

# Sentiment Analysis of the Fifth District Manufacturing and Service Surveys

---

Santiago M. Pinto

**T**he Richmond Fed conducts monthly surveys of business conditions in the Fifth Federal Reserve District in order to obtain timely information about economic conditions and to provide context to data obtained from other sources. The survey instruments allow respondents to enter free-form comments. This article employs basic text analytic techniques to quantify the sentiment embodied in those survey comments.

An important portion of the information collected and received by regional Reserve Banks is communicated in an unstructured or textual form. The qualitative data conveyed through surveys, or gathered at roundtable meetings with business firms or Bank directors, are very valuable pieces of information for the Banks. This information is generally used to corroborate and provide context to other sources of data. However, the data also reflect sentiment or attitudes derived from economic conditions, a perspective that constitutes a key determinant of firms' and households' economic decisions as supported by an extensive academic literature.<sup>1</sup>

Quantifying and measuring sentiment is not straightforward. Recent development of text analytic tools, however, could be useful. Dif-

---

■ Federal Reserve Bank of Richmond; [santiago.pinto@rich.frb.org](mailto:santiago.pinto@rich.frb.org). The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

<sup>1</sup> See, for instance, Bram and Ludvigson (1998), Souleles (2004), Barsky and Sims (2012), among many others. The importance of gauging sentiment has recently been highlighted in a speech by Richmond Fed President Thomas Barkin (see Barkin 2019). Barkin not only describes his view on how confidence affects investment decisions by businesses and consumers' expenditures on big-ticket items, but he claims that these reactions have become a lot more sensitive over time.

ferent applications of text analytic techniques are becoming widespread in government agencies, academia, and the private sector as a way to uncover some of the information hidden in unstructured and textual resources. These techniques are useful not only because they could potentially quantify qualitative data, but also because they could uncover novel information hidden in unstructured and textual resources. The availability of relevant and timely economic indicators at the local level is generally limited. The development of a systematic approach that uses text analytic tools to examine and evaluate the information content of a variety of sources, including media and qualitative surveys, could offer new opportunities to better understand changes in local economic conditions and predict economic sentiment.

Using very simple text analytic tools, this article extracts and analyzes the sentiment expressed in comments provided by participants in two surveys conducted by the Richmond Fed: the Fifth District Survey of Manufacturing Activity and the Fifth District Survey of Service Sector Activity. Specifically, the article first develops a set of sentiment indicators that intend to capture the “emotions” reflected in the open-ended comments. The indicators are intended to track three categories of sentiments: negative, positive, and uncertain. Second, to evaluate the information content of the indicators, the article contrasts the sentiment measures against responses to other questions included in the surveys. This kind of exercise is meaningful because these other questions are supposed to specifically inquire about monthly changes in business conditions experienced by survey participants. Third, the article examines the evolution of the sentiment indicators over time and compares their behavior to an indicator of economic activity reported by the Richmond Fed, the manufacturing composite diffusion indices (DIs). Fourth, the article also shows that this methodology can be employed to identify the extent to which responses by individual survey participants show a systematic pattern. For instance, based on the sentiment implicit in their written comments, this approach can identify those respondents who are systematically positive, negative, or uncertain.

This approach, of course, has its limitations. As with any other method, it is subject to bias and misinterpretation, and the results should always be contrasted against other methods and data. However, the analysis of qualitative data may help enhance the predictive accuracy and corroborate the information provided by other more traditional sources.

The article is organized as follows. Section 2 briefly reviews the text analytic methodology and its application, focusing on sentiment analysis. Section 3 applies these techniques to the survey comments of

the Richmond Fed Surveys and discusses the main findings. Finally, Section 4 summarizes the conclusions of the analysis and highlights other potential applications of the present approach.

## 1. QUALITATIVE DATA COLLECTED BY THE FEDERAL RESERVE BANK OF RICHMOND

All twelve Reserve Banks have regional economics departments that collect, analyze, and publish regional and national data. These data are both quantitative (such as the unemployment rate, employment growth rate, housing prices, etc.) and qualitative, from conversations with representatives from different sectors of the local economy and from surveys. The most visible use of the analysis is to give the president of each Reserve Bank a summary of regional economic conditions, information that is later shared at the Federal Open Market Committee (FOMC) meetings and made available to policymakers, consumers, and businesses. The information collected and disseminated in this way constitutes an additional instrument to evaluate economic conditions: it not only provides context for data obtained from other sources, it is useful to confirm developing trends and understand their effect on the broader economy.<sup>2</sup>

As part of these efforts, the Richmond Fed conducts several monthly surveys that collect qualitative information on business activity. The two largest ones in terms of number of participants are the Fifth District Survey of Manufacturing Activity (the “Manufacturing Survey”) and the Fifth District Survey of Service Sector Activity (the “Service Survey”). In order to identify the factors that drive current and expected business conditions in real time, the surveys ask participants a number of questions concerning changes in various measures of activity. Most of the questions in the surveys are qualitative in nature, since respondents are only required to report whether they experienced an increase, decrease, or no change in each economic variable from the preceding month or if they expect to observe similar changes six months ahead.<sup>3</sup>

For example, participants in the Manufacturing Survey are asked, among other questions, whether employment, orders, or shipments decreased, did not change, or increased from the previous month and

---

<sup>2</sup> The article by Macheras et al. (2015) explains in more detail how and why regional economic conditions may help policymakers understand economic changes observed at the macro level.

<sup>3</sup> The use of the term “qualitative” is common in the literature to refer to directional changes rather than quantitative changes in a specific variable. The term “qualitative” is also used in the present article to refer to textual data.

how they expect those variables to change in the next six months. The Service Survey includes questions that overlap with those asked in the Manufacturing Survey (such as changes in employment, wages, and local economic conditions), in addition to a few other specific questions (such as changes in revenue and product demand). The qualitative information collected through these surveys is later aggregated and combined into several DIs.<sup>4</sup> For the Manufacturing Survey, the Richmond Fed also reports a composite DI defined as the weighted sum of three individual DIs: employment, shipments, and orders.<sup>5</sup>

The survey also allows participants to provide feedback through open-ended textual comments. The comments are not only valuable because they offer information about emerging topics and trends, but they also indicate the respondents' perceptions or sentiment regarding the surrounding economic environment during a given time period.

The present analysis uses basic text mining techniques to closely examine the survey comments submitted by the surveys' participants during the period April 2002 to December 2018.<sup>6</sup> The analysis intends to evaluate the sentiment implicit in those comments, examine how sentiment changes over time, and evaluate the connection between sentiment and participant responses to the other questions included in the survey.

## 2. TEXT ANALYTICS

### **What can text analytics do and how does it work?**

Text analytics has several different uses and applications. For instance, it can be used to find hidden connections, patterns, and models in plain language narratives or unstructured data. It might be useful to detect emerging areas of concern or interest in specific target groups. Alternatively, it could be used to find trending themes by identifying topic

---

<sup>4</sup> DIs are used and reported by various agencies and organizations, such as the BLS, the Institute of Supply Management (ISM), and the University of Michigan Surveys of Consumers. The diffusion index calculated by the Richmond Fed is simply the difference between the proportion of those that report an increase and those that report a decrease. For additional background information on the structure and information content of the surveys and DIs, see Price and Watson (2014), Waddell (2015), Pinto et al. (2015), and Lazaryan and Pinto (2017).

<sup>5</sup> The panel is unbalanced. The subset of respondents may change from one period to the next. Approximately 45 percent of 200 contacts respond to the Manufacturing Survey in a typical month. The numbers are similar for the Service Survey. Also, panel members may drop out of the survey or they may be removed because they have not responded for an extended period of time.

<sup>6</sup> The input, or corpus, to be analyzed is the entire database of comments from these two surveys.

areas that are either novel or are growing in importance, or it could be used to consistently track concepts that are generally difficult to quantify (such as risk or uncertainty). Specific text analytic tools include text clustering (the classification and grouping of documents according to similarity measures), content categorization (assignment of text documents into predefined categories and building models), concept extraction, entity extraction (identifying named text features, such as people, organizations, places, etc.), entity relation modeling (learning relations between named entities), text summarization, and sentiment analysis.

Text analytic methods generally involve four steps. The first step consists of selecting the input or sources to be analyzed, usually referred to as “corpus.” The input can be any textual data, such as open-ended questions in surveys, a collection of documents, or transcribed minutes from a meeting. The second is a preprocessing step that involves the implementation of several methods and techniques to simplify the data. The process includes the extraction and identification of individual words (usually referred to as “tokenization” of the textual document), word stemming and lemmatization, the recognition of names, entities, places, and dates, and the removal of common or “stop” words that do not provide any meaning to the text (e.g.: “the,” “at,” “in,” and “with”). Stemming is a process through which words are reduced to their roots or stems. For example, the words “fox” and “foxes” may be reduced to the root “fox.” Lemmatization also tries to group words, but the process is somewhat more complicated because it attempts to associate words according to their meanings. For example, the lemma of the words “paying,” “paid,” and “pay” is “pay.” The objective of both stemming and lemmatization is to match and group words in order to reduce the size of the data and, consequently, reduce processing time and memory. The third step is the analysis. At this stage, the goal is to extract features from the documents, define a model based on those features, and train the model with a subsample of the data. Lastly, the fourth step consists of the validation of the results from the analysis. The validation is both internal, i.e., using available data not employed to construct the model, and external, i.e., using other available data sources and methods.

