al 2017). The present study discards the first 200 iterations in each quarter, and uses the later samples to approximate the posterior distribution, which are expected be more reliable.

This approach is most closely related to the rolling LDA methodology in Bittermann and Rieger (2022). However, this paper diverges by pre-processing the text data after creating the time slices. In addition, this approach does not incorporate information from previous quarters into subsequent quarters to allow the model to focus on identifying topics within each time slice without being constrained by the structure or influence of adjacent time slices.

Section 4: News Sentiment Data

This section describes the data used to compile the news sentiment index. The primary goal is to transform the text data into numerical data in order to assess its potential as a macroeconomic indicator.

Section 4.1: Newspaper Articles

The news sentiment index is constructed with 38,974 news articles that were published by four major news outlets in the Bahamas over the period 2015Q1-2024Q1. Specifically, they are the Tribune, the Nassau Guardian, The Bahama Journal and Bahamas Eyewitness News. To collect articles that are relevant to the Bahamian economy, the entire business section of each news website is web scraped. Additional keywords collect other pertinent articles that are categorized in different sections of the news outlet. The keywords are “econ”, “employ”, “tour”, “real estate”, “construction”, “price”, “inflation”, “invest”, “capital”, “BISX[3]”, “bank”, “fiscal” and “financial”. The keyword “tour” encompasses related terms such as “tourism” and “tourist,” while “econ” captures terms like “economy” ,“economic,” and “economist.” Similarly, “employ” includes “employment,” “employed,” “unemployed,” and “unemployment.” At the end of the webscraping process, articles that had a title but no content are removed, along with duplicate articles.

Section 4.2: Pre-processing Text Data

Conducting any analysis with unprocessed text data would be computationally inefficient given its unstructured and high-dimensional nature (Gentzkow et al. 2019). As a result, it is

[2] A member of the Markov Chain Monte Carlo (MCMC) class of algorithms that is used to draw samples from a probability distribution (Steyvers & Griffiths, 2007)

[3] BISX is the acronym for the Bahamas International Securities Exchange

common to reduce the size and complexity of the text data before applying any natural language processing algorithm. The present study synthesizes common pre-processing steps as outlined in Valtonen et al., (2024). The purpose of these steps can be divided into three categories: noise reduction, complexity reduction, and vectorization, all of which are essential to prepare the data for topic modelling.

Removing Noise

The first stage of pre-processing removes irrelevant or non-informative elements from the text, referred to as "noise." These characters or words do not contribute to the context or sentiment of the text. So, to start, non-word characters are deleted such as punctuation and numbers, URLs, white spaces, and hashtags. Next, stop words, which are frequently occurring words that add little contextual meaning, such as “the,” “and,” and “a”, are eliminated to prevent them from overshadowing more important words in the text. For this study, the stop-word list is a combination of the default stop-word lists found in the R packages "tm" and "SnowballC." In addition, the list is tailored to the Bahamian news articles, so common but non-informative words like “Bahamian,” “Bahamas”, the names of news websites, and names are included in the list. Finally, the list contains numerical count words such as "hundred," "million," and "billion" because they do not add semantic value to the text. The full list of stop-words can be found at the end of the Appendix.

Reducing Complexity

The second stage of pre-processing reduces the complexity and dimensionality of the text. First, every word is transformed to lower case so that the LDA model only has to analyse one consistent case. Next, the text is tokenized, which means that the continuous bodies of text are split up into smaller discrete units or individual words. Tokenization helps to simplify the text structure, which makes it easier to analyse. To further decrease the dimensionality of the text, lemmatization is used. Text data often contains multiple variations of the same word and lemmatization reduces the number of variations by returning each word to its dictionary form, or "lemma," based on its part of speech (POS) and context. For example, the words “better” and “best” could be converted to the word “good.”

By reducing the variety of word forms, the LDA algorithm can quickly recognize the patterns and relationships between words. Although lemmatization is computationally intensive, it preserves the context of words, making it preferable for maintaining the meaning of the news articles (Schofield and Mimno, 2016). A spell check is also applied to the lemmas to ensure correct spelling. After, POS tagging extracts and retains the nouns, which typically captures the main idea of sentences and thus are useful for topic modelling.

Vectorization

As previously stated, LDA groups words into topics based on how often they co-occur within the same context. These co-occurrences are calculated using the numerical representation of the text data, which is generated through a process called vectorization. While several vectorization methods exist, the present study adopts the document term matrix (DTM) approach. In this type of matrix, rows correspond to documents (in this case, news articles), while the columns represent unique words, and the cells contain the frequency of each word within a specific document. In essence, the DTM summarizes the distribution of words across all documents, allowing the LDA model to identify relationships between words and assign them to appropriate topics.

By generating a separate DTM for each quarter, the LDA algorithm tracks the shifts in word frequency patterns over time, which enables the detection of enduring or emerging topics. For example, the word "virus" might have appeared infrequently in DTMs before 2020 but would appear more frequently in DTMs after 2020, especially in response to the COVID-19 pandemic, thus signalling the appearance of a new topic.

Section 4.4: Analysis of Topics

Analysing the evolution of the LDA model's topics over time provides insight into the subjects or events that dominated Bahamian news and potentially influenced economic conditions. Doms and Morin (2004) posited that consumers updated their expectations about the economy during periods of intense news coverage because the cost of acquiring information decreases. As such, evaluating the topics that the media frequently reports to the public is essential to understand which events are perceived as most economically important and how they might influence public sentiment. To determine the prevalent news topics, the present study follows the methodology of Zhu et al. (2016) and calculates the topic intensity for each quarter. Topic intensity is derived directly from LDA results, which provide the proportions of each topic within the document collection. Summing these proportions for each quarter calculates the intensity of each topic over time. In essence, if the media publishes a large number of news articles on a particular topic in a given quarter, its intensity will be high. Monitoring topic intensity reveals the rise and fall of certain topics, which can assist legislators in real time policymaking because it identifies emerging concerns in the media, which would allow them to respond swiftly to issues before it impacts the behaviour of the public. Although the present dataset of just under 39,000 articles is relatively small, the analysis still captures notable shifts in media focus. A heat map (Figure 1) visualizes these intensities. The horizontal axis represents the time slices while the vertical axis denotes each of the 19 topics. Darker colours correspond to higher topic intensity, which means that those topics dominated the news.

For example, in the fourth time slice, or 2015Q4, the darkest cell is topic 13, which contains words such as "china, bank, project, resort, liquidator, export, court, construction, contractor, and chapter". Topic 13 is most likely about the construction of the Baha Mar Resort, which was widely discussed in the news at the time (Jett, 2016). These results indicate that LDA effectively captures the relevant themes within a given time period. The model is also able to track enduring topics as evidenced by the fifth time slice or 2015Q1. One of the darkest cells in the quarter is topic 6 appears to allude to Baha Mar again, as it consists of terms such as "project, construction, china, open, developer, export, contractor, resort, bank, and employee". The LDA model is also able to detect entirely new topics such as the landfall of Hurricane Dorian, as evidenced by topic 8, which is the darkest cell in quarter 19 or 2019Q3. This topic is composed of words such as "hurricane, storm, island, effort, relief, disaster, recovery, rebuild, home, and water".

Further, quarter 22, or 2020Q2 marks the onset of the COVID-19 pandemic, with topic 6 consisting of words like "tourism, hotel, travel, visitor, airport, open, plan, resort, industry, and protocol," and topic 11 containing terms such as "case, health, test, virus, resident, border, testing, flight, officials, and travel". As the pandemic persisted, it continued to dominate news coverage, as seen in quarter 24 (2020Q4), where topic 18—composed of terms such as "case, health, number, death, virus, spread, hospital, officials, healthcare, and measure”—was the most intense. By quarter 26 (2021Q1), topic 11 included words like “vaccine, health, case, vaccination, test, emergency, death, dose, testing, and virus”. Providing a complete list of words from each of the 19 topics across quarters would be overwhelming, so word clouds (Figures A5 to A7) provide a sample of key topics. These clouds show only the top 10 words per topic to avoid cluttering.

This analysis of topic intensity provides the foundation for the next step, which is quantifying the sentiment of these news topics. By recognizing the dominant topics, the factors behind shifts in the sentiment index can be identified.

