CEWP 22-01

Using Natural Language Processing to Measure COVID19-Induced Economic Policy Uncertainty for Canada and the US*

Shafiullah Qureshi — Carleton University
Ba Chu — Carleton University
Fanny S. Demers — Carleton University
Michel Demers — Carleton University

January 18, 2022

CARLETON ECONOMICS WORKING PAPERS

Carleton University

Department of Economics

1125 Colonel By Drive
Ottawa, Ontario, Canada
K1S 5B6

Using Natural Language Processing to Measure COVID19-Induced Economic Policy Uncertainty for Canada and the US*

Shafiullah Qureshi, Ba Chu, Fanny S. Demers, Michel Demers†

Carleton University & Ottawa-Carleton Graduate School of Economics

January 18, 2022

Abstract

We develop an economic policy uncertainty (EPU) index for Canada and the US using natural language processing (NLP) methods. Our EPU-NLP index is based on an application of several algorithms, including a rapid automatic keyword extraction algorithm (RAKE), a combination of the RoBERTa and the Sentence-BERT algorithms, a PyLucene search engine, and the GrapeNLP local grammar engine. For comparison purposes, we also develop an index based on a strictly Boolean method. We find that the EPU-NLP index captures COVID-19 related uncertainty better than the Boolean index. Using a structural VAR approach, we found that an economic policy uncertainty shock with EPU-NLP results in larger declines in real GDP, employment, industrial production and the TSX index than with EPU-Boolean for Canada. Similar results were also found for the US: an EPU-NLP shock led to larger declines in industrial production, employment, real personal consumption expenditure, and S&P500 than EPU-Boolean. The SVAR model showed an abrupt contraction in economic variables both for Canada and the US in line with the COVID-19 impact. Moreover, an uncertainty shock (with the EPU-NLP) caused a much larger contraction in economic variables for the period including the COVID-19 pandemic, than for the period before COVID-19.

*This research is partially funded by a Faculty of Public Affairs Research Engagement Grant.

†Department of Economics, Carleton University, Ottawa, Canada. S. Qureshi also at Department of Economics,NUML, Islamabad, Pakistan. Emails: shafiullah.qureshi@carleton.ca suqureshi@numl.edu.pk,ba.chu@carleton.ca, fanny.demers@carleton.ca, michel.demers@carleton.ca

1 Introduction

The sudden incursion of the COVID-19 pandemic and the world-wide recession that followed have generated great interest in measuring the resulting uncertainty and its impact on macroeconomic variables. An increase in uncertainty has been shown to have a very important impact on economic decisions, particularly, on investment decisions, if firms face irreversibility, ((Bernanke, 1983), (Demers, 1991),(Dixit and Pindyck, 1994)), fixed costs ((Caballero and Engel, 1999)) or financial constraints, and also on consumption decisions when consumers are risk averse, prudent or face binding budget constraints. Obtaining a measure of the degree of uncertainty is important for assessing its macroeconomic impact and for guiding policymakers in making appropriate monetary and fiscal policy decisions.

Furthermore, policy itself may lead to uncertainty. Thus, for example, (Altug et al., 2009)investigate the impact of tax-policy uncertainty on the dynamic investment decisions of the firm. Several authors have given priority to developing an index to measure uncertainty. One prominent example is the forward-looking Baker-Bloom-Davis newspaper-based economic policy uncertainty index (Baker et al., 2016). Other notable examples are the model-based uncertainty measures of (Jurado et al., 2015) for the US and (Moran et al., 2020) for Canada.With the COVID-19 shock as a backdrop, (Altig et al., 2020) note that while model-based measures have the benefit of being well grounded in a model in which the role and the nature of uncertainty is well-defined, such measures are essentially backward looking, and are based on the premise that the underlying model has not changed and that the statistical relationship among variables is still the same even after large and unprecedented shocks. Furthermore, the macroeconomic variables (leading indicators) in the underlying model are only available with a lag, and hence, not available in real time. In the wake of the COVID-19 shock, (Altig et al., 2020) thus point to the importance of having alternative measures of uncertainty that are *forward looking* and available in *real time*.