### **Sentiment analysis**

Sentiment analysis uses the tools of text analytics to measure and classify the emotional content of unstructured textual data. This classification has typically been used to analyze opinions and product ratings, to inform political strategy, and used in research methods to quantify

qualitative data.<sup>7</sup> The goal of this approach is essentially to map a piece of text to a specific sentiment category, such as positive, negative, or uncertain. Different techniques are generally employed to construct this mapping. Some of them are based on predefined dictionaries (the lexical or “bag of words” approach), while others rely on machine learning algorithms. See, for example, Hansen et al. (2018). They all, however, share the general principles.

The lexical or “bag of words” approach assigns textual data to each sentiment category using a predefined dictionary or list of words typically associated with those categories. Sentiment is then determined by the frequency of words in each category found in the text. However, relying exclusively on these kinds of dictionaries may lead to errors and misinterpretations. In general, the task of classifying text according to its sentiment is a lot more complicated because the meaning of a word may depend on the context and the specific combination of words found in an expression. For instance, if a word that reflects a positive sentiment is combined with a word that has a negative connotation, then the overall sentiment becomes negative. Other factors, such as sarcasm or slang, may complicate even more the analysis based on dictionaries. The approach followed later in the article (explained in Section 3) extends the “bag of words” approach by incorporating short expressions associated with different tonalities and by implementing general linguistic rules to deal with some of the problems described above.<sup>8</sup>

---

<sup>7</sup> Recent work in economics and finance has used text analytic tools to develop various indicators of economic activity. See, among others, Nyman et al. (2018), Thorsrud (2018), and Calomiris and Mamaysky (2019).

<sup>8</sup> Provalis Research, vendor for QDA Miner and WordStat text analytic software, provides a general sentiment dictionary in a website download. The WordStat Sentiment Dictionary was created by combining negative and positive words from three dictionaries: the Harvard IV TagNeg dictionary of negative words, the Martindale Regressive Imagery dictionary, and the Pennebaker Linguistic and Word Count dictionary. The dictionary building utility program in WordStat was then used to expand the word list, generating over 9,500 negative and nearly 4,700 positive word patterns. The word lists themselves do not measure sentiment; rather, sentiment is determined by applying two linguistic rules. Negative sentiment is measured by “negative words not preceded by a negation (no, not, never) within four words in the same sentence” and “positive words preceded by a negation within four words in the same sentence.” Positive sentiment can be measured similarly but is not as predictive. Improving the accuracy of a sentiment dictionary requires additional “training” of the generic dictionary to customize for a particular domain or body of content. For additional information, see “Sentiment Dictionaries” (Provalis Research) at <https://www.provalisresearch.com/products/content-analysis-software/wordstat-dictionary/sentiment-dictionaries> (accessed November 1, 2018). Two examples of dictionaries customized for specific domains include the Loughran and McDonald financial sentiment dictionary (for more information, see Loughran and McDonald [2011] and Loughran and McDonald [2015]) and the Lexicoder Sentiment Dictionary (see Young and Soroka [2012]) for the analysis of political news. The developers of the two dictionaries took different approaches toward achieving a greater accuracy of sentiment analysis.

The exercise developed in the present article is closely related to the work by Shapiro et al. (2018), in which the authors examine sentiment embodied in the news media. They use text analytic techniques to construct sentiment indices intended to capture the opinions expressed in economic and financial newspaper articles and to determine the writer's attitude toward certain issues. To develop their indices, the paper uses a proprietary machine learning predictive model developed by a company called Kanjoya.<sup>9</sup> They next analyze the information content of these measures by examining their correlation with different indicators of business economic conditions and their predictive accuracy. They find not only a strong contemporaneous correlation between sentiment and key business cycle variables, but also that sentiment helps in forecasting inflation and the federal funds rate.

### 3. SENTIMENT ANALYSIS OF THE FIFTH DISTRICT SURVEYS OF ECONOMIC ACTIVITY

The main objective of the exercise is to construct different measures that capture the sentiment and opinions embodied in the open-ended comments offered by survey participants and to examine how sentiment changes over time. To assess the information content of these measures, I compare them to the participants' responses to other questions in the survey that are supposed to track monthly changes in economic activity.<sup>10</sup>

#### **Preliminary analysis: views of participants who write comments**

Before proceeding with the textual analysis, and to understand the limitations and scope of the methodology described in the next section, it should be noted that not all respondents choose to write comments. In fact, during the period under consideration, on average, 26 percent of survey participants in the Service Survey and 30 percent in the Manufacturing Survey offer written comments. To draw meaningful conclusions from the textual analysis of comments, it is important to understand the behavior of participants who take the time and effort

---

<sup>9</sup> See Shapiro et al. (2018) for a thorough description of their methodology.

<sup>10</sup> The present analysis should simply be regarded as an exercise that shows the potential use of text mining techniques. Applying these techniques would probably make more sense when dealing with large bodies of text rather than with the surveys mentioned above, since they only target a limited number of participants. However, even for small samples, it is still valuable to develop a methodology, using some of these techniques, that systematically and consistently examines the qualitative data collected by the Richmond Fed.

to offer such information. Specifically, is the group of participants who write comments biased in a particular direction or can this subset of participants be regarded as a representative subsample?

One way of examining the differential behavior across groups is by determining the extent to which writing comments covaries with responses to the other questions included in the surveys. To do this, I compare the behavior of the two groups by evaluating how they respond to the question on changes in current employment.<sup>11</sup> Figure 1 shows the monthly difference between the employment DIs calculated using responses from each group of survey participants (i.e., [ $DI_{\text{no comments}} - DI_{\text{comments}}$ ]) along with its HP-filtered trend (solid line). The values reported in the figure combine responses from the two surveys: Service and Manufacturing. The series do not seem to indicate systematically different behavior between the two groups until approximately October 2014. While until October 2014 the difference between DIs indicated that those who write comments assessed economic conditions more negatively than those who do not write comments (i.e., the employment DI calculated using responses from the group of participants who write comments is lower than the employment DI calculated using responses from the group who don't write comments), the difference has become negative from that time period onward.

In order to examine the extent to which this kind of behavior differs across the service and manufacturing sectors, I perform the same exercise using data from each survey separately. The results are plotted in Figure 3. The figure shows periods in which the series move together (from April 2007 until April 2012) and periods in which they behave differently (from the beginning of the sample until April 2007, and from April 2012 until the end of the sample). In those periods when the series do not coincide, the survey participants who write comments in the Service Survey tend to be relatively less optimistic about economic conditions than those who write comments in the Manufacturing Survey. However, beginning in September 2017, the pattern has changed: those who write comments in the Service Survey become increasingly pessimistic, while survey participants in the Manufacturing Survey tend to show the opposite behavior.

Overall, the latter exercises suggest that the conclusions obtained from the sentiment analysis performed on survey comments should be interpreted with caution. Specifically, the conclusions could be biased because the analysis relies on information provided by a subsample of survey participants, whose incentives to report written comments might

---

<sup>11</sup> Only the analysis that considers the employment question is reported here. Similar conclusions can be drawn by comparing responses to other survey questions.



be effectively driven by their own perceptions of economic conditions (as indicated by their responses to other questions in the survey).

### Sentiment analysis: methodology

The first step of the analysis is to preprocess the textual survey data following the steps described in Section 2.<sup>12</sup> Next, I construct different sentiment indicators by extending the lexical or “bag-of-words” methodology discussed previously. The approach involves the following steps. First, I define the set of sentiment categories  $I = \{negative, positive, uncertain\}$ , where  $i \in I$  is a representative element of this set. Second, I analyze the text and detect the list of words that belong to each of the categories based on a predefined dictionary.<sup>13</sup> Third, in addition to identifying such words, I categorize text according to the use of different short expressions that commonly reflect certain types of emotions.<sup>14</sup>

Fourth, I define several linguistic rules that take into account the context of words to assess sentiment. The idea is that sentiment is not simply determined by the frequency of words present in the dictionary. For instance, positive words are assumed to reflect positive sentiment if their meanings are not modified by the presence of other words. Specifically, positive sentiment is captured by positive words “not near” a negation (such as “no,” “not,” and “never”) or “not near” a negative word and by negative words “near” a negation or “near” another negative word.<sup>15</sup> Negative sentiment can be measured using similar rules. In this way, negative sentiment would be described by the presence of negative words not near negations or other negative words and by positive words near negations or negative words. According to this approach, comments like “*We **do not** think it is a cause for **concern**,*” “*We **have not** had a **problem** hiring entry level staff,*” or “*Customer traffic is **not bad***” would be classified as positive, and comments like “*From September 10 to mid-October, business **was not** at all **good**,*”

---

<sup>12</sup> These steps are common to most every analysis performed on textual data. This stage essentially entails the identification and removal of frequently used words that appear in a content set and do not have sentiment connotations. The removal of words with many occurrences reduces the “noise” in the subsequent sentiment analysis.