Figure 1: Heat Map of Topic Intensities

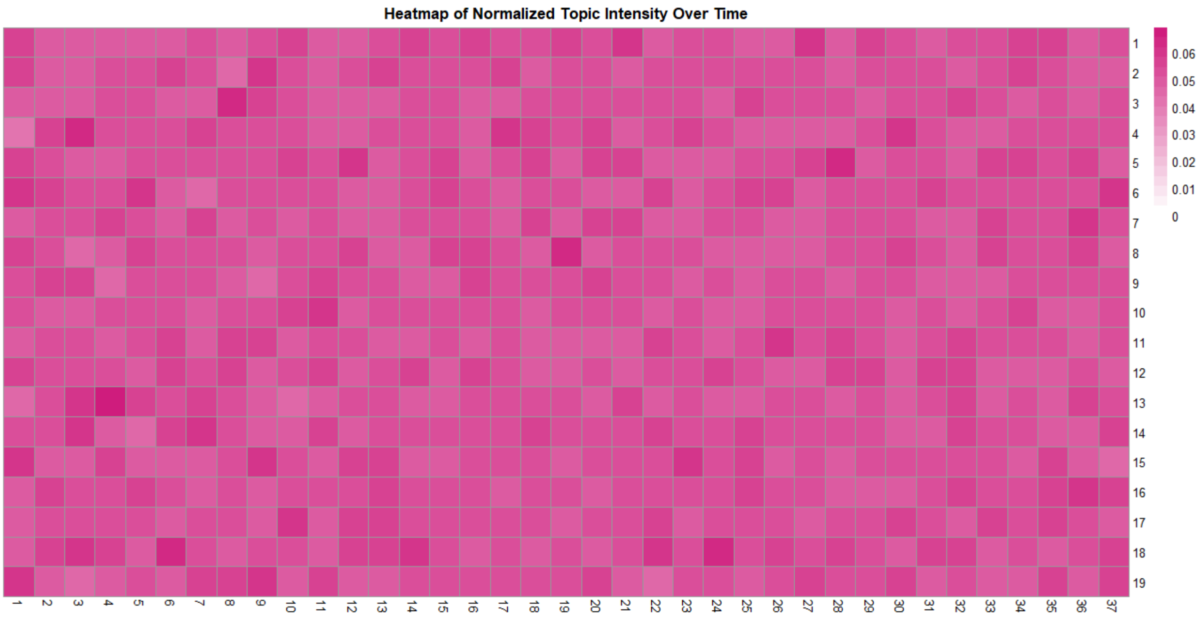

Source: Author’s Calculations

Section 5: Calculating the Sentiment Index

The present study uses the sentiment dictionary developed by Barbaglia et al. (2022) to quantify the sentiment associated with the topics. The dictionary contains sentiment scores for a wide range of terms related to economic and financial concepts. The dictionary classifies words as positive, neutral, and negative. Scores range between -1 and 1, with negative words receiving scores below zero and positive words receiving scores above zero. For example, the dictionary gives the word “crisis” a score of -0.75, while the word “devaluation” receives a score of -0.6. On the positive side, the word “favourable” is scored a 0.5, while the term “innovation” is given a 0.65. To begin the compilation process of the index, a sentiment score is assigned to every word in each topic in each quarter using the dictionary. To adjust the scores, the sentiment score is multiplied by the probability of the term appearing in the topic, which is estimated by LDA. This probability indicates how strongly a term is associated with a topic, thereby suggesting the term’s importance within its topic. Terms with higher beta values are better representatives of the topic. In the word clouds provided in the appendix (Figures A5 –A7), one can see that the larger words have a higher likelihood of belonging to a topic.

Scaling the sentiment contribution of each term according to its importance within the topic allows the most important words to have the biggest impact on the overall sentiment score of the topic. This method along with tracking topic intensity could benefit policymakers by making them aware of topics that are most likely to cause sentiment swings. Thus, policymakers would be able to address the factors contributing to negative sentiment swings directly. In the next step, the adjusted scores for each term are summed together to calculate the overall sentiment score for each topic. Finally, the topic sentiment scores are aggregated to create the sentiment index for each quarter. The final index is displayed in Figure 2 and in Figure 3 where it is plotted against the GDP growth rate. Returning to Huang et al. (2018), this method calculates relative sentiment instead of absolute sentiment because the sentiment scores are calculated relative to the topics and word usage prevalent in each quarter, which accounts for any changing dynamics of language and sentiment exhibited by Bahamian news outlets over time.

Perhaps the main benefit of this index is that it can be constructed quicker than traditional economic data, because news sentiment is based on articles published within each quarter. Although it cannot replace official statistics, the sentiment index provides a timely snapshot of economic conditions, potentially serving as a useful tool for policymakers between data releases.

Figure 2: Sentiment Index 2015Q1-2024Q1

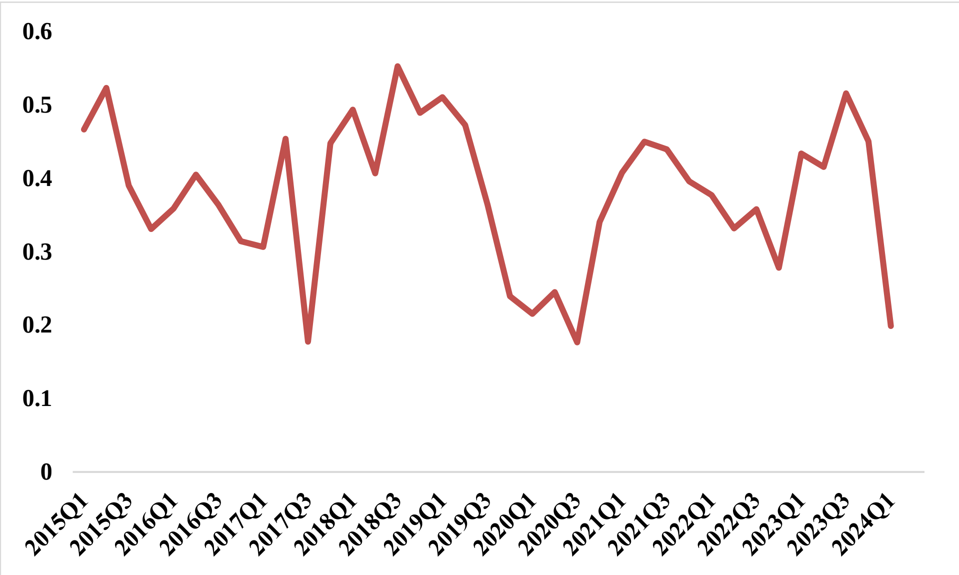

Source: Author's Calculations

Section 6: Nowcasting Exercises

The rest of this paper uses the short run ARDL and smooth local projection IRFs to evaluate the potential of the news sentiment index as a viable macroeconomic indicator in the short run.

Section 6.1: Data

This paper utilizes several macroeconomic variables in both of the short-term forecasts. The independent variable of interest is the newly constructed news sentiment index. The dependent variable is real GDP. A quarterly real GDP series is obtained from the Bahamas National Statistical Institute (BSNI), which is only available for the period 2015Q1-2023Q4. As a result, the final quarter from the sentiment index must be excluded. Undoubtedly, 36 quarters is a small sample size for a time series estimation. However, GDP is probably one of the best proxies of economic activity, so it is reasonable to compare the index to GDP to assess the viability of news sentiment as an economic indicator. In terms of data consideration, only a few predictors are included to avoid multicollinearity and any other problems that could arise when a regression has too many explanatory variables relative to the sample size.

The control variables are collected from the Central Bank of the Bahamas' Quarterly Statistical Digest (QSD). Choosing quality over the quantity when it comes to regressors, the selected variables proxy important facets of the Bahamian economy. Visitor arrivals represent the Bahamian tourism sector, which contributes greatly to the Bahamian economy. Exports proxy the country's external sector. Meanwhile, including the retail price index incorporates the influence of inflation on consumer spending and GDP. Finally, the government expenditure variable captures the impact of the fiscal sector on the economy. The variables are logged with the exception of the sentiment and retail price indexes. As is common in time series estimations, the Augmented Dickey Fuller (ADF) test determines if the variables are stationary at least at the first difference. Table B2 of the appendix presents the results of the stationarity tests, which are that all of the variables are stationary at the first difference. Other summary statistics are provided in table B3 of the appendix.

Section 6.2: Autoregressive Distributed Lag Model

It is essential to begin this exercise by testing for a cointegrating relationship among the variables. Engle and Granger (1987), Pesaran, Shin and Smith (2001) and Nkoro and Uko (2016) all emphasized the importance of using the error correction form of a regression specification if the variables are cointegrated in order to estimate both the short run and long run dynamics of a model. The present study tests for a long run relationship using the bounds testing approach, originally proposed by Pesaran, Shin and Smith (2001). However, given the relatively small sample size of 36 observations, the present study utilizes the critical values calculated by Narayan (2005) which are tailored for sample sizes ranging from 30 to 80 observations. Consequently, the bounds testing procedure is robust for studies with small samples, thereby providing, unbiased long-term estimates even if one or more regressors are potentially endogenous.

The results, presented in table B4 in the appendix, indicate no cointegration among the variables, as the F-Statistic lies below the lower bound critical values. This is not surprising based on the volatility of the news sentiment, which is evident in Figure 2. It is highly reactive to current events, political changes, market fluctuations. Such variability suggests that news sentiment does not exhibit the stable, long-term relationship with GDP that is required for cointegration. Therefore, this study proceeds with the standard short run ARDL. In addition, following Nkoro and Uko (2016), the first difference of each variable must be used to avoid spurious results. The disadvantage is that long run dynamics cannot be captured. However, this is not a concern for the present analysis, as the goal is to be able to predict current real GDP. For the regression, the Akaike Criterion (AIC) suggests one lag for each variable. Section C1 of the appendix presents the model equation.

Diagnostics

Diagnostic tests such as homoscedasticity, normality, model misspecification, autocorrelation and stability tests are conducted to ensure that the ARDL is stable and correctly specified. Table C2 and Figure C1 in the appendix report the outcomes of the tests. The residuals are normally distributed. Further, there is no evidence of serial correlation, heteroscedasticity or model misspecification. In addition, the cumulative sum of the recursive residuals (CUSUM) plots indicate that the parameters are stable.