As mentioned above, an important and very widely used measure of uncertainty is the Economic Policy Uncertainty (EPU) Index developed by (Baker et al., 2016, henceforth,BBD). Being forward-looking in nature, the BBD-EPU newspaper-based index has been found by various authors to successfully capture uncertainty, especially policy uncertainty. Currently, an index is available for 26 countries (including the US and Canada). The use of this index is so widespread that data providers such as Bloomberg, FRED, Haver and Reuters also make the EPU available for users on their website. Their index has also been used in numerous economics articles since its development. We describe in detail the development of the BBD-EPU index in section 2 below. Let us simply note here, however, that the BBD-EPU (at least the one for the US) was very human-input intensive and expensive to develop.

In this paper, we suggest an alternative newspaper-based, and (almost entirely) computer-based, approach to developing an EPU index directly related to COVID-related uncertainty for Canada and the US, by appealing to Natural Language Processing (NLP) techniques ((Gentzkow et al., 2019)). These techniques are widely used by software engineers, but have not yet received much attention in economics. Our index circumvents the necessity to rely very heavily on human resources, is less expensive and faster to obtain. These attributes make it useful for developing EPUs for country-specific policy categories and subcategories, for developing monthly or daily EPUs, and also for EPUs for countries not yet having their own BBD-EPU. We use a "text mining" approach (an artificial intelligence (AI) technology,) that uses NLP to transform unstructured data (such as ordinary texts) into structured data (i.e., data or texts that are organized into categories, such as, username, user ID, address, etc.) that in turn permit computers to understand, interpret, and classify human language.

Our method differs markedly from a Boolean method. In contrast to the latter, our method is capable of capturing *contextual* and *implied meanings* of EPU-related terms thanks to our use of the RoBERTa (Liu et al., 2019) algorithm[1] which we combine with its specialization for semantic searches, SBERT (Reimers and Gurevych, 2019). To ensure greater accuracy and robustness with respect to capturing the contextual meaning of words, we also use an additional independent NLP algorithm, namely GrapeNLP, developed by (Sastre, 2011), which is based on Unitex-GramLab (**?**) (Paumier et al., 2009).

In order to highlight the important difference between a Boolean method and the EPU-NLP, we also develop an alternative, strictly Boolean, index (EPU-Boolean) which we compare with our NLP-based one. We show that the EPU-NLP is better able to track COVID-19-related uncertainty than the EPU-Boolean.

[1]RoBERTa is the "robustly optimized" version of (Devlin et al., 2018)'s seminal neural-network-based BERT (*Bidirectional Encoder Representations from Transformers*). We describe these algorithms below.

We also compare the EPU-NLP with other leading uncertainty indices such as the BBD-EPU, BBD's equity market volatility index (EMV), and the Chicago Board Options Exchange's (CBOE) volatility (VIX) index, and find that it is closely correlated with them. In addition, we also develop a weekly EPU-NLP index for Canada (whereas only a monthly EPU index is currently available), and for the US, and we show that this index traces well several notable uncertainty-generating events.

We then conduct a structural vector autoregression (SVAR) analysis to observe the impact of a one standard-deviation (SD) EPU-NLP shock on some macroeconomic variables for Canada and the US. We also compare the impact of a one-SD EPU-NLP shock on pre-COVID-19 data (Jan. 2015-Dec. 2019) with its impact on the data range including COVID-19 data (Jan. 2015-October 2020). The SVAR results show that a one-SD shock in the EPU-NLP index provokes a larger contraction in real GDP and other macroeconomic variables for Canada and the US than a one-SD shock in the EPU-Boolean index. Moreover, a EPU-NLP shock results in a stronger decline in these variables for the span of time that includes the COVID-19 pandemic than the one that excludes it.

The remainder of the paper proceeds as follows. In Section 2 we explain in detail the development of the BBD-EPU. In Section 3, we present the stages of the construction of the EPU-NLP and describe the algorithms that were used in its development. Section 4 presents the SVAR results, and section 5 concludes. Lastly, in the Appendix, we provide additional figures and tables.