<sup>13</sup> The methodology uses the dictionary constructed by Loughran and McDonald (2011) as the starting point. The dictionary is modified and trained for the specific corpus under study.

<sup>14</sup> An explanation of the methodology, including examples from the survey comments, is described in the Appendix (see Section B).

<sup>15</sup> In the present exercise, a word is defined to be “near” another word if they are within five words of each other (before or after), in the same sentence. Some of these rules are variations of those suggested by Provalis Research.

“The market is **not as strong** at retail as last fall,” or “things are **not as good as everyone thinks**” would be categorized as negative (i.e., “not good”). Uncertainty is simply assessed by determining the presence of words or expressions generally associated with this sentiment.

Finally, I analyze the mix of positive, negative, and uncertain “words” and assess the overall sentiment embodied in the text by calculating three types of indicators. To calculate the first sentiment measure, I sum the number of survey comments (or occurrences) assigned to each sentiment category. To do this, I define a case-specific indicator function that is equal to one when a comment from a survey participant (i.e., a case) contains at least one expression that belongs to the previously defined categories (negative, positive, or uncertain), and I then sum over all the indicator functions.<sup>16</sup>

The second measure of sentiment is based on the number of words in each category showing up in the comments. According to this indicator, the sentiment of a comment would depend on the relative frequency of words. Compared to the previous measure, this one reflects more accurately differences in the intensity of each sentiment expressed in the textual data. However, it does not contain information about the spread of the sentiment among respondents. In other words, it could be possible for a few comments to drive the sentiment in a specific time period if those comments include many words associated with the respective categories.

For the third measure, I construct an indicator that also uses the frequency of words in each category but normalized by the total number of words (the values are expressed as a rate every 10,000 words). The results of the analysis are summarized in the following sections.

## Results of the analysis

### *Sentiment and responses to questions on business activity*

In this section, I determine the extent to which the sentiment embodied in the written comments is associated with changes in business conditions. To establish this relationship, I compare the sentiment of the comments to the responses offered by survey participants to other questions included in the Fifth District Surveys. As mentioned earlier, these questions ask participants to determine if a specific variable has changed from the previous month (current changes) or is expected to change in the next six months (expected changes). The set of possible

---

<sup>16</sup> A specific comment may be assigned to more than one category.

responses  $J = \{1, 2, 3\}$ , where “(1)” is decrease, “(2)” is remain unchanged, and “(3)” is increase, and  $j \in J$  is a representative response from set  $J$ .

In the first place, I evaluate the sentiment of survey comments in conjunction with the responses to the question on current changes in employment.<sup>17</sup> Figure 2 summarizes the association between sentiment categories and responses to the employment question. The tables on the left (top and bottom) are constructed by counting the cases (or respondents) in each sentiment group  $i$  who respond  $j$  to changes in employment. The table on the top left reports the column percentages, i.e., the percentage of cases in each category  $i$  as a proportion of those who respond  $j$ . The table on the bottom left shows the row percentages, i.e., the number of cases that respond  $j = 1, 2, 3$ , as a proportion of cases in each category  $i$ . The tables in the middle show similar percentages constructed using the frequency of words, and the table on the right uses the frequency of words as a proportion of total words.

Consider the top left table. The largest percentage of cases in the negative category is observed when the response is (1) or “decrease,” with 59 percent, and then smallest when the response is (3) or “increased,” with 47 percent. When the response is (2) or “remain unchanged,” the percentage is 51, in between the other two. For the positive category, response (3) has the largest percentage and response (1) has the smallest. For the uncertainty category, the maximum is reached when the response is (2). The table on the bottom left shows that the percentage of those who report (1) is highest for the negative category (19 percent), the percentage of those who report (2) is highest for the uncertainty category (69 percent), and the percentage who report (3) is highest for the positive category (22 percent).

The tables constructed using word frequencies, both tables in the middle and the table on the right, show identical results.<sup>18</sup> Finally, Figure 4 in Appendix A shows the results from a similar analysis performed separately for the Manufacturing and Service Surveys.<sup>19</sup> From that table, it can be concluded that the sentiment indicators accurately reflect the opinion of those participating in the two separate surveys.

In general, the tables suggest that the sentiment indicators based on textual data accurately reflect participants’ perceptions about eco-

---

<sup>17</sup> The question about current changes in employment is common to both the Manufacturing and Service Surveys. The present analysis combines the information from both surveys in order to work with a larger sample size.

<sup>18</sup> Similar conclusions are obtained using expected changes in employment.

<sup>19</sup> Only the tables using frequency of words (rate per 10,000) are reported in Figure 4 in Appendix A. The tables calculated using case occurrences and frequency of words show the same conclusions.

conomic conditions. In other words, the information offered by these indicators seems to be consistent with other information conveyed by survey participants, in this case, the information revealed by their responses to the question that asks about changes in employment.

Finally, I examine the correspondence between the sentiment categories and other questions included in the surveys, such as current changes in: (i) local economic conditions (Figure 5 in Appendix A; the tables are constructed using data from the Manufacturing and Service Surveys since this question is common to both), (ii) shipments (Figure 6; data from the Manufacturing Survey), (iii) orders (Figure 7; data from the Manufacturing Survey), (iv) demand (Figure 8; data from the Service Survey), and (v) revenues (Figure 9; data from the Service Survey). The results confirm the conclusions from the previous analysis that compares sentiment and employment changes and further validate the measures of sentiment introduced earlier.

### *Changes in sentiment by month*

Figures 10, 11, and 12 display the monthly evolution of the measures of sentiment. Figure 10 shows the changes in the negative, positive, and uncertain indicators, calculated as the number of cases or respondents assigned to each sentiment category. The series reported in the graph are simply the percentage of cases in each category. Figure 11 shows the evolution of sentiment indicators that include the frequency of words in each category. The series, as before, are expressed as the percentage of words in each category at each period of time. Finally, Figure 12 shows the frequency of words in each category, normalized by the total number of words in each period (the numbers are expressed as a rate per 10,000 words). All figures include the series' twelve-month moving averages (solid lines).

The following observations are worth pointing out from the graphs. First, the behavior of all the sentiment indicators is similar in all three figures. Moreover, the category representing negative sentiment is relatively more important than the other two categories. However, the value of the information offered by these sentiment indicators is not determined by the level of such measures but by how these measures change in time, reflecting changing views and perceptions about the evolution of the economy.

Second, in all cases, the negative sentiment indicator reaches its maximum (within the sample considered in the analysis) at the end of 2008, and declines thereafter until April 2010. This series reaches a new peak in the second half of 2013 and later steadily declines until August 2017. Since August 2017, negative sentiment has been increasing.

Third, the series that reflect positive sentiment evolve in the opposite way. In fact, the correlations between the negative and positive sentiment indicators are -0.71, -0.89, and -0.53 in Figures 10, 11, and 12, respectively.<sup>20</sup>

Fourth, the indicator of uncertainty shows a somewhat different behavior. The uncertainty measure rises prior to the Great Recession, reaching a peak in the middle of 2007. After a brief decline it rises again reaching another peak at the end of 2012. Since then, the indicator has been declining, except for a short period of time from mid-2016 to approximately September 2017 in which it slightly increased.<sup>21</sup>

Next, I construct a sentiment indicator that aggregates the individual information described above. Specifically, the sentiment indicator is defined as the difference between negative and positive sentiment, i.e., [*Negative* – *Positive*]. This means that higher values of this indicator would be associated with higher overall negative perceptions and views about the economy. The evolution of this indicator is depicted in Figure 15.<sup>22</sup> A striking feature of this series is that negative sentiment has been steadily increasing since mid 2017, reaching in December 2018 similar levels as those observed during late 2012 and the beginning of 2013.

It is likely that certain factors affect and drive sentiment differently in the manufacturing and service sectors. I therefore evaluate the extent to which the sentiments associated with the comments included in the two surveys, the Manufacturing (M) and Service (S) Surveys, differ. Figures 17, 18, and 19 display the evolution of the sentiment indicators constructed using frequency of words (rate per 10,000 words) for each survey. The correlation between each sentiment indicator across surveys is positive but low (0.06 for negative sentiment, 0.04 for positive sentiment, and 0.17 for the uncertainty category), suggesting there could be factors affecting sentiment in each sector differently.

Finally, I calculate the sentiment indicator introduced earlier ([*Negative* – *Positive*] using frequency of words normalized by the total number of words in each period), but only for the Manufacturing Survey, and I compare the evolution of this indicator to the composite DI described in Section 1. The series are plotted in Figure 16. The

---

<sup>20</sup> The entire correlation matrix is shown in Figure 14 in Appendix A.

<sup>21</sup> An enlarged version of the series showing the behavior of the uncertainty indicator is shown in Figure 13.