Section 6.3: Impulse Response Function Estimation by Local Projection

This paper generates impulse response functions (IRFs) using the local projection (LP) approach to investigate the full short run dynamics between news sentiment and real GDP. First proposed by Jorda (2005), the standard LP method involves generating a sequence of separate predictive regressions at different horizons for a variable of interest. Then, a shock in one of the explanatory variables is introduced into the system (which does not have any specified underlying multivariate dynamic system), and the IRFs are estimated directly from the regression coefficients at each horizon. In this case, a shock is defined as a temporary increase in the news sentiment index that is orthogonal to current and lags of all other variables i.e. news sentiment becomes more positive. LPs have gained popularity due to numerous advantages such as robustness to misspecification, and the ability to accommodate non-linear variables or parameters.

However, this flexibility comes at an efficiency cost, and the impulse response estimator may suffer from excessive variability, especially at longer horizons (Barnichon and Brownlees, 2019; Jorda, 2023). This is why the present study utilizes the B-splines [4]smooth local projection (SLP) method by Barnichon and Brownlees (2019). The SLP variation involves fitting the impulse response coefficients as a linear combination of B-spline basis functions and then shrinking these coefficients toward a polynomial. This smoothing technique preserves the flexibility of the LP, while reducing the variance of the IRFs. It also aids with the interpretation of the IRFs. Further, the shrinkage estimator helps to prevent overfitting, which is a concern in smaller samples. See section C2 in the Appendix for more information on the model equations

To improve the reliability of the IRFs, this paper adopts the authors' usage of the Newey-West HAC estimator to create confidence intervals for the IRFs that are robust to autocorrelation and heteroscedasticity in the residuals. Further, the Akaike selection criteria decided that the optimal lag length is one lag. Figure 4 displays the impulse response function generated via the smooth local projection method.

[4] B-splines are piecewise polynomial functions defined by knots, with each basis function being composed of several polynomial segments (Barnichon and Brownlees 2019)

Section 7: Results

Figure 3: Sentiment Index vs. Real GDP Growth Rate

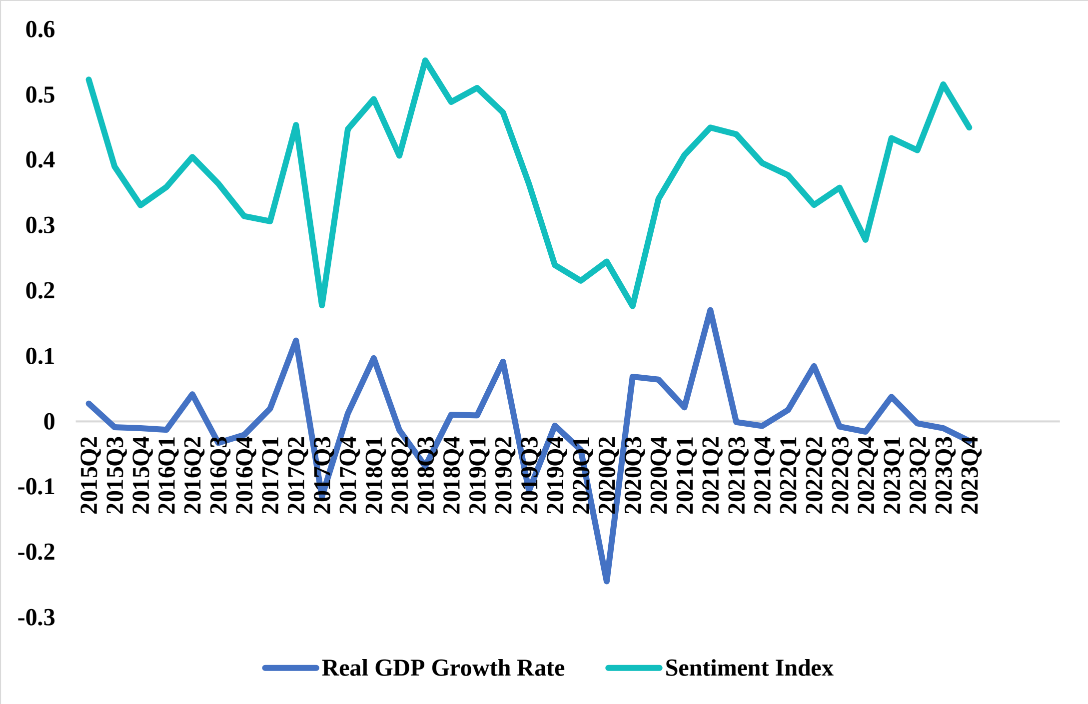

Source: Bahamas National Statistical Institute and Author's Calculation

Section 7.1: Analysis of the Sentiment Index

Figure 2 shows that the news sentiment index fluctuated significantly over the period. However, figure 3 suggests a concordance between index and the real GDP growth rate. Table B5 in the appendix substantiates this supposition, indicating a positive correlation of 0.50 between the two variables.

Several noticeable co-movements occur during economic downturns. For example, both GDP and the news sentiment index decreased in 2017Q3, when a major hurricane made landfall, and throughout 2020Q1-Q3, during the height of the COVID-19 pandemic. These co-movements occur during economic upturns as well. For example, GDP and the sentiment index rose together in 2020Q4 as the Bahamian economy began to rebound from the pandemic. This trend is also present during periods of strong GDP growth, such as in 2016Q2, 2017Q2, 2021Q2, and 2023Q1 when the news sentiment index also increased.

Demonstrating the predictive capability of the news sentiment index, there are instances where changes in the sentiment index are an early indicator of turnarounds in economic activity. For example, GDP expanded in 2019Q2, while news coverage instead became more negative. However, by 2019Q3, there was a reduction in GDP. Similarly, the index rose in 2020Q2, which contrasted a decline in GDP, but by 2020Q3, GDP grew.

Overall, the movement of the news sentiment index appears to coincide with pronounced expansions and contractions in real GDP, which suggests that it could assist with nowcasting economic activity. This variable could also potentially serve as leading indicator since changes in sentiment appear to foreshadow a future shift in GDP, even when the two variables initially appear to be moving in opposite directions. However, there are also cases where the sentiment index movements diverged from the movements in GDP, with no clear lead-lag relationship, which could suggest that short-term optimism or pessimism in news coverage may not always reflect current economic conditions such as 2023Q3, where the index rose but GDP fell. Thus, further analysis is required to determine the true nowcasting potential of the news sentiment index. In the following sections, the results of ARDL and the IRF provide more insight into the behaviour of the sentiment index relative to the behaviour of real GDP.

Section 7.2: Analysis of ARDL Results

Table C1 of the appendix presents the outcome of the ARDL model, which reveals that a one-unit increase in the news sentiment index is associated with a 16.6%[5] increase in GDP *ceteris paribus*. This coefficient is statistically significant at the 5% level. This positive relationship persists into the first lag, where a one-unit increase in news sentiment index is related to a 14.9%[6] uptick in real GDP, which is also statistically significant at the 5% level. The news sentiment index appears to have a quantifiable and statistically significant impact on real GDP, not only contemporaneously but also in the following period. These findings indicate that short-term optimism in news coverage benefits economic activity. Further, they corroborate with the previous section results, which implied that shifts in media sentiment could foreshadow economic performance at least up until the next quarter.

In addition to the news sentiment index, two control variables—visitor arrivals and government expenditure—are found to have significant influence on GDP. In particular, a one percent increase in visitor arrivals is associated with an increase of 0.04991%in the dependent variable. This coefficient is statistically significant at the 1% level, which is unsurprising since tourism contributes considerably to economic activity. However, the relationship between visitor arrivals and GDP becomes insignificant statistically after one lag, though the sign of the relationship remains positive. A review of the data indicates that the variable combines both stopover[7] and non-stopover visitors. Stopover visitors, who spend more time and money onshore, likely drive the initial boost in GDP. In contrast, day visitors and cruise passengers likely contribute less to economic activity. The mix of visitor types, combined with tourism seasonality, may explain why the impact is not sustained across quarters.

Meanwhile, a one percent increase in government expenditure is linked to a 0.1276% rise in GDP. This result is also statistically significant at the 1% level. After one lag, a one percent rise in government expenditure is associated with a statistically significant 0.1086% increase in GDP. This is also in line with expectations given that government spending is a component of GDP. In contrast, total exports and the retail price index do not show a statistically significant relationship with GDP in the short run. Given that the Bahamian economy's reliance on tourism, the impact of total exports on GDP is less significant. As for inflation, The Bahamas imports over 80% of its goods, primarily from the United States, making it a

[5] The percentage change = $(e^{0.1539} - 1)\times100\approx16.6\%$
[6] The percentage change = $(e^{0.1384} - 1)\times100 \approx14.9\%$
[7] A visitor who stays for 24 hours or more

price-taker in global markets. As a result, domestic inflation is largely imported (Branch et al., 2016), and external factors may exert a stronger influence on local inflation rates than domestic economic activity, weakening the connection between the retail price index and GDP.