2 The development of the Baker-Bloom-Davis EPU (BBD-EPU)

As (Baker et al., 2016) explain extensively in their paper, the construction of their index was done in 2 stages over 2 years, and involved a great deal of human resources[2]. The authors first developed a 65-page guideline over a 6-month period. Then, under close supervision by the authors and on the basis of the guideline, different teams of students were trained as readers (auditors), and were given the task of sifting through 12,000 newspaper articles to identify those that contained three terms, namely one from each of the following three sets: E (economic *or* economy); U (uncertain *or* uncertainty); and *at least one* policy-related term from the third set P (Congress, deficit, Federal Reserve, legislation, regulation or White House.) ((Baker et al., 2016), p.1594). Their index was based on the frequency of articles in 10 leading U.S. newspapers[3] containing these keywords, where the articles were rated by the human auditors as either $EPU^H = 1$ or $EPU^H = 0$ depending on whether they contained the three categories of terms or not. When coding an article as $EPU^H = 1$ the human-auditor would also record the exact policy terms that appeared in the passage related to economic policy uncertainty, thus permitting the authors to develop a set of 15 terms with a high frequency of occurrence in the articles. Using permutations of 4 policy terms out of these 15 terms, they generated a computer assignment of either $EPU^c = 0$ or $EPU^c = 1$ for about 32,000 permutations (i.e., $\frac{15!}{(15-4)!} \approx 32,000$). They then proceeded to compare the human-based records with the computer-based ones in order to eliminate the false positives and false negatives generated by the computer-based records, and to choose the set of terms that minimized the false positives plus the false negatives (of the computer-based records). They found a correlation of 0.89 between their human and computer generated indices for the 1989-2012 period. They also developed a general overall economic uncertainty index (EU), where they drop the requirement for the third policy related term, as well as (in the case of the US) *specialized* EPU indices for 11 different policy categories (and sub-categories) such as fiscal policy, tax policy, monetary policy, healthcare policy, national security policy, etc. They developed both a monthly and a daily index for the US, but only a monthly index in the case of other countries, including Canada.

As the authors themselves indicate, the development of this index was very intensive in terms of human-input, and required substantial resources (i.e, it was expensive). In addition, one of the reasons for the false positives and false negatives generated during their computer-based records is due to the Boolean nature of their procedure, and to words being evaluated out of context by their computer-based method. In the case of

[2] "We spent six months developing an audit process designed to evaluate and refine our U.S. EPU indexes and another 18 months running a large-scale human audit study. During the latter phase, student teams working under our close supervision read and coded articles drawn from eight newspapers..."(BBD, 2016, p. 1608)

[3] These are: USA Today, Miami Herald, Chicago Tribune, Washington Post, Los Angeles Times, Boston Globe, San Francisco Chronicle, Dallas Morning News, New York Times, and Wall Street Journal.

the BBD-EPU developed for the US, they are able to sift out the false positives and negatives by comparing the computer generated records with their very meticulously developed human-based records.[4]

It should be added that while the BBD-EPU developed for the US (and for other English-speaking countries) benefited greatly from their human-based audit as a verification process for their computer-generated records, this does not seem to have been quite the case for the BBD-EPU's for countries where the native language is other than English. For these, there was a more cursory verification process: "To help develop suitable E, P, and U term sets, we consulted persons with native-level fluency and economics expertise in the relevant language and country." In fact, only two newspapers per country were used and the search terms were often simply translations of the English search words (given in Table 2 below), and the selection of articles was entirely Boolean in nature.(See the On-line Appendix to (Baker et al., 2016))

3 Constructing the EPU-NLP Index: Data, Methodology and Algorithms

We describe here the procedure that we used to develop our EPU-NLP index. We note that we are using two NLP techniques to refine our search for relevant articles, the first being the application of RoBERTa/SBERT (mentioned in step 3 below), followed by the GrapeNLP approach in step 5. Our motivation for choosing the BERT family of algorithms and GrapeNLP is that *each* of these techniques was (separately) very successful in a competition held at Kaggle whose goal was to extract summary tables from a large data set of COVID-19-related articles (500,000 research articles, about COVID-19, SARS-CoV-2, and related coronaviruses). In particular, (Sastre et al., 2020), one of the winning entries, used GrapeNLP to find the impact of temperature and humidity on the spread of the virus. Here, we use both approaches consecutively as an attempt at greater accuracy and robustness.

We base ourselves on articles gathered from eight newspapers from Canada and seven newspapers for the US from January 2015 to October 2020. For Canada, these were: The Calgary Herald, The Financial Post, The Montreal Gazette, The National Post, The Ottawa Citizen, The Toronto Star and The Vancouver Sun. For the US, we used USA Today, The Los Angeles Times, The Wall Street Journal, The Dallas Morning News, The Miami Herald, and the New York Times. (See also Table 2 and Table 3). We first enumerate the six steps that we followed and then explain the procedure in greater detail below.