<sup>22</sup> It should be considered that, as mentioned earlier, negative words tend to be more preponderant in comments than positive words. Also, changes in the positive and negative sentiment indicators may individually offer valuable information, each one correlated with different set of variables. Future work will evaluate the information content of each one of the sentiment series.

left axis indicates the units of the manufacturing sentiment indicator, and the right axis the units of the composite index.<sup>23</sup> The series, as expected, have a negative correlation (the correlation between the two (smoothed) series is -0.51). However, it is interesting to note that since approximately October 2017 both series have been increasing.<sup>24</sup> This means that during this period both negative sentiment, and favorable business conditions, captured by the level of the composite DI, have been rising. A similar behavior is only briefly observed in 2004, at least during the sample period considered in the present analysis.

Understanding the factors driving the behavior of the series is, of course, crucial in order to make sense of the conveyed information. A complete investigation is relegated for future research. However, by performing a very preliminary analysis, I was able to identify a positive association between stock market volatility and our indicator of negative sentiment.<sup>25</sup> Specifically, the correlation between the Chicago Board Options Exchange Market Volatility Index (VIX) and the negative sentiment indicator (smoothed) series is 0.44 during the sample period considered in the analysis.<sup>26</sup> The series are plotted in Figure 20.<sup>27</sup>

Now, a final comment regarding the extent to which a methodology like the one developed in this paper could help Reserve Banks in their efforts to evaluate economic conditions. The alignment of the information provided by the sentiment indicator with other qualitative measures, such as the composite DI, would help confirm the Banks' view about economic conditions. It should not be interpreted, how-

---

<sup>23</sup> The range of the DI is  $[-100, 100]$ .

<sup>24</sup> The correlation between the (smoothed) series is 0.84 during the period October 2017 to December 2018.

<sup>25</sup> Note that, in principle, the series are supposed to capture changes in sentiment and economic conditions in the Fifth District. A thorough analysis would require the identification of regional and national factors associated with the evolution of those variables. The work by Lazaryan and Pinto (2017), for instance, studies the extent to which the composite DI is associated with regional and national economic variables. A similar analysis could be performed using the negative sentiment indicator developed in this paper.

<sup>26</sup> The VIX indicator is constructed using a number of options included in the S&P 500 index and is supposed to capture the stock market's expectation of volatility over the next thirty days. While the correlation between the VIX and composite diffusion index (smoothed) series during the sample period under consideration is -0.68, the correlation has become positive since the beginning of 2017.

<sup>27</sup> In Pinto et al. (forthcoming), we construct a measure of uncertainty and apply the methodology using data from the Survey of Consumers conducted by the University of Michigan. While the correlation between our measure of uncertainty and consumer confidence (measured by the Index of Consumer Sentiment) is generally negative, as expected, they both tend to rise during the period 2009–14. To some extent, such behavior is similar to the one highlighted above when comparing the evolution of the negative sentiment indicator and the composite DI (even though such behavior is observed at different periods).

ever, that when these indicators move in opposite directions (providing perhaps conflicting evidence about the state of the economy) that the methodology is flawed. In fact, these kinds of scenarios could simply reveal the fact that sentiment gives us different information, not captured by other data, and further exploration would be necessary. The sentiment indicator, as a result, is used in this context as a way to corroborate information obtained from other qualitative assessments.

### *Sentiment and survey respondents*

A similar analysis can be carried out to identify respondents who systematically show a negative, positive, or uncertain sentiment. Note that the surveys conducted by the Richmond Fed have a panel structure. A list of contacts, developed throughout the years and representative of the Fifth District industry composition, receives online surveys every month. Using this panel of respondents, the methodology can determine the extent to which some contacts are systematically more pessimistic or optimistic than others. Understanding the systematic behavior of individual participants and identifying those contacts who consistently express a specific sentiment (positive, negative, or uncertain) would provide a much more accurate assessment and interpretation of the monthly responses by correcting any bias in the results due to sample selection. As an illustration, Figures 21, 22, and 23 list contacts, in decreasing order, according to the sentiment generally communicated through their survey comments.<sup>28</sup>

## 4. CONCLUSIONS

The present article illustrates the use of basic text analytic tools by evaluating the sentiment of survey comments collected by two surveys conducted by the Richmond Fed: the Fifth District Manufacturing and Service Surveys. First, the article constructs several indicators that intend to capture the sentiment embodied in the open-ended comments written by survey participants. Second, in order to evaluate the information content of these indicators, the article contrasts the sentiment measures against responses to other survey questions. This exercise is useful since the other survey questions are meant to specifically track monthly changes in business conditions experienced by survey participants. Finally, the article analyzes the evolution of the sentiment indi-

---

<sup>28</sup> I have carried out similar sentiment analysis by industry NAICS code and state. However, due to small sample sizes, the conclusions tend to be very imprecise.

cators and compares their behavior to an indicator of economic activity reported by the Richmond Fed, the composite DI.

Sentiment as measured in the paper (defined as the difference between negative and positive sentiment) generally aligns well with other assessment measures of qualitative data, such as the composite DI. However, there are instances in which these measures convey conflicting information. For example, the sentiment indicator and the composite DI have both been increasing since approximately October 2017. Such behavior has only been briefly observed in 2004.

The fact that sentiment might not fully align with other assessments does not necessarily imply that the methodology is flawed. It could simply mean that sentiment is capturing different information. In this way, the sentiment indicator could be used as a tool to corroborate other information collected by the Bank. When sentiment and diffusion indices head in opposite directions, for example, we would be less confident about what the qualitative surveys are telling us, requiring further exploration.

Different factors could potentially play a role in explaining the behavior of sentiment. While a thorough investigation of such determinants is beyond the scope of the present paper, a preliminary analysis allows us to identify a positive correlation between stock market volatility and the negative sentiment indicator.

It should be emphasized that the present exercise is simply a first attempt to evaluate sentiment in survey comments. A more rigorous analysis is definitely required in order to apply this method for other purposes, such as assessing the level of uncertainty in the economy or drawing conclusions about individuals' expectations. However, the preliminary results indicate that this kind of analysis is promising.

There are many other potential applications of text analytics. Some of these applications are meaningful not only to extract information from the surveys conducted by the Richmond Fed, but also to gain insights from the rest of the qualitative data communicated to the Bank. For instance, these tools could be used to uncover recurrent and emerging issues, identify trends, or consistently track the evolution of certain topics (such as "tariffs," "labor market," "inflation," etc.). The use of text mining techniques by regional Reserve Banks is not as widespread as in other sectors of the economy. However, regional Reserve Banks can definitely benefit from these methods both in academic research and policymaking. Unstructured data provide an additional source of information that, jointly with other data collected by the Banks, could offer a more complete description and understanding of the changes taking place in the economy.