Section 7.3: Analysis of Impulse Response Function Response Results

The ARDL cannot report on the relationship between GDP and the sentiment index beyond the first quarter due to the lack of cointegration. So, the IRF provides insight into how GDP behaves in the slightly longer short run, specifically 10 quarters after a change in news sentiment. Figure 4 shows that GDP sharply rises in the second quarter following the shock in the first quarter, mirroring the ARDL results. As mentioned in the introduction, consumer sentiment can link news sentiment and GDP. Thus, it can be theorized that good news boosts consumer and business confidence, which could lead to higher spending and investment, which could translate to higher GDP. In the third, fourth, fifth and sixth quarters, GDP's response remains positive but gradually decays with each quarter. By the seventh quarter, there is a noticeable decline in the GDP. The response of GDP suggests that the shock in news sentiment has a significant and immediate impact on economic activity, with effects that diminish after six quarters, or 14 months.

These results align with other empirical studies such as Nguyen and La Cava (2020) who found that the impact of a shock in news sentiment on their economic activity indicators persisted for 12 months. In Shapiro et al. (2022), the boost in consumption and output, and the decrease in inflation peaked between 12-18 months before diminishing. Interestingly, the IRFs computed by Park et al (2024) indicated that the shock in their uncertainty index, led to the longest lasting effect, which was a decline in revenue for 18 months. This result is also consistent with main findings of Barbaglia et al. (2024) which implied that that sentiment measures are significant predictors for GDP growth at horizons ranging from 30 days to a year before the official release.

However, it is unusual that the response of GDP did not include a return to baseline. This paper provides two theories. First, since news is published daily, its tone can change frequently. As consumers and businesses react to new information, their expectations and behaviours adjust accordingly as postulated by Doms and Morin (2004). Thus, it is possible that new news alters the influence of old news, potentially diminishing or completely reversing the initial expansionary impact on GDP, particularly if the new news is more pessimistic. Another theory is that Bahamian consumers overreact to the initial good news by engaging in unsustainable spending or investing, which eventually negatively affects GDP and results in a decline. Bordalo et al., (2022) argued that overreaction to good news has historically contributed to excess volatility and boom-bust cycles in stock prices, credit, investment, and even to financial crises. However, more analysis is necessary to quantify the direct impact of news sentiment on consumption and investment behaviour in the Bahamas in order to confirm these theories.

Overall, the news sentiment index could help to 'nowcast' Bahamian economic activity, especially in the absence of other economic statistics. Given that these results are in line with previous literature, it is likely that they are credible, despite the sample size limitation.

Figure 4: Response of GDP to a Shock in News Sentiment Index

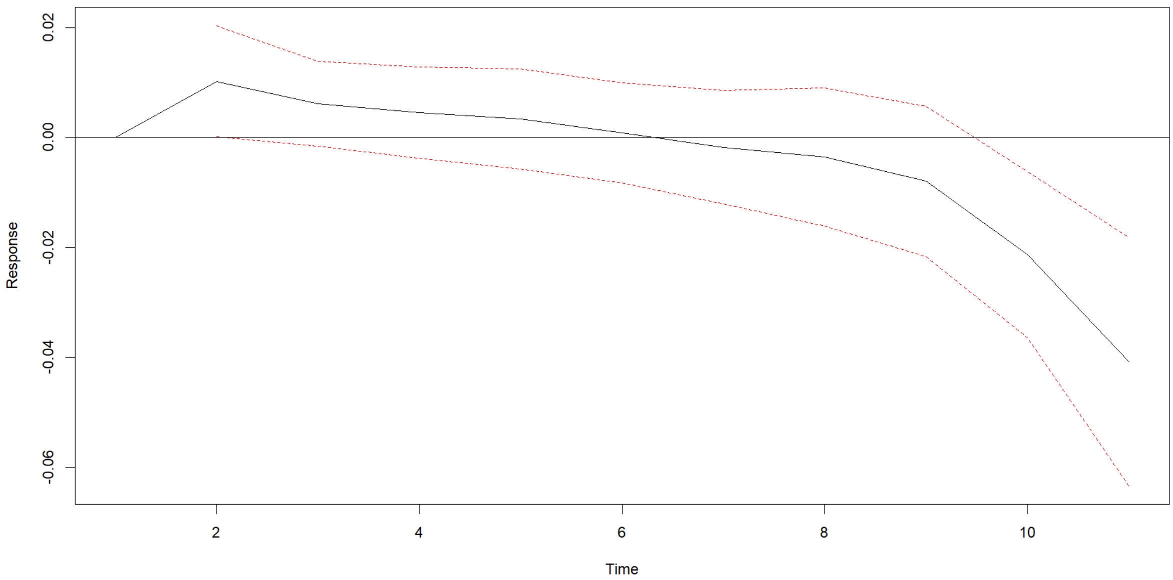

Source: Author's Calculation

Section 8: Conclusion

Monitoring the economy in real time has become a priority for economists and policymakers. The COVID-19 pandemic underscored the need for timely and accurate economic indicators. Text data such as economic and financial news has been identified as a source of high frequency data that can be utilized to forecast economic conditions in the short term. The tone or sentiment of the news has been found to be particularly useful in these 'nowcasting' exercises. As a result, various methods have been developed to convert news sentiment into a numerical economic indicator.

The present study utilizes a time varying, tone adjusted LDA, which is an unsupervised machine-learning algorithm, to organize Bahamian news articles into their main themes. A separate LDA is estimated for each quarter to account for the temporal dynamics of news data and to capture the relevant topics and by extension, relative sentiment from 2015Q1 to 2024Q1. Using a sentiment dictionary, each term within a topic is scored, and these scores are scaled according to the term's importance within the topic. The final sentiment index is composed of the aggregated scores across topics and quarters. The use of LDA and the adjusted sentiment index provide insights that could assist policymakers in identifying the key factors driving news sentiment.

The performance of the news sentiment index in predicting economic activity is assessed using the ARDL and impulse response functions. Despite the challenges posed by a limited number of observations, the results from both models indicate that the sentiment index is a strong indicator of GDP. More specifically, news sentiment becoming more positive is associated with an immediate and significant expansion in GDP that persists for up to 14 months. This finding aligns with existing literature and reinforces the potential of news sentiment as a leading economic indicator.

To conclude, this study demonstrates the advantages of incorporating real-time text data into economic forecasting models. The LDA model efficiently organizes the news articles and the news sentiment index can be compiled faster than other macro indicators. Moving forward, the development of new sentiment-based economic indicators and the refinement of older ones would likely enhance policymakers’ ability to monitor and respond to economic changes in a timely manner.

References

Aguilar, Pablo, Corinna Ghirelli, Matias Pacce, and Alberto Urtasun. “Can News Help Measure Economic Sentiment? An Application in COVID-19 Times.” *SSRN Electronic Journal*, 2020. https://doi.org/10.2139/ssrn.3673825.

Ash, Elliott, and Stephen Hansen. “Text Algorithms in Economics.” *Annual Review of Economics* 15, no. 1 (September 13, 2023): 659–88. https://doi.org/10.1146/annurev-economics-082222-074352.

Bahamas National Statistical Institute. “Advanced Estimates for Quarterly Gross Domestic Product:2023,” April 23, 2024. https://www.bahamas.gov.bs/wps/wcm/connect/723ca137-0c93-4d83-9bf9-5694bab711d4/BNSI+QUARTERLY+GDP+2023.pdf?MOD=AJPERES.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. “Measuring Economic Policy Uncertainty*.” *The Quarterly Journal of Economics* 131, no. 4 (November 1, 2016): 1593–1636. https://doi.org/10.1093/qje/qjw024.

BańBura, Marta, Domenico Giannone, and Lucrezia Reichlin. “Nowcasting.” In *The Oxford Handbook of Economic Forecasting*, edited by Michael P. Clements and David F. Hendry, 1st ed., 193–224. Oxford University Press, 2012. https://doi.org/10.1093/oxfordhb/9780195398649.013.0008.

Barbaglia, Luca , Luca Tiozzo Pezzoli, Sergio Consoli, Elisa Tosetti, and Sebastiano Manzan. “Economic Lexicon.” JRC Publications Repository, December 21, 2022. https://data.jrc.ec.europa.eu/dataset/1c054ef4-561a-464a-9077-3f6b09630da2.

Barbaglia, Luca, Sergio Consoli, and Sebastiano Manzan. “Forecasting GDP in Europe with Textual Data.” *Journal of Applied Econometrics*, January 13, 2024, jae.3027. https://doi.org/10.1002/jae.3027.

Barnichon, Regis, and Christian Brownlees. “Impulse Response Estimation by Smooth Local Projections.” *The Review of Economics and Statistics* 101, no. 3 (July 2019): 522–30. https://doi.org/10.1162/rest_a_00778.

Barsky, Robert B, and Eric R Sims. “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence.” *American Economic Review* 102, no. 4 (June 1, 2012): 1343–77. https://doi.org/10.1257/aer.102.4.1343.

Bittermann, André, and Jonas Rieger. “Finding Scientific Topics in Continuously Growing Text Corpora,” September 9, 2022. https://doi.org/10.23668/PSYCHARCHIVES.8168.

Blei, David M. “Probabilistic Topic Models.” *Communications of the ACM* 55, no. 4 (April 2012): 77–84. https://doi.org/10.1145/2133806.2133826.

Blei, David M., and John D. Lafferty. “Dynamic Topic Models.” In *Proceedings of the 23rd International Conference on Machine Learning - ICML ’06*, 113–20. Pittsburgh, Pennsylvania: ACM Press, 2006. https://doi.org/10.1145/1143844.1143859.