1. Use the *RAKE* algorithm to search for frequently used economy, uncertainty and policy related words in the newspapers.

2. Select articles that contain the words obtained from step 1 using a Python filter. (1,182,945 articles for Canada and 720,266 for the US.)

3. Use a combination of RoBERTa and SBERT to filter out those articles having a cosine-similarity score of 0.75 or more (We remain with 622,948 articles for Canada and 379,166 for the US at this stage.)

4. Use a Python-based Apache Lucene search engine (PyLucene) to shortlist further articles selected from step 3 based on relevant search words.

5. Develop a "local grammar" with Unitex/GramLab on the basis of the keywords obtained from the previous steps. We use this local grammar with the GrapeNLP python package developed by (Sastre, 2011) (We finally remain with 18,526 articles for Canada and 18, 032 articles for the US.)

6. We calculate the EPU-NLP index following the method indicated in BBD. ((Baker et al., 2016))

[4]Even so, the following quote from their paper illustrates some of the difficulties they faced in the context of their computer-generated index. "We also experimented with compound text filters, for example, adding {government AND tax} to the baseline term set. Somewhat to our surprise, we were unable to develop simple compound text filters that achieved a materially lower gross error rate than our baseline term set." (p.1609) Thus, as the authors note, adding "tax" to the set of policy terms led to a false positive by choosing an article that "...includes remarks about taxable and tax-exempt securities, but it contains no discussion of policy matters."(Footnote 19, p. 1609) As we will see below, our NLP-based index would be able to avoid such an article, since our method chooses articles based on an exact match of particular word-groups and would overlook articles only containing "taxable" and tax-exempt" if these two word groups are not among the ones specified.

To highlight the benefits of using the NLP approach, we also constructed another index on the basis of the same data set, using a strictly Boolean approach, assigning a $EPU^{Boolean} = 1$ or $= 0$ depending on the presence or absence of uncertainty related keywords,(but without the ability to discern the context in which these keywords appear.) We now explain in greater detail the RAKE, RoBERTa/SBERT and GrapeNLP algorithms that were used in these steps as well as the calculation method of the EPU-NLP in the last step.

As we indicated above, to ensure greater robustness in our results, we are using two NLP techniques consecutively to refine our search for relevant articles, the first being the application of RoBERTa/SBERT as mentioned in step 3, followed by the GrapeNLP approach in step 5.

3.1 The RAKE (Rapid Automatic Keyword Extraction) algorithm

RAKE is a language-independent, unsupervised ML algorithm developed by Rose et al. (2010). A "keyword" (also called a "token") is defined as a sequence of one or more words. This algorithm splits text into a list of keywords, by using "stop words" (like "the", "a", "at", "for", "above", "on", "is", "all"), as well as punctuation (comma, semicolon (;), quotes (", ') etc.) as a means of separating one string of contiguous words from another. These strings are candidates for keywords. The algorithm then creates a table of "co-occurrences" of each of the words in these strings (i.e., words that occur together within the string). For example, the words economic and crisis may occur together in some strings. The algorithm assigns a score to each word (w) on the basis of its frequency of occurrence within the entire text (freq(w)) and also on the basis of its "degree" (deg(w)), that is, depending on whether the word appears in conjunction with another word. The score assigned to a single word is the ratio of deg(w)/(freq(w)) and the score assigned to a keyword string is the sum of the score assigned to the words composing it. Thus, for example, if the word economic appears twice in the text, once with the word "crisis", and the other time with the word "turmoil," it gets a degree (deg(economic)) score of 2+2=4 because it appears twice and also because it appears in conjunction with two other words that are part of the keywords. If the words crisis and turmoil appear only once each, and only in conjunction with the word economic, then the "keyword" economic crisis receives a score of:

$$deg(economic)/freq(economic) + deg(crisis)/freq(crisis) = (4/2) + (2/1) = 4$$

We built an initial short list of simple keywords, such as uncertain, economic, recession, Covid-19 and coronavirus, as a means of initializing the RAKE algorithm. We fed these words into RAKE in order to observe their frequency of occurrence in the articles and also to observe other uncertainty-related keyword sequences. We thus produced a list of bi-grams (two-word groups), and tri-grams which had a very high frequency of occurrence and a high score. For example, we found that the following terms appeared very frequently in our data set: *coronavirus crisis* (3218 times), *virus crisis* (3290 times), *economic uncertainty* (447 times), *economic crisis* (2029 times), *remains unclear* (434 times), *job losses* (2009 times), *consumer confidence* (684 times), *virus lockdown* (2309 times), *global recession* (1081times), *make ends meet* (583 times). (These frequently-observed terms also helped us in developing a "local grammar" (see (**?**)) as we will see below.