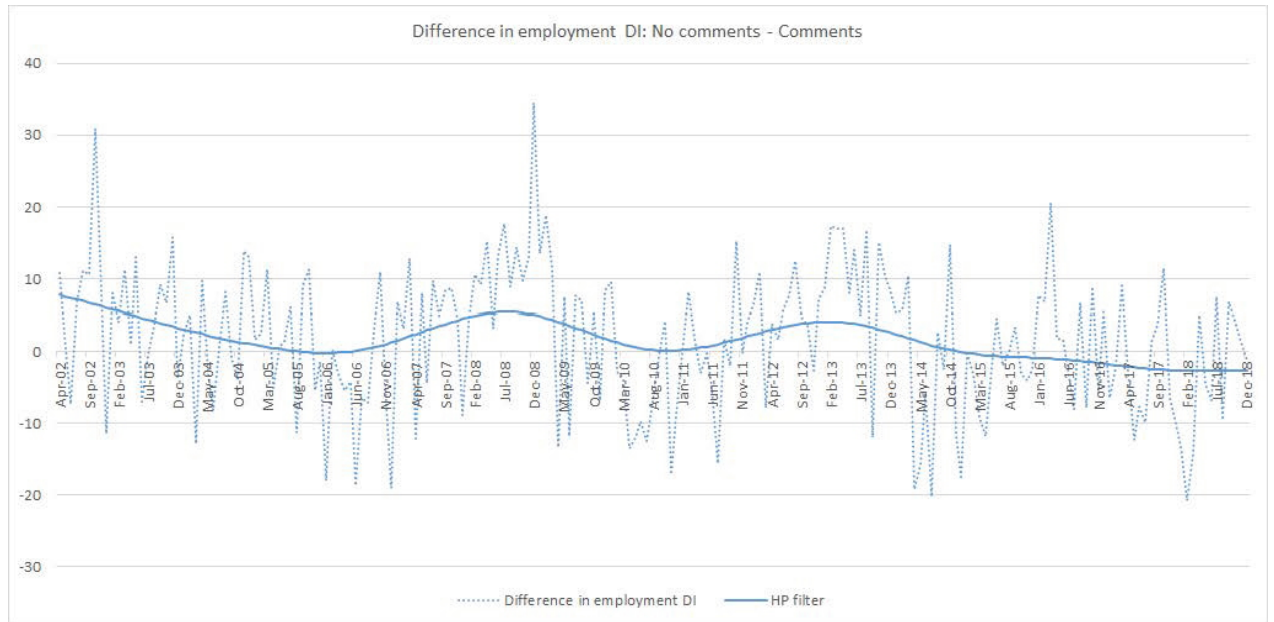


---

**APPENDIX: APPENDIX A**

**Figure 1 Difference in Employment DI: No Comments vs. Comments**

---



---

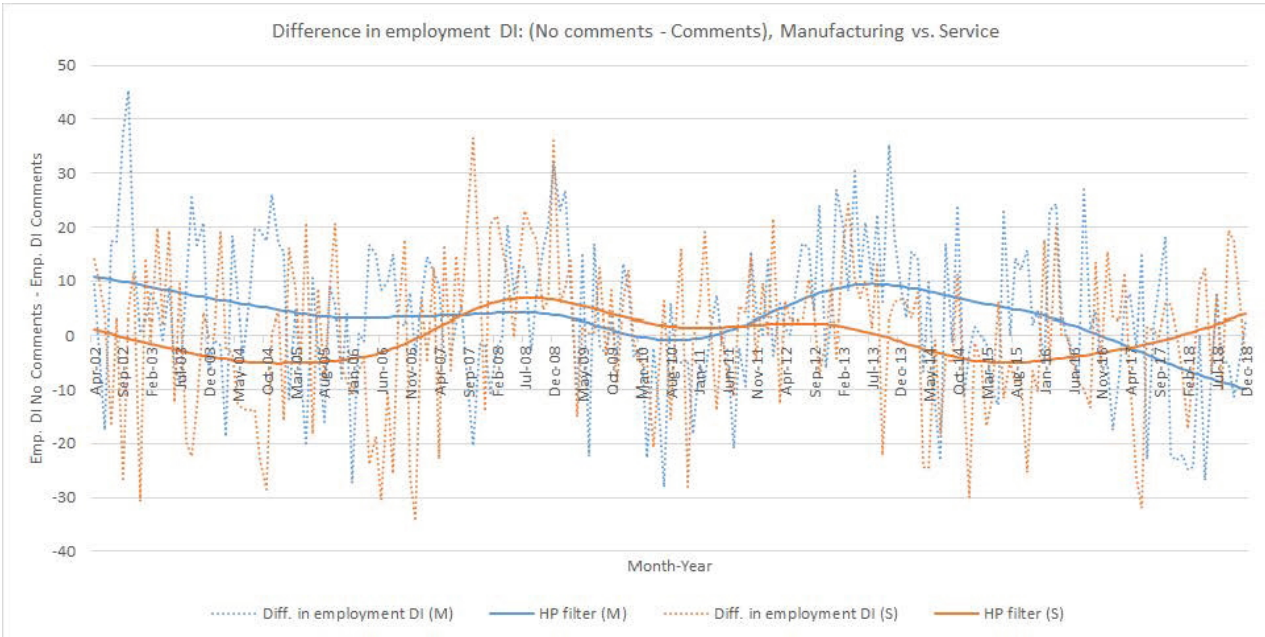
Notes: The figure shows the monthly difference between the employment DI calculated for those who don't submit written comments and those who submit written comments ( $DI_{no\ comments} - DI_{comments}$ ). The information used to calculate the DIs includes responses from both the Manufacturing and Service Surveys.

**Figure 2 Sentiment and Changes in Current Employment**

Case occurrence (column)				Frequency (column)				Frequency (rate per 10,000 words)			
SENTIMENT	1	2	3	SENTIMENT	1	2	3	SENTIMENT	1	2	3
NEGATIVE	59%	51%	47%	NEGATIVE	67%	59%	52%	NEGATIVE	761.01	632.08	537.80
POSITIVE	30%	34%	41%	POSITIVE	26%	32%	40%	POSITIVE	297.61	342.75	415.66
UNCERTAINTY	11%	14%	12%	UNCERTAINTY	7%	9%	8%	UNCERTAINTY	81.30	98.45	80.55
Case occurrence (row)				Frequency (row)							
SENTIMENT	1	2	3	SENTIMENT	1	2	3				
NEGATIVE	19%	64%	17%	NEGATIVE	19%	65%	16%				
POSITIVE	14%	64%	22%	POSITIVE	13%	64%	23%				
UNCERTAINTY	14%	69%	17%	UNCERTAINTY	14%	70%	17%				

Notes: (1) decrease, (2) no change, (3) increase. The table is constructed using the combined data from the Manufacturing and Service Surveys.

**Figure 3 Difference in Employment DI: No Comments vs. Comments – Manufacturing and Service Surveys**



Notes: The figure shows the monthly difference between the employment DI calculated for those who don't submit written comments and those who do submit written comments ( $DI_{no\ comments} - DI_{comments}$ ). M: Manufacturing, S: Service.

**Figure 4 Sentiment and Changes in Current Employment: Manufacturing and Service Surveys**

Frequency (rate per 10,000 words)						
	Service			Manufacturing		
SENTIMENT	1	2	3	1	2	3
<b>NEGATIVE</b>	498.94	420.77	347.33	562.27	491.96	404.08
<b>POSITIVE</b>	275.48	312.78	361.36	239.47	254.05	294.35
<b>UNCERTAINTY</b>	20.57	32.87	32.74	22.62	26.98	24.44

Notes: (1) decrease, (2) no change, (3) increase.

**Figure 5 Sentiment and Changes in Current Local Economic Conditions**

Case occurrence (column)				Frequency (column)				Frequency (rate per 10,000 words)			
SENTIMENT	1	2	3	SENTIMENT	1	2	3	SENTIMENT	1	2	3
NEGATIVE	55%	51%	44%	NEGATIVE	66%	59%	48%	NEGATIVE	727.61	640.56	481.92
POSITIVE	29%	35%	42%	POSITIVE	24%	33%	43%	POSITIVE	267.82	361.80	437.34
UNCERTAINTY	16%	14%	13%	UNCERTAINTY	10%	8%	9%	UNCERTAINTY	107.86	89.00	87.28
Case occurrence (row)				Frequency (row)							
SENTIMENT	1	2	3	SENTIMENT	1	2	3				
NEGATIVE	26%	49%	26%	NEGATIVE	28%	49%	24%				
POSITIVE	19%	47%	34%	POSITIVE	17%	46%	36%				
UNCERTAINTY	26%	46%	28%	UNCERTAINTY	27%	45%	28%				

Notes: (1) decrease, (2) no change, (3) increase. The values are calculated using the combined data from the Manufacturing and Service Surveys.

**Figure 6 Sentiment and Changes in Current Shipments**

Case occurrence (column)				Frequency (column)				Frequency (rate per 10,000 words)			
SENTIMENT	1	2	3	SENTIMENT	1	2	3	SENTIMENT	1	2	3
NEGATIVE	56%	55%	48%	NEGATIVE	67%	64%	56%	NEGATIVE	719.60	724.79	571.44
POSITIVE	30%	32%	39%	POSITIVE	25%	28%	36%	POSITIVE	267.07	321.24	370.54
UNCERTAINTY	14%	12%	13%	UNCERTAINTY	9%	8%	8%	UNCERTAINTY	94.65	91.08	82.64
Case occurrence (row)				Frequency (row)							
SENTIMENT	1	2	3	SENTIMENT	1	2	3				
NEGATIVE	36%	38%	27%	NEGATIVE	36%	36%	27%				
POSITIVE	30%	35%	34%	POSITIVE	29%	34%	37%				
UNCERTAINTY	36%	34%	29%	UNCERTAINTY	36%	34%	30%				

Notes: (1) decrease, (2) no change, (3) increase. The values are calculated using data from the Manufacturing Survey.

**Figure 7 Sentiment and Changes in Current Orders**

Case occurrence (column)				Frequency (column)				Frequency (rate per 10,000 words)			
SENTIMENT	1	2	3	SENTIMENT	1	2	3	SENTIMENT	1	2	3
NEGATIVE	56%	56%	48%	NEGATIVE	67%	65%	55%	NEGATIVE	718.98	743.42	557.75
POSITIVE	30%	32%	39%	POSITIVE	24%	28%	37%	POSITIVE	261.47	318.02	382.36
UNCERTAINTY	14%	12%	13%	UNCERTAINTY	9%	8%	8%	UNCERTAINTY	98.02	87.69	82.26
Case occurrence (row)				Frequency (row)							
SENTIMENT	1	2	3	SENTIMENT	1	2	3				
NEGATIVE	38%	35%	27%	NEGATIVE	38%	35%	27%				
POSITIVE	32%	33%	36%	POSITIVE	29%	31%	39%				
UNCERTAINTY	38%	31%	31%	UNCERTAINTY	39%	31%	30%				

Notes: (1) decrease, (2) no change, (3) increase. The values are calculated using data from the Manufacturing Survey.