Blei, David M., Andrew Y Ng, and Michael I. Jordan. “Latent Dirichlet Allocation.” *Journal of Machine Learning Research* 3 (March 1, 2003): 993–1022. https://www.jmlr.org/papers/volume3/blei03a/blei03a.pdf.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. “Overreaction and Diagnostic Expectations in Macroeconomics.” *Journal of Economic Perspectives* 36, no. 3 (August 1, 2022): 223–44. https://doi.org/10.1257/jep.36.3.223.

Branch, Sharon, Yvonne Cooper, and Martiniqua Moxey. “An Empirical Analysis Of The Nexus Between Private Sector Credit, Economic Growth, Government Expenditure, Interest Rate And Inflation: Case Of The Bahamas (1989-2014),” 2016. https://api.semanticscholar.org/CorpusID:195815290.

Cao, Juan, Tian Xia, Jintao Li, Yongdong Zhang, and Sheng Tang. “A Density-Based Method for Adaptive LDA Model Selection.” *Neurocomputing* 72, no. 7–9 (March 2009): 1775–81. https://doi.org/10.1016/j.neucom.2008.06.011.

Central Bank of The Bahamas. "Quarterly Statistical Digest May 2024," May 21, 2024. https://www.centralbankbahamas.com/viewPDF/documents/2024-05-23-12-39-17-CBOB-Quarterly-Statistical-DigestMay-2024Final.pdf.

Consoli, Sergio, Luca Barbaglia, and Sebastiano Manzan. "Fine-Grained, Aspect-Based Sentiment Analysis on Economic and Financial Lexicon." *Knowledge-Based Systems* 247 (July 2022): 108781. https://doi.org/10.1016/j.knosys.2022.108781.

Deveaud, Romain, Eric SanJuan, and Patrice Bellot. "Accurate and Effective Latent Concept Modeling for Ad Hoc Information Retrieval." *Document Numérique* 17, no. 1 (April 30, 2014): 61–84. https://doi.org/10.3166/dn.17.1.61-84.

Doms, Mark E., and Norman Morin. "Consumer Sentiment, the Economy, and the News Media." *Finance and Economics Discussion Series* 2004, no. 51 (September 2004): 1–70. https://doi.org/10.17016/FEDS.2004.51.

Engle, Robert F., and C. W. J. Granger. "Co-Integration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55, no. 2 (March 1987): 251. https://doi.org/10.2307/1913236.

Gan, Jingxian, and Yong Qi. "Selection of the Optimal Number of Topics for LDA Topic Model—Taking Patent Policy Analysis as an Example." *Entropy* 23, no. 10 (October 3, 2021): 1301. https://doi.org/10.3390/e23101301.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. "Text as Data." *Journal of Economic Literature* 57, no. 3 (September 1, 2019): 535–74. https://doi.org/10.1257/jel.20181020.

Griffiths, Thomas L., and Mark Steyvers. "Finding Scientific Topics." *Proceedings of the National Academy of Sciences* 101, no. suppl_1 (April 6, 2004): 5228–35. https://doi.org/10.1073/pnas.0307752101.

Huang, Melody Y., Randall R. Rojas, and Patrick D. Convery. "News Sentiment as Leading Indicators for Recessions." arXiv, 2018. https://doi.org/10.48550/ARXIV.1805.04160.

Jett, Jennifer. "Bahamas Mega-Resort Stands Empty, Casualty of Dispute With Chinese Partners." *The New York Times*, February 2, 2016, sec. Real Estate. https://www.nytimes.com/2016/02/03/realestate/commercial/bahamas-china-baha-mar.html.

Jordà, Òscar. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95, no. 1 (February 1, 2005): 161–82. https://doi.org/10.1257/0002828053828518.

———. "Local Projections for Applied Economics." *Annual Review of Economics* 15, no. 1 (September 13, 2023): 607–31. https://doi.org/10.1146/annurev-economics-082222-065846.

Larsen, Vegard, and Leif Anders Thorsrud. "Business Cycle Narratives." *SSRN Electronic Journal*, 2018. https://doi.org/10.2139/ssrn.3130108.

Narayan, Paresh Kumar. "Fiji's Tourism Demand: The ARDL Approach to Cointegration." *Tourism Economics* 10, no. 2 (June 2004): 193–206. https://doi.org/10.5367/000000004323142425.

———. "The Saving and Investment Nexus for China: Evidence from Cointegration Tests." *Applied Economics* 37, no. 17 (September 20, 2005): 1979–90. https://doi.org/10.1080/00036840500278103.

Nkoro, Emeka, and Kelvin Uko. "Autoregressive Distributed Lag (ARDL) Cointegration Technique: Application and Interpretation." *Journal of Statistical and Econometric Methods* 5, no. 4 (2016): 63–91. https://www.scienpress.com/Upload/JSEM/Vol%205_4_3.pdf.

Nguyen, Kim, Gianni La Cava, and Reserve Bank of Australia. "Start Spreading the News: News Sentiment and Economic Activity in Australia." RBA Research Discussion Papers. Reserve Bank of Australia, December 23, 2020. https://doi.org/10.47688/rdp2020-08.

Papanikolaou, Yannis, James R Foulds , Timothy N. Rubin, and Grigorios Tsoumakas . "Dense Distributions from Sparse Samples: Improved Gibbs Sampling Parameter Estimators for LDA." *Journal of Machine Learning Research* 18 (2017): 1–58. https://jmlr.csail.mit.edu/papers/volume18/16-526/16-526.pdf.

Park, Sunghwa, Hyunsok Kim, Janghan Kwon, and Taeil Kim. "Construction of Korean Shipping Industry Risk Sentiment Index Using News Articles." *International Journal of Shipping and Transport Logistics* 1, no. 1 (2023): 1. https://doi.org/10.1504/IJSTL.2023.10050213.

Pesaran, M. Hashem, Yongcheol Shin, and Richard J. Smith. "Bounds Testing Approaches to the Analysis of Level Relationships." *Journal of Applied Econometrics* 16, no. 3 (May 2001): 289–326. https://doi.org/10.1002/jae.616.

Schofield, Alexandra, and David Mimno. "Comparing Apples to Apple: The Effects of Stemmers on Topic Models." *Transactions of the Association for Computational Linguistics* 4 (December 2016): 287–300. https://doi.org/10.1162/tacl_a_00099.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J. Wilson. "Measuring News Sentiment." *Journal of Econometrics* 228, no. 2 (June 2022): 221–43. https://doi.org/10.1016/j.jeconom.2020.07.053.

Steyvers, Tom, and Mark Griffiths. "Probabilistic Topic Models ." In *Handbook of Latent Semantic Analysis*, 1st ed. Psychology Press, 2007. https://doi.org/10.4324/9780203936399.

Valtonen, L., Saku J. Mäkinen, and Johanna Kirjavainen. "Advancing Reproducibility and Accountability of Unsupervised Machine Learning in Text Mining: Importance of Transparency in Reporting Preprocessing and Algorithm Selection." *Organizational Research Methods* 27, no. 1 (January 2024): 88–113. https://doi.org/10.1177/10944281221124947.

Zhu, Maoran, Xiaopeng Zhang, and Hongwei Wang. "A LDA Based Model for Topic Evolution: Evidence from Information Science Journals." In *Proceedings of 2016 International Conference on Modeling, Simulation and Optimization Technologies and Applications (MSOTA2016)*. Xiamen, China: Atlantis Press, 2016. https://doi.org/10.2991/msota-16.2016.12.

Appendix

Section A1: The Latent Dirichlet Allocation Model

This section provides additional information about the statistical notation of LDA. The table below summarizes the key elements.

Notation	Definition
K	Number of Topics
N_d	Number of word tokens
D_d	A collection of documents
Φ_k	Distribution of words within topic k.
d	A document
θ_d	Topic Distribution for Document d
β	Dirichlet prior hyperparameter on Φ for word-topic distribution.
α	Dirichlet prior hyperparameter on θ for topic k
w_{dn}	A word
x_{dn}	Topic Assignment for word w_{dn}

In more detail, this model uses Bayesian inference techniques to decompose the text data into its topics, as mentioned in Section 4. The process repeats for each quarter.

A corpus is a collection of D documents denoted by $D = \{d_1, . . . , d_D\}$. A document is denoted by d and is composed of N words so $d = (w_{d1}, . . . , w_{dn})$. w_{dn} is the nth word in document d. Each document also consists of K topics. The generative process of LDA is as follows:

1. For each topic k = 1, ..., K, where K is the total number of latent topics:
   - Word probability distributions are chosen from a Dirichlet distribution with hyper parameter β, denoted as: $\phi_k \sim Dir(\beta)$. Here, β is set to 50/K[8]
   - ϕ_k represents the distribution of words within topic k.
2. For each document d:
   - Topic mixtures are chosen from a Dirichlet distribution with hyperparameter α, denoted as: $\theta_d \sim Dir(\alpha)$. In this study, α is set to 1 [9]
   - θ_d represents the distribution of topics within document d.
3. For each word:
   - A topic assignment (i.e. x_{dn}) for word w_{dn} is sampled from the multinomial distribution defined by θ_d, denoted $x_{dn} \sim Mult(\theta_d)$, $x_{dn} \in \{1, . . . , K\}$. This step assigns a topic to each word in the document.
   - After determining the topic assignment, the actual word w_{dn} is drawn from the multinomial distribution corresponding to the word distribution within the assigned topic denoted as $w_{dn} \sim Mult(\phi_k)$, where k is the topic assigned in the previous step.