3.2 The BERT, RoBERTa and SBERT algorithms

In step 3 of our procedure, we use a combination of RoBERTa and SBERT in order to develop *sentence embeddings* and further refine our selection of articles. Word or sentence embedding is a technique in ML that is used to map words or phrases into vectors of real numbers. We develop a list of *queries* (full sentences) on the basis of the RAKE results (see Table 3 and Table 4), which are then processed by the SBERT and RoBERTa algorithms to extract sentence embeddings. These algorithms are extensions of the neural-network-based BERT (*Bidirectional Encoder Representations from Transformers*) algorithm developed by (Devlin et al., 2018) at Google. One of the particularities of the BERT algorithm is that it is *bi-directional.* That is, it can better detect the context within which a word occurs by taking into account words that appear both *before* and *after* the keyword (i.e., both to its left and to its right), and will perform a different embedding depending on the context.Thus, for example, the following sentences given below will be embedded (encoded) differently in view of the different contextual meaning of the word *taxing*:

Bicycling up this steep hillside is very taxing on one's legs.
The new policy involves taxing the rich at a higher rate.

Another similar example would be the different contextual meaning of the word *bank* in the following sentences*:*

> *We went to the river <u>bank</u>.*
> *The <u>bank</u> rate is expected to go up.*

In this respect, it can resolve ambiguities related to words having different meanings in different contexts and therefore can better avoid false-positives or false-negatives, and better select the relevant articles. The BERT algorithm is pre-trained by using a technique called *Masked Learning Modeling* which essentially "hides" (i.e., “masks”) 15% of the keywords in each query by replacing them with another token or mask, and requiring the algorithm to predict the true keywords.[5] The RoBERTa algorithm developed by (Liu et al., 2019) is a "robustly optimized" version of the BERT algorithm that uses a much larger training set (160G instead of 16G). The RoBERTa algorithm presents several other advantages over the BERT algorithm. Thus, for example, it uses dynamic "masking" as opposed to the static masking used in BERT.[6] The dynamic masking pattern of RoBERTa implies that the pattern of masking changes with each sequence that is fed into the algorithm, whereas in BERT, the pattern is set at the initial sequence and remains fixed throughout. The RoBERTa algorithm requires fine-tuning depending on the application at hand. Hence, in the current application, the RoBERTa model was fine tuned by using unlabeled newspaper articles and using the "hugging face" open source NLP platform.

The SBERT algorithm (Reimers and Gurevych, 2019) is a further refinement of BERT that is better suited for semantic searches.[7]The SBERT algorithm uses two identical ("siamese") sub-networks (both of them RoBERTa algorithms) so as to compare two sentence (or document) embedding representations (i.e., paragraphs/corpus of sentences and queries) using the cosine-similarity measure to assess the degree of similarity between newspapers articles and queries. The cosine-similarity measure is given by:

$$\sigma\left(A,B\right)=\frac{A\cdot B}{\|A\|_2\|B\|_2}=\frac{\sum\limits_{i=1}^{n}A_iB_i}{\left(\sqrt{\sum\limits_{i=1}^{n}A_i^2}\right)\left(\sqrt{\sum\limits_{i=1}^{n}B_i^2}\right)}$$

where A and B are n dimensional vectors, and $\|.\|_2$ is the L^2 norm.The cosine-similarity measure $\sigma(A,B)$ takes the value 1 when the two vectors are exactly the same and the value -1 when they are completely dissimilar. Comparing this similarity measure to cross-entropy or to mean-squared error MSE (which uses Euclidean distance as its measure of closeness), this measure has the advantage of being dependent only on the direction of the vectors and not on their magnitudes, and hence, is independent of the scaling of the two vectors. [8] In the third step of our procedure, we pair the articles obtained in step 2 with sentences from the queries that we developed on the basis of the RAKE results (Table 3) and feed them into SBERT. After a pooling process[9] needed to transform the contextual embeddings obtained from each RoBERTa sub-network into vectors of fixed length, SBERT calculates the cosine-similarity measure between the articles and the queries. We choose articles that have a cosine-similarity measure of 0.75 or greater.