**Figure 8 Sentiment and Changes in Current Demand**

Case occurrence (column)				Frequency (column)				Frequency (rate per 10,000 words)			
SENTIMENT	1	2	3	SENTIMENT	1	2	3	SENTIMENT	1	2	3
NEGATIVE	54%	50%	44%	NEGATIVE	61%	58%	47%	NEGATIVE	713.74	629.65	475.69
POSITIVE	32%	36%	43%	POSITIVE	30%	33%	45%	POSITIVE	348.18	360.65	459.93
UNCERTAINTY	15%	14%	13%	UNCERTAINTY	9%	9%	8%	UNCERTAINTY	108.42	96.97	85.82
Case occurrence (row)				Frequency (row)							
SENTIMENT	1	2	3	SENTIMENT	1	2	3				
NEGATIVE	23%	43%	34%	NEGATIVE	24%	44%	32%				
POSITIVE	18%	40%	43%	POSITIVE	17%	37%	45%				
UNCERTAINTY	22%	43%	36%	UNCERTAINTY	23%	42%	35%				


Notes: (1) decrease, (2) no change, (3) increase. The values are calculated using data from the Service Survey.

**Figure 9 Sentiment and Changes in Current Revenues**


Case occurrence (column)				Frequency (column)				Frequency (rate per 10,000 words)			
SENTIMENT	1	2	3	SENTIMENT	1	2	3	SENTIMENT	1	2	3
NEGATIVE	57%	51%	46%	NEGATIVE	63%	57%	49%	NEGATIVE	704.22	622.60	505.76
POSITIVE	31%	35%	41%	POSITIVE	29%	34%	43%	POSITIVE	323.86	367.74	444.66
UNCERTAINTY	12%	14%	13%	UNCERTAINTY	8%	10%	9%	UNCERTAINTY	87.10	105.24	91.65
Case occurrence (row)				Frequency (row)							
SENTIMENT	1	2	3	SENTIMENT	1	2	3				
NEGATIVE	31%	38%	31%	NEGATIVE	31%	39%	30%				
POSITIVE	24%	36%	40%	POSITIVE	23%	36%	41%				
UNCERTAINTY	25%	41%	34%	UNCERTAINTY	25%	42%	34%				

Notes: (1) decrease, (2) no change, (3) increase. The values are calculated using data from the Service Survey.


**Figure 10 Sentiment by Month: Case Occurrence**



**Figure 11 Sentiment by Month: Word Frequency**




**Figure 12 Sentiment by Month: Word Frequency (Rate per 10,000 Words)**



Notes: The indicators shown in Figures 10, 11, and 12 are calculated using the combined data from the Manufacturing and Service Surveys.



**Figure 13 Uncertainty by Month: Word Frequency (Rate per 10,000 Words)**




Notes: The indicators are calculated using the combined data from the Manufacturing and Service Surveys.

Figure 14 Correlation between Sentiment Indicators

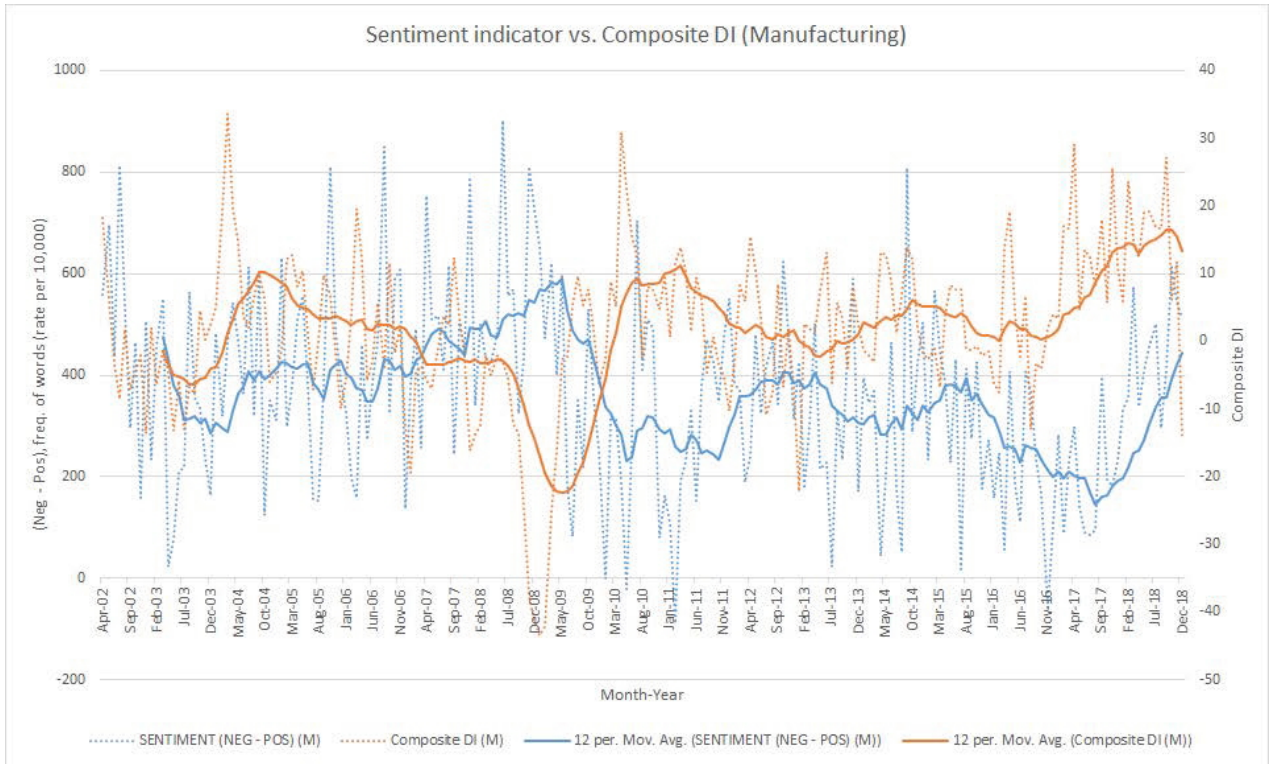
	Case Neg.	Case Pos.	Case Unc.	Freq. Neg.	Freq. Pos.	Freq. Unc.	Freq. rate Neg.	Freq. rate Pos.	Freq. rate Unc.
Case Neg.	1.00								
Case Pos.	-0.71	1.00							
Case Unc.	-0.53	-0.22	1.00						
Freq. Neg.	0.81	-0.65	-0.33	1.00					
Freq. Pos.	-0.65	0.80	-0.07	-0.89	1.00				
Freq. Unc.	-0.43	-0.22	0.87	-0.36	-0.09	1.00			
Freq. rate Neg.	0.70	-0.63	-0.21	0.83	-0.79	-0.21	1.00		
Freq. rate Pos.	-0.57	0.69	-0.04	-0.84	0.92	-0.05	-0.53	1.00	
Freq. rate Unc.	-0.38	-0.26	0.84	-0.31	-0.14	0.97	-0.05	-0.01	1.00

**Figure 15 Sentiment Indicator**



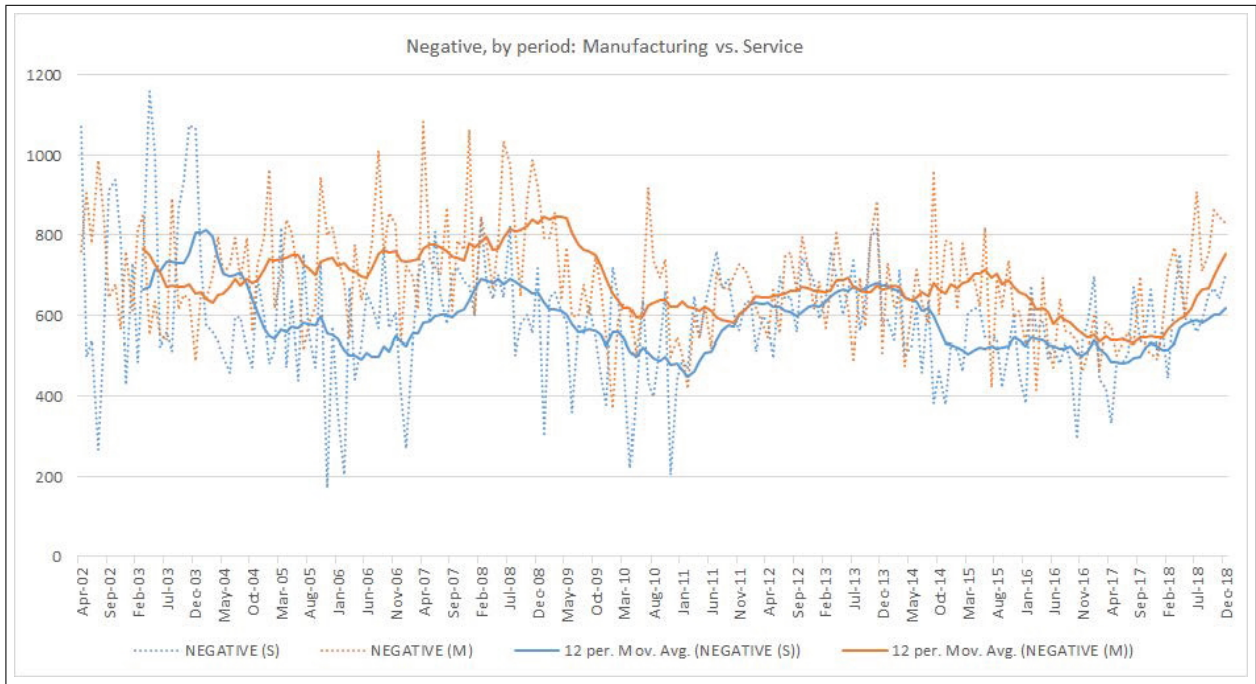
Notes: The indicators are calculated using the combined data from the Manufacturing and Service Surveys.