[8] This value was adopted from Griffiths and Steyvers (2004). It is a relatively small β, so words are more evenly spread across topics, meaning that words have a greater likelihood of belonging to multiple topics.
[9] This value was adopted from Griffiths and Steyvers (2004). It is a higher α, which means that documents are likely to contain a higher mixture of topics, without an extreme bias toward a single topic.

Section A2: Topic Selection Analysis

The following figures present the results of the diagnostic tests at various topic counts across different quarters. The methods by CaoJuan (2009) and Griffiths (2004) maximize their respective metrics, while the methods by Devaud et al. (2014) and Arun et al. (2010) minimize theirs. As shown in the figures below, the maximizing methods peak at 19 topics, whereas the minimizing methods reach their lowest point. This indicates that the model's performance likely does not improve beyond 19 topics.

Figure A1: Selection of Number of Topics: Quarter 5

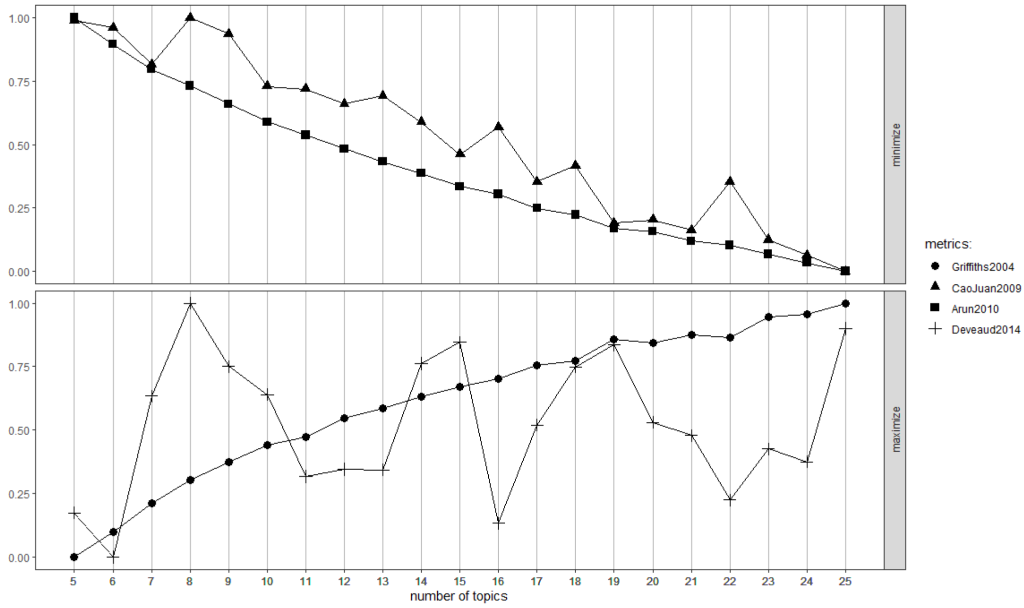

Source: Author's Calculations

Figure A2: Selection of Number of Topics: Quarter 15

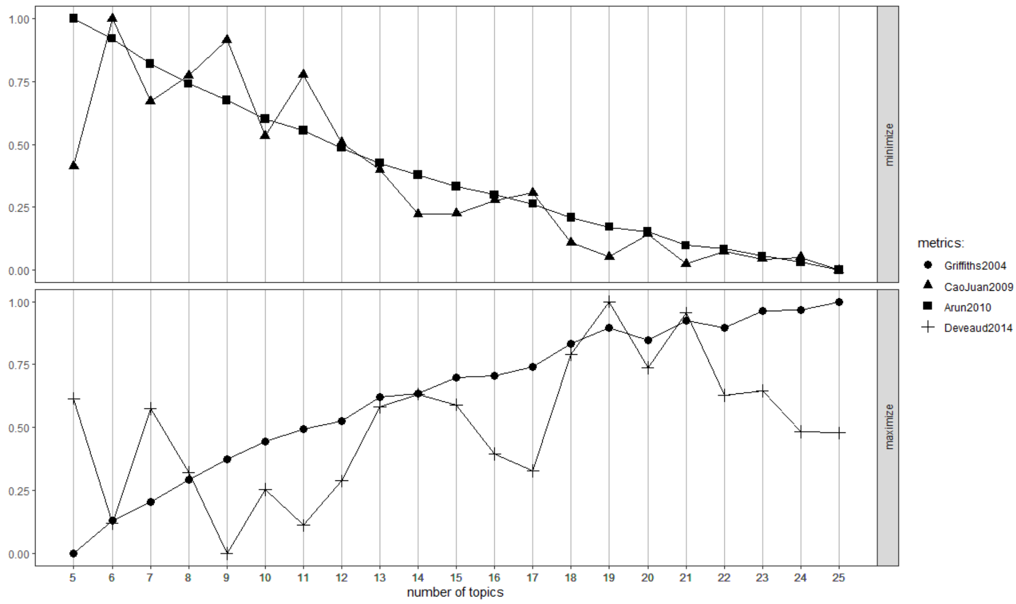

Source: Author's Calculations

Figure A3: Selection of Number of Topics: Quarter 20

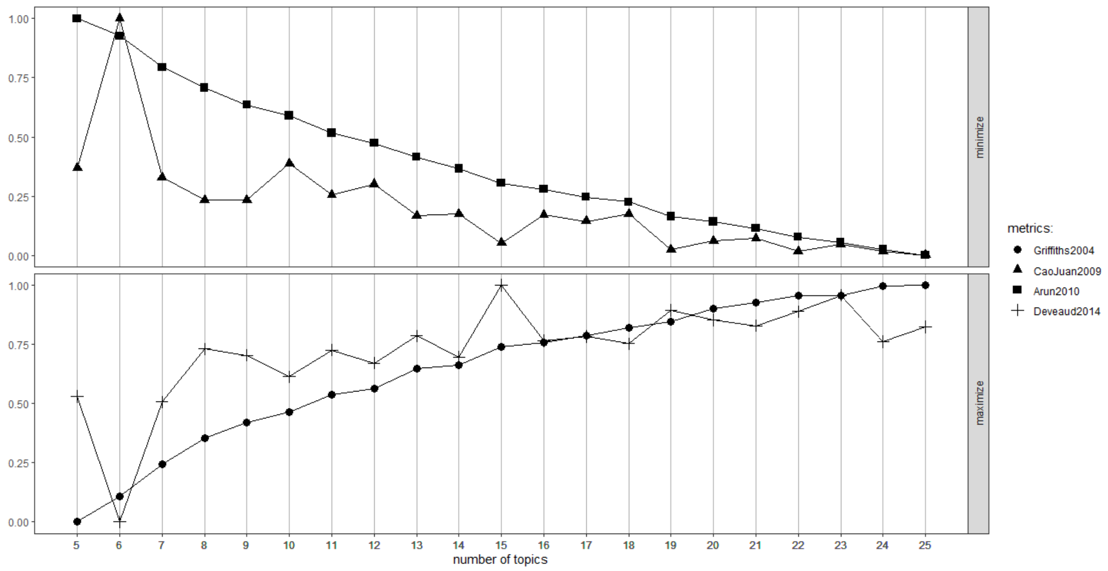

Source: Author's Calculations

Figure A4: Selection of Number of Topics: Quarter 30

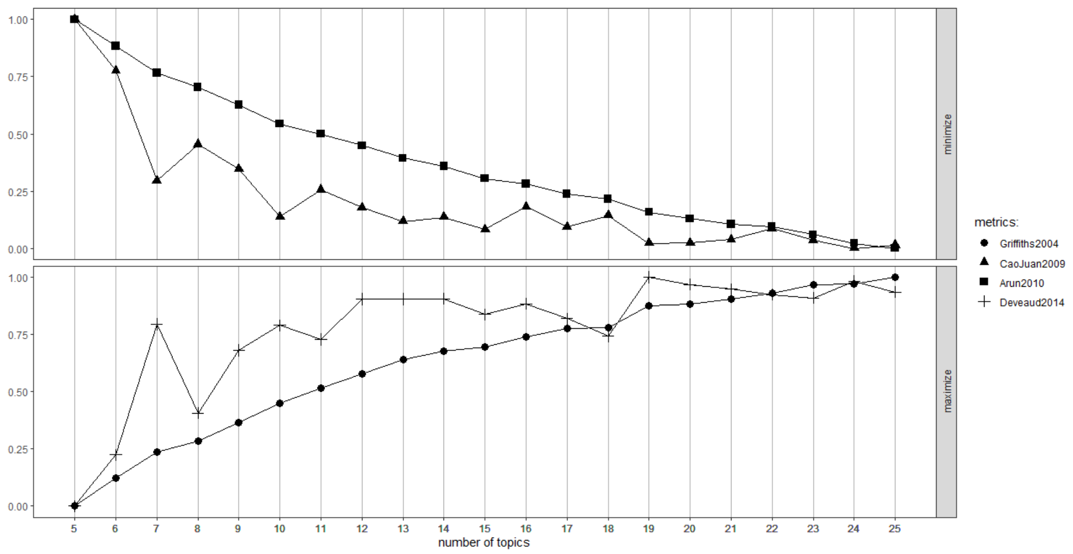

Source: Author's Calculations

Figure A5: Word Cloud for 2015Q4 Topic 13

Figure A6: Word Cloud for 2019Q3 Topic 8

Figure A7: Word Cloud for 2020Q2 Topic 11:

Table B1: Data Sources

Data	Source
News Articles	The Tribune, The Nassau Guardian, Bahamas Eyewitness News, The Bahama Journal
News Sentiment Index	Author's Calculation
Real GDP	Bahamas National Statistical Institute
Visitor Arrivals	Central Bank of the Bahamas Quarterly Statistical Digest Table 8.4
Retail Price Index (November 2014=100)	Central Bank of the Bahamas Quarterly Statistical Digest Table 8.1
Total Exports	Central Bank of the Bahamas Quarterly Statistical Digest Table 7.2
Government Expenditure	Central Bank of the Bahamas Quarterly Statistical Digest Table 5.1

Section B1: Data and Preliminary Analysis

Table B2: Data Transformation

Variable	Expected Sign	Transformation	ADF Test Statistic at Level	ADF Test Statistic at 1st Difference
GDP	N/A	Logged	0.08759757	-2.999728
Sentiment Index	+	None	-0.3564509	-3.139569
Retail Price Index	-	None	-1.345337	-2.473494
Exports	+	Logged	0.3258849	-4.884816
Government Expenditure	+	Logged	0.6111358	-11.58499
Visitor Arrivals	+	Logged	-0.06601925	-3.47789

***Critical Values:**
1%: -2.62
5%: -1.95
10%: -1.61
Source: Author's Calculations in R

Table B3: Summary Statistics

Variable	Mean	Minimum	Maximum	Median	Skewness	Standard Deviation	Kurtosis
GDP	8.000054	7.677397	8.121141	8.007938	-1.786307	0.09649274	6.188308
Sentiment Index	0.3858169	0.1764257	0.5523291	0.4000816	-0.4923529	0.09715626	2.576216
Retail Price Index	2.158555	-1.113646	6.647469	1.791896	0.4552793	1.871977	2.848812
Exports	11.70001	10.92292	12.4178	11.69259	-0.07321596	0.3233077	3.035543
Government Expenditure	13.45623	13.08135	13.9516	13.4382	0.6334001	0.2166542	2.926906
Visitor Arrivals	13.81792	8.29355	14.77811	14.27573	-2.639212	1.376592	9.504507

Source: Author's Calculations in R

Table B4: Bounds Test Result

	F - Statistic	Case	Lower Bound	Upper Bound
10% Critical Value	1.512934	3	2.508	3.763
5% Critical Value	1.512934	3	3.037	4.443
1% Critical Value	1.512934	3	4.257	6.040

Source: Author's Calculations in R

Table B5: Correlation Matrix

Real Gross Domestic Product	Real Gross Domestic Product	Sentiment Index	Visitor Arrivals	Total Exports	Government Expenditure	Retail Price Index
GDP	1.00000000	0.498339512	0.9163302	0.6646315	0.057872307	0.4779326
Sentiment Index	0.49833951	1.000000000	0.3657556	0.4920448	0.006065076	0.2576317
Visitor Arrivals	0.91633023	0.365755585	1.0000000	0.5054679	-0.185591665	0.3906340
Total Exports	0.66463152	0.492044821	0.5054679	1.0000000	0.446687929	0.6708908
Government Expenditure	0.05787231	0.006065076	-0.1855917	0.4466879	1.000000000	0.3374831
Retail Price Index	0.47793260	0.257631728	0.3906340	0.6708908	0.337483140	1.0000000

Source: Author's Calculations in R

Section C: Nowcasting Models

Section C1: The ARDL Specification

The standard ARDL specification is as follows:

$$ARDL(p,q): Y_t = \beta_0 + \sum_{i=1}^{p} \beta_i Y_{t-i} + \sum_{i=1}^{q} \delta_i X_{t-i} + \varepsilon_i$$

Lagged Y's represent the autoregressive nature, while lagged X's denote the explanatory variables. 'p' is the number of lags for the dependent variable, while 'q' is the number of lags for the independent variables.

The specification of Narayan (2004) was used to tailor this specification to the present study. Thus, the ARDL is as follows:

$$\Delta lnGDP_t = \beta_0 + \beta_1 lnGDP_{t-1} + \beta_2 SentimentIndex_{t-1} + \beta_3 \ln VisitorArrivals_{t-1} + \beta_4 \ln TotalExports_{t-1} + \beta_5 \ln GovernmentExpenditure_{t-1} + \beta_6 RetailPriceIndex_{t-1} + \varepsilon_i$$

Δ denotes the first difference operator; ε_i is the white noise residual. β_0 is the constant term. Meanwhile, β_1 through β_7 characterize the coefficients of the short run dynamics of the model.

Table C1: ARDL Results

Residuals:

Min	1Q	Median	3Q	Max
-0.05794	-0.01533	0.00362	0.01906	0.04289

Coefficients	Estimate	Std. Error	t-value	Pr(>|t|)
(Intercept)	-1.551e-03	5.148e-03	-0.301	0.766063
GDP_1	-2.776e-01	1.792e-01	-1.549	0.135545
Sentiment Index	1.539e-01	7.195e-02	2.139	0.043826 *
Visitor Arrivals	4.991e-02	5.444e-03	9.168	5.72e-09 ***
Total Exports	3.488e-02	2.348e-02	1.486	0.151599
Government Expenditure	1.276e-01	2.872e-02	4.442	0.000205 ***
Retail Price Index	-7.838e-05	5.588e-03	-0.014	0.988935
Sentiment Index_1	1.384e-01	6.276e-02	2.205	0.038225 *
Visitor Arrivals_1	1.621e-02	1.048e-02	1.547	0.136017
Total Exports_1	-1.212e-02	2.625e-02	-0.462	0.648747
Government Expenditure_1	1.086e-01	3.100e-02	3.501	0.002018 **
Retail Price Index_1	7.878e-03	5.440e-03	1.448	0.161696

Residual Standard Error: 0.0299 on 22 degrees of freedom
Multiple R-squared: 0.8968, **Adjusted R-squared**: 0.8453
F-statistic: 17.39 on 11 and 22 DF, **P-value**: 2.239e-08

Significance codes: 0 '*' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1**

Source: Author's Calculations in R

Figure C1: CUSUM Tests

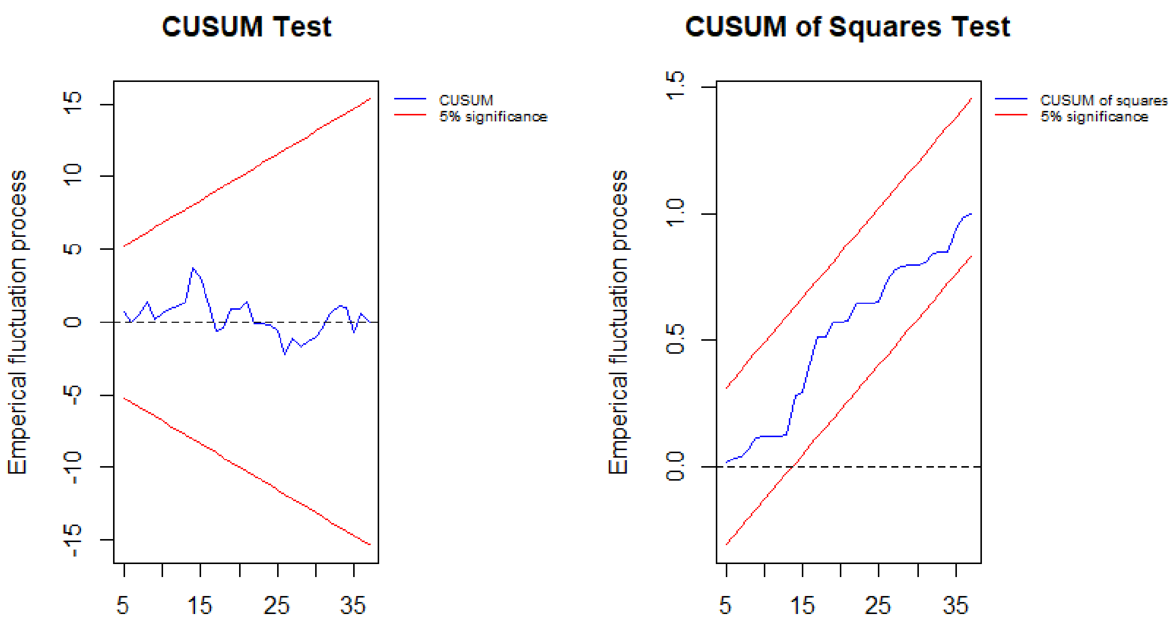

Source: Author's Calculations in R

Table C2: ARDL Diagnostics

Test	P-Value	Interpretation
Ramsay Reset Test	0.7186	No evidence of model misspecification
Breusch-Godfrey Test	0.3472	No evidence of serial correlation
ARCH LM Test	0.5904	No evidence of heteroscedasticity
Jarque-Bera Test	0.7842	Residuals are normally distributed

Source: Author's Calculations in R

Section C2: Smooth Local Projection Impulse Response Functions

This section discusses the details regarding the impulse response functions. First, it is important to discuss the specification of the standard LP IRFs in order to derive the smoothed version. Adapted from Jorda (2005), the standard IRFs can be calculated using following set of regressions:

$$GDP_{t+h} = \alpha_{(h)} + \beta_{(h)} SentimentIndex_t + \sum_{i=1}^{p} \gamma_{(h)}\, w_{it} + \varepsilon_{(h)t+h}$$

The key notation is presented in the table below:

h	The number of horizons, which is 10 quarters.
$\alpha_{(h)}$	The intercept term specific to horizon h.
$\beta_{(h)} SI_t$	The coefficient that represents the impulse response of GDP to a shock in the sentiment index.
$\sum_{i=1}^{p} \gamma_{(h)}\, w_{it}$	w_{it} denotes the control variables. $\gamma_{(h)}$ represents their coefficients up to lag p
$\varepsilon_{(h)t+h}$	The error term

Adapted from Barnichon and Brownlees (2019), the SLP IRFs can be derived by approximating, $\beta_{(h)}$ with a B-spline basis function (i.e. $\beta_{(h)} \approx \sum_{k=1}^{K} b_k B_k\,(h)$) to obtain the following specification:

$$GDP_{t+h} \approx \sum_{k=1}^{K} a_k B_k\,(h) + \sum_{k=1}^{K} b_k B_k\,(h) SentimentIndex_t + \sum_{i=1}^{p} \sum_{k=1}^{K} c_{ik} B_k(h) w_{it+}\, \varepsilon_{i,t+h}$$

The notation for the SLP estimation can be described in further detail:

$B_k(h)$	The k_{th} B-spline basis function evaluated at horizon h
a_k	The coefficient for the k_{th} B-spline basis function for the intercept term.
b_k	The coefficient for the k_{th} B-spline basis function for the shock variable SI_t
c_{ik}	The coefficient for the k_{th} B-spline basis function for the control variable w_{it} at lag i.

Figure C2: Impulse Response Function Generated Using Standard Local Projection Method

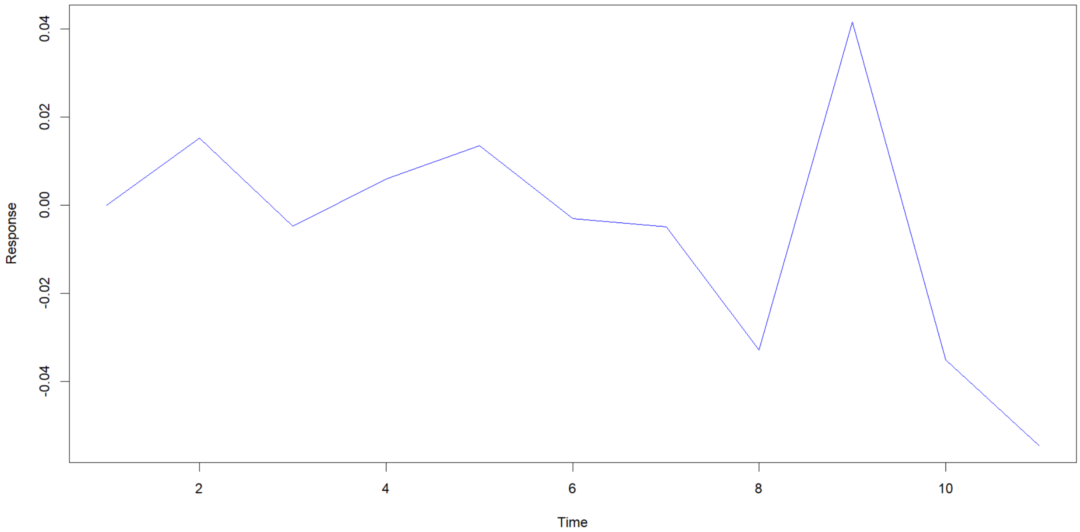

Figure C3: Impulse Response Function Generated Using Standard Local Projection versus Impulse Response Function Generated Using Smooth Local Projection

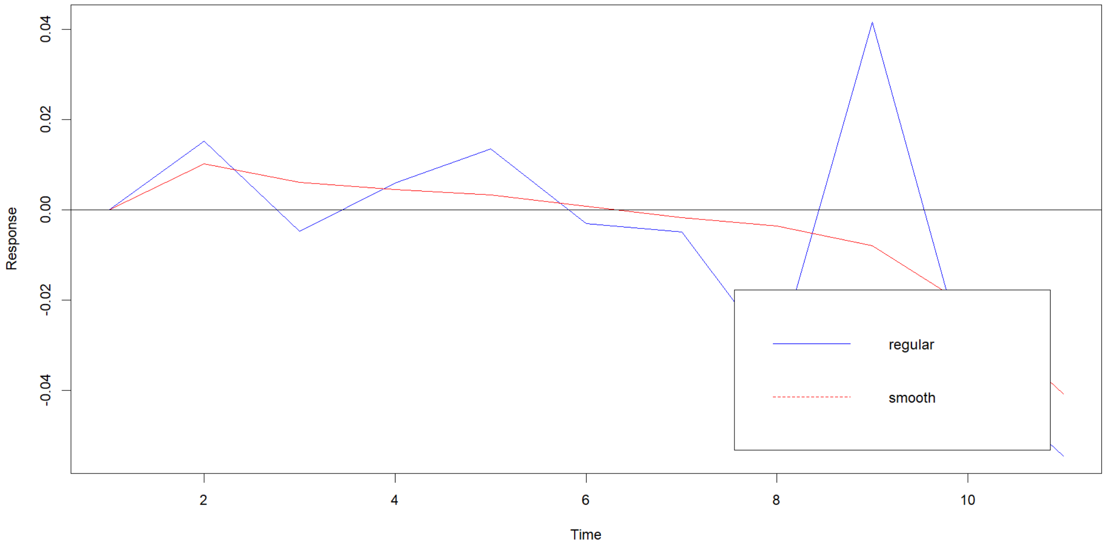

Stop Words

- and
- a
- in
- to
- for
- by
- i
- me

- my
- myself
- we
- our
- ours
- ourselves
- you
- your
- yours
- yourself
- yourselves
- he
- him
- his
- himself
- she
- her
- hers
- herself
- it
- itself
- they
- them
- their
- theirs
- themselves
- what
- which
- who
- these
- are
- was
- being
- having
- does
- doing
- ought
- im
- youre
- hes
- shes
- were
- theyre
- ive
- youve
- weve
- theyve
- id

- youd
- hed
- shed
- wed
- ill
- youll
- hell
- shell
- well
- day
- theyll
- isnt
- wasnt
- hasnt
- hadnt
- dont
- wont
- wouldnt
- shant
- mustnt
- let
- whose
- heres
- wheres
- whys
- hows
- a
- an
- the
- if
- or
- as
- until
- of
- by
- for
- about
- against
- between
- into
- below
- to
- from
- up
- out
- then
- once
- here

- there
- where
- why
- how
- both
- each
- few
- more
- other
- some
- only
- minister
- pastor
- letter
- editor
- also
- will
- million
- thousand
- billion
- hundred
- said
- through
- however
- week
- people
- year
- country
- tuesday
- wednesday
- thursday
- saturday
- sunday
- today
- evening
- night
- since
- bahama
- bahamas
- bahamian
- say
- ms
- percent
- now
- either
- despite
- notwithstand
- nothwithstanding

- nevertheless
- oppenheimer
- standalone
- eyewitness
- tribune
- guardian
- thing
- cooper
- not
- down
- john
- mary
- jane
- robert
- susan
- miller
- matheo
- smith
- davis
- philip
- michael
- jermaine
- campbell
- lincoln
- johnson
- Monday
- Tuesday
- Friday
- Saturday
- year
- years
- months
- month
- cent
- sector
- days
- week
- weeks
- millions
- part
- date
- term
- time
- whom
- this
- that
- those
- am

- is
- be
- been
- being
- have
- has
- had
- having
- do
- did
- doing
- would
- should
- could
- youre
- hes
- shes
- its
- theyre
- youve
- weve
- theyve
- id
- youd
- hed
- shed
- wed
- theyd
- ill
- youll
- hell
- shell
- well
- isnt
- arent
- werent
- havent
- doesnt
- wont
- shouldnt
- cant
- cannot
- mustnt
- lets
- thats
- whos
- heres
- theres
- whens
- wheres
- whys

- hows
- but
- if
- or
- because
- while
- at
- by
- for
- with
- during
- before
- after
- above
- below
- to
- from
-
- out
- on
- off
- over
- under
- again
- further
- when
- where
- why
- how
- all
- any
- both
- each
- more
- most
- some
- such
- no
- nor
- own
- same
- so
- than
- too
- very
- didnt
- couldnt
- whose
- whats
- although
- period

- year
- family
- third
- second
- yesterday
- weekend
- wednesday
- thursday
- saturday
- morning
- afternoon
- etc
- though
- regarding
- bahama
- bahamas
- bahamian
- can
- mr
- ms
- mrs
- dr
- island
- nevertheless
- oppenheimer
- standalone
- cooper
- chester
- carl culmer
- butler
- pinder
- johnson
- coleby- davis
- Halikitis
- ryan pinder
- arinthia komolafe
- joseph
- Diggiss
- lightbourne
- keith bell
- new providence
- island
- san salvador
- exuma
- bimini
- long island
- cat island
- ragged island
- paradise island
- harbour island
- crooked island