[5]To be more precise, 10% of the masked terms are replaced with randomly selected keywords, 10% are replaced with the true word and 80% are replaced with the token [MASK].

[6]It eliminates the "next sentence prediction" feature of BERT and also uses a byte-level byte-pair encoding with a vocabulary of 50,000 sub-word units (as opposed to a character-level byte-pair-encoding with a vocabulary of 30,000 units in BERT).

[7]It is also much faster than BERT. The authors note that finding the most similar pairs of sentences in a set of 10,000 sentences took 65 hours with BERT, but only 5 seconds (and an additional 0.01 second to compute the cosine similarity) with SBERT.

[8]As (Barz and Denzler, 2020) emphasize, this is an important advantage especially when dealing with relatively small data sets. In addition, the loss function based on the cosine-similarity measure, $L_\sigma = 1 - \sigma\left(A,B\right)$, is bounded between (0,2) as opposed to the loss functions dependent on cross-entropy and MSE which can take arbitrarily large values. As (Barz and Denzler, 2020) note, the L^2 normalization serves as a regularizer, dispensing with the need for an additional regularization hyper-parameter that would have to be tuned for each data set.

[9]We use the RoBERTa model to map tokens in a sentence to the contextual word embeddings from RoBERTa. The next layer in our model consists of averaging (“mean-pooling”) all contextualized word embeddings obtained from RoBERTa. In other words, each sentence is passed first through the word_embedding_model (in RoBERTa) and then through the pooling_model to give fixed-sized vectors. Vectors of fixed length are required by SBERT.

3.3 GrapeNLP grammar

In step 5, we use GrapeNLP grammar to further refine our choice of articles. We follow (Sastre et al., 2020) who used this approach (developed in (Sastre, 2011)) to extract research papers from the *COVID-19 Open Research Dataset,* comprising 500,000 research articles about COVID-19 and other similar viruses. Their aim was to identify articles that were directly related to the impact of temperature or humidity on the spread of the virus.[10] We build a local grammar using the Unitex grammar editor[11] developed by (Paumier et al., 2009). We then use GrapeNLP to convert the grammar to a form that may be processed by Python. Let us first illustrate this approach by using an example given in Sastre. We reproduce it in Figure 1 below.[12]

Figure 1: Example: A GrapeNLP grammar diagram

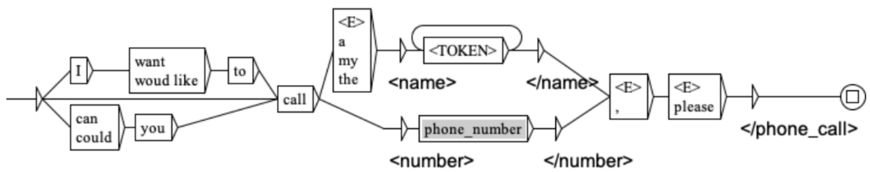

The example shown in Figure 1, is a grammar that is built to recognize sentences that may be used by someone requesting to make a phone call. The sentence may take different forms: I (want/would like) (to) (call) (<E>a/my/the) (TOKEN) or (phone number). The box containing an <E> is optional, in the sense that none of the terms in that box need be present (e.g. I want to call 911) Boxes without an <E> are compulsory and at least one of the words in that box must be present. The box labeled TOKEN may contain any name (for example, Mary, mother, emergency, etc.) while the box (phone number) is a subgrammar (i.e., another grammar that is evoked by this one) which recognizes phone numbers and which may contain symbols such as + in the case of country codes, or parentheses, etc. Alternatively, the sentence might take the form "Could you call my sister, please". We again note that the comma and the word "please" are optional. For our purpose of finding phrases expressing policy or COVID-19 related uncertainty in the newspaper articles, we adapt this methodology to find phrases such as "economic uncertainty caused by the coronavirus lockdown" or more complex ones such as "During the prolonged period