**Figure 16 Sentiment Indicator vs. Manufacturing Composite DI**




Notes: The sentiment indicator and the composite DI are calculated using data from the Manufacturing Survey.


**Figure 17 Sentiment by Month: Negative, Manufacturing vs. Service (rate per 10,000 words)**




**Figure 18 Sentiment by Month: Positive, Manufacturing vs. Service (rate per 10,000 words)**



**Figure 19 Sentiment by Month: Uncertainty, Manufacturing vs. Service (rate per 10,000 words)**

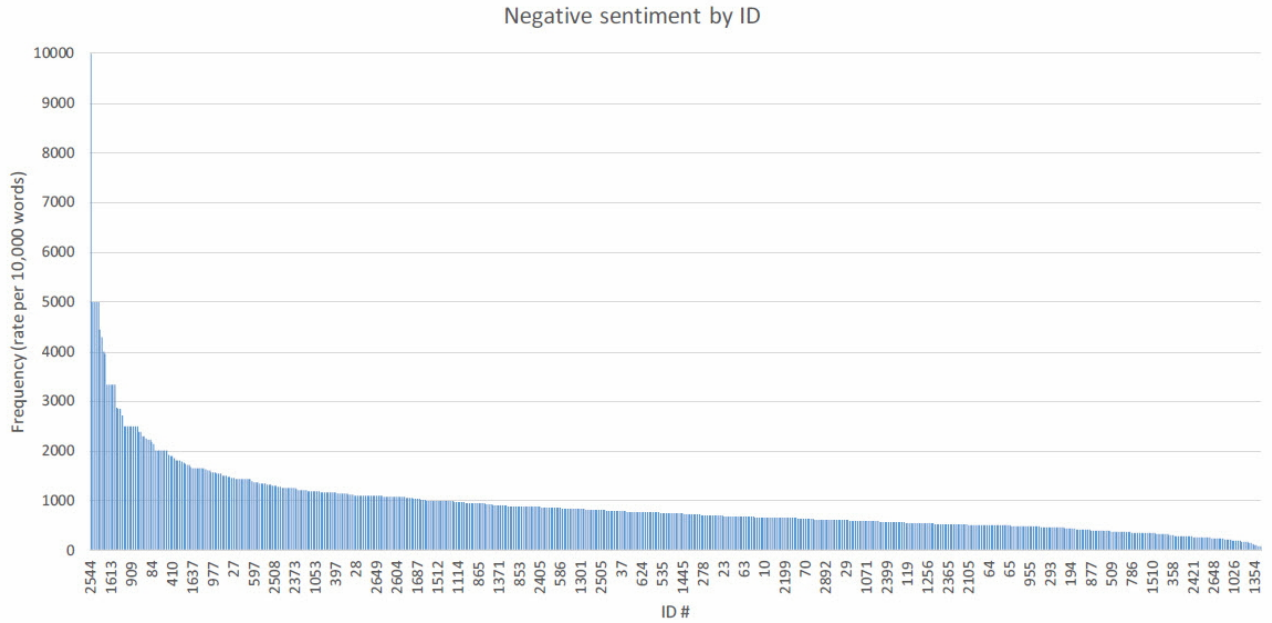


**Figure 20 Sentiment Indicator vs. VIX**




Notes: Left axis: Sentiment indicator. Right axis: VIX.

**Figure 21 Sentiment by ID: Negative**




**Figure 22 Sentiment by ID: Uncertainty**





**Figure 23 Sentiment by ID: Positive Sentiment**



---

---

**APPENDIX: APPENDIX B****1. METHODOLOGY**

The tables in Figures 24 and 25 show a sample of the words included in the dictionary used in the analysis. Note that the words themselves are not directly associated with the respective sentiment category. By using predefined rules, I assume that a sentence expresses a specific sentiment depending on how the words are combined. Figure 26 shows different examples of sentences categorized as “negative,” “positive,” or “uncertain” using the rules described in the text.



Figure 25 Words (continued)

TYPE OF WORD	WORD	FREQUENCY	% TOTAL	NO. CASES	% CASES	TF * IDF
POSITIVE WORDS	CONSTRUCT	397	0.17%	367	0.77%	838.8
POSITIVE WORDS	COOL	22	0.01%	21	0.04%	72.8
POSITIVE WORDS	CORRECT	35	0.02%	34	0.07%	110.1
POSITIVE WORDS	DEAL	51	0.02%	49	0.10%	152.4
POSITIVE WORDS	DEBENT	16	0.01%	16	0.03%	36.6
POSITIVE WORDS	EARN	19	0.01%	18	0.04%	65.0
POSITIVE WORDS	ENCOURAGE*	18	0.01%	18	0.04%	61.6
POSITIVE WORDS	ENERGY	272	0.12%	260	0.55%	618.4
POSITIVE WORDS	ENOUGH	50	0.02%	55	0.12%	170.4
POSITIVE WORDS	EXCEL*	19	0.01%	17	0.04%	65.5
POSITIVE WORDS	EXPAND	68	0.03%	67	0.14%	193.9
POSITIVE WORDS	EXTEND	31	0.01%	31	0.07%	98.8
POSITIVE WORDS	EXTREM*	74	0.03%	73	0.15%	208.3
POSITIVE WORDS	FAR	15	0.01%	15	0.03%	52.5
POSITIVE WORDS	FAST	22	0.01%	22	0.05%	73.4
POSITIVE WORDS	FASTER	13	0.01%	13	0.03%	48.3
POSITIVE WORDS	FAVOR*	30	0.01%	29	0.06%	95.5
POSITIVE WORDS	FREE	21	0.01%	21	0.04%	70.5
POSITIVE WORDS	FRESH	12	0.01%	11	0.02%	43.6
POSITIVE WORDS	GAIN	51	0.02%	49	0.10%	152.4
POSITIVE WORDS	GOOD	601	0.26%	529	1.11%	1174.4
POSITIVE WORDS	GREAT*	137	0.06%	133	0.28%	349.9
POSITIVE WORDS	HA	1209	0.53%	1081	2.27%	1087.2
POSITIVE WORDS	HARD*	12	0.01%	12	0.03%	43.2
POSITIVE WORDS	HELP	194	0.09%	187	0.39%	466.7
POSITIVE WORDS	HIGHER	267	0.12%	244	0.51%	657.3
POSITIVE WORDS	HIGHEST	35	0.02%	31	0.07%	111.5
POSITIVE WORDS	HOME	260	0.11%	248	0.46%	608.2
POSITIVE WORDS	HOPE	169	0.07%	162	0.34%	417.1
POSITIVE WORDS	IMPROVE*	33	0.01%	32	0.07%	104.7
POSITIVE WORDS	INTEREST*	267	0.12%	262	0.54%	639.2
POSITIVE WORDS	LAISSEZ*	59	0.03%	57	0.12%	169.5
POSITIVE WORDS	LEADERSHIP	33	0.01%	30	0.06%	105.6
POSITIVE WORDS	MAINTAIN	47	0.02%	47	0.10%	141.3
POSITIVE WORDS	MIND	13	0.01%	12	0.03%	46.8
POSITIVE WORDS	MODEST	18	0.01%	18	0.04%	61.6
POSITIVE WORDS	MORE*	38	0.02%	37	0.08%	118.2
POSITIVE WORDS	OPPORT*	48	0.02%	46	0.10%	144.7
POSITIVE WORDS	OPTIM*	88	0.04%	87	0.18%	240.9
POSITIVE WORDS	PARTNER	14	0.01%	10	0.02%	51.5
POSITIVE WORDS	PATIENT	56	0.02%	52	0.11%	165.8
POSITIVE WORDS	PERMIT	12	0.01%	12	0.03%	43.2
POSITIVE WORDS	PREMIUM	31	0.01%	31	0.07%	98.8
POSITIVE WORDS	PROFIT*	198	0.09%	188	0.40%	475.0
POSITIVE WORDS	PROGRESS	16	0.01%	16	0.03%	36.6
POSITIVE WORDS	PROMISE*	12	0.01%	11	0.02%	43.6
POSITIVE WORDS	PROJECT	14	0.01%	14	0.03%	49.4
POSITIVE WORDS	RELIEF*	21	0.01%	20	0.04%	70.9
POSITIVE WORDS	REMODEL	18	0.01%	18	0.04%	61.6
POSITIVE WORDS	REPAIR	14	0.01%	13	0.03%	49.9
POSITIVE WORDS	RICH*	29	0.01%	28	0.06%	93.7
POSITIVE WORDS	RISE*	20	0.01%	20	0.04%	67.5
POSITIVE WORDS	ROBUST	27	0.01%	27	0.06%	87.6
POSITIVE WORDS	SAFE*	33	0.01%	33	0.07%	104.2
POSITIVE WORDS	SECURE*	41	0.02%	39	0.08%	126.5
POSITIVE WORDS	SHARE	59	0.03%	49	0.10%	152.4
POSITIVE WORDS	SKILL	178	0.08%	162	0.34%	439.3
POSITIVE WORDS	SOFT	114	0.05%	111	0.23%	300.1
POSITIVE WORDS	SOFTEN	47	0.02%	46	0.10%	141.7
POSITIVE WORDS	SOFTER	13	0.01%	13	0.03%	46.3
POSITIVE WORDS	SOLID	18	0.01%	18	0.04%	61.6
POSITIVE WORDS	SPECIAL	16	0.01%	16	0.03%	55.6
POSITIVE WORDS	STAINLESS	15	0.01%	15	0.03%	52.5
POSITIVE WORDS	STRENGTH*	24	0.01%	24	0.05%	79.1
POSITIVE WORDS	STRONG*	457	0.20%	418	0.88%	939.8
POSITIVE WORDS	SUCCESS*	22	0.01%	22	0.05%	73.4
POSITIVE WORDS	SUPPORT	47	0.02%	47	0.10%	141.3
POSITIVE WORDS	SURE*	73	0.03%	71	0.15%	206.3
POSITIVE WORDS	SUPPORT*	34	0.01%	34	0.07%	107.0
POSITIVE WORDS	TALENT*	27	0.01%	27	0.06%	87.6
POSITIVE WORDS	THANK	26	0.01%	26	0.06%	84.8
POSITIVE WORDS	THESE*	31	0.01%	30	0.06%	92.2
POSITIVE WORDS	TRACTION	24	0.01%	24	0.05%	79.1
POSITIVE WORDS	TRAIN	51	0.02%	45	0.09%	154.2
POSITIVE WORDS	TRAVEL	50	0.02%	42	0.09%	152.7
POSITIVE WORDS	TRUE	22	0.01%	22	0.05%	73.4
POSITIVE WORDS	UNDERSTAND	15	0.01%	14	0.03%	53.0
POSITIVE WORDS	WARM*	127	0.06%	124	0.26%	328.2
POSITIVE WORDS	WELL	154	0.07%	148	0.31%	386.1
POSITIVE WORDS	WIN	19	0.01%	18	0.04%	65.0
POSITIVE WORDS	WISE*	11	0.00%	9	0.02%	41.0
POSITIVE WORDS	WORK	566	0.25%	485	1.02%	1127.4
UNCERTAINTY	ALMOST	74	0.03%	74	0.16%	207.8
UNCERTAINTY	APPEAR	102	0.04%	99	0.21%	273.6
UNCERTAINTY	CAUTIOUS	13	0.01%	13	0.03%	46.3
UNCERTAINTY	DEPEND	83	0.04%	83	0.17%	229.0
UNCERTAINTY	DIFFER	35	0.02%	35	0.07%	109.7
UNCERTAINTY	GUESS	22	0.01%	22	0.05%	73.4
UNCERTAINTY	MIGHT	19	0.01%	19	0.04%	64.6
UNCERTAINTY	PREDICT	28	0.01%	28	0.06%	90.5
UNCERTAINTY	RISK	21	0.01%	21	0.04%	70.5
UNCERTAINTY	SOMEWHAT	67	0.03%	67	0.14%	191.0
UNCERTAINTY	UNCERTAIN	47	0.02%	47	0.10%	141.3
UNCERTAINTY	UNKNOWN	29	0.01%	29	0.06%	93.2

Notes: The symbol “\*” is a wildcard that substitutes for several alternative forms of a word or expression. TF IDF: term frequency x inverse document frequency

**Figure 26 Examples**

Negative	
Not good	Bad
Floods in the region won't help either	Apparel industry continues to suffer
govt furlough does not help consumer confidence problem	We have had several plants and retail stores to close in the past few weeks
One local car dealership closed - another not well	December sales were off over 15% as bad as I have seen
Seasonal--economy fair-not robust	Winter has been extremely tough on our business
From Sept. 10 to mid-October business was not at all good	lack of consumer confidence
Positive	
Not bad	Good
There should be no further decline	demand for our services continues to be strong
High gas and raw goods prices have not dampened the buying spirit	Business is very good.
No adverse signs yet to our business	Better than it's been since 2006
No noticeable loss of business because of gas prices	Weather has improved and we are very busy
inflation and rising raw material prices are not too bad	Our property management business is growing market share and doing well
Uncertainty	
<b>Appear</b>	
Prices appeared fairly stable in April/sales continued strong	
Demand appears to be returning for permanent placement activity	
Energy costs appear stable	
Richmond economy appears steady	
Office leasing appears to be picking up slightly	
<b>Depend</b>	
We are very dependent on Medicaid funding.	
Several capital projects are being planned at this time dependent on the market and economy continuing to remain strong - include lumber storage building	
Increased profits depend largely on increased volume of service	
Depends on gas prices	
Increase in reimbursements depends on cotntracts	
<b>Somewhat</b>	
Business slows down somewhat for hot summer months	
This month somewhat reflects last month	
Recovering oil prices are helping us somewhat	
Holidays should help somewhat	
Business somewhat better	

---

**REFERENCES**

- Barkin, Thomas I. 2019. “Confidence, Expectations and Implications for Monetary Policy.” Speech at the Global Interdependence Center’s Rocky Mountain Economic Summit, Victor, Idaho, July 11.
- Barsky, Robert B., and Eric R. Sims. 2012. “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence.” *American Economic Review* 102 (June): 1343–77.
- Bram, Jason, and Sydney Ludvigson. 1998. “Does Consumer Confidence Forecast Household Expenditure? A Sentiment Index Horse Race.” Federal Reserve Bank of New York *Economic Policy Review* 4 (June): 59–78.
- Calomiris, Charles W., and Harry Mamaysky. 2019. “How News and Its Context Drive Risk and Returns around the World.” *Journal of Financial Economics* 133 (August): 299–336.
- Hansen, Stephen, Michael McMahon, and Andrea Prat. 2018. “Transparency and Deliberation within the FOMC: a Computational Linguistics Approach.” *Quarterly Journal of Economics* 133 (May): 801–70.
- Lazaryan, Nika, and Santiago Pinto. 2017. “Using the Richmond Fed Manufacturing Survey to Gauge National and Regional Economic Conditions.” Federal Reserve Bank of Richmond *Economic Quarterly* 103 (First-Fourth Quarter): 81–136.
- Loughran, Tim, and Bill McDonald. 2011. “When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks.” *Journal of Finance* 66 (February): 35–65.
- Loughran, Tim, and Bill McDonald. 2015. “The Use of Word Lists in Textual Analysis.” *Journal of Behavioral Finance* 16 (March): 1–11.
- Macheras, Ann, Santiago Pinto, Jessie Romero, and Pierre-Daniel G. Sarte. 2015. “Why Does the Fed Study Regional Economics?” Federal Reserve Bank of Richmond *Economic Brief* 15-01 (January.)
- Nyman, Rickard, Sujit Kapadia, David Tuckett, David Gregory, Paul Ormerod, and Robert Smith. 2018. “News and Narratives in Financial Systems: Exploiting Big Data for Systemic Risk Assessment.” Bank of England Staff Working Paper 704 (January).

- Pinto, Santiago, Sonya Waddell, and Pierre-Daniel G. Sarte. 2015. "Monitoring Economic Activity in Real Time Using Diffusion Indices: Evidence from the Fifth District." Federal Reserve Bank of Richmond *Economic Quarterly* 101 (Fourth Quarter): 275–301.
- Pinto, Santiago, Pierre-Daniel G. Sarte, and Robert Sharp. Forthcoming. "The Information Content and Statistical Properties of Diffusion Indices." *International Journal of Central Banking*.
- Price, David A., and Aileen Watson. 2014. "The Richmond Fed Manufacturing and Service Sector Surveys: A User's Guide." Federal Reserve Bank of Richmond *Economic Brief* 14-03 (March).
- Shapiro, Adam Hale, Moritz Sudhof, and Daniel Wilson. 2018. "Measuring News Sentiment." Federal Reserve Bank of San Francisco Working Paper 2017-01 (June).
- Souleles, Nicholas S. 2004. "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys." *Journal of Money, Credit and Banking* 36 (February): 39–72.
- Thorsrud, Leif Anders. 2018. "Words Are the New Numbers: A Newsy Coincident Index of the Business Cycle." *Journal of Business & Economic Statistics*, <https://doi.org/10.1080/07350015.2018.1506344>.
- Waddell, Sonya. 2015. "Predicting Economic Activity through Richmond Fed Surveys." Federal Reserve Bank of Richmond *Econ Focus* (Fourth Quarter): 36–39.
- Young, Lori, and Stuart Soroka. 2012. "Affective News: The Automated Coding of Sentiment in Political Texts." *Journal of Political Communication* 29 (April): 205–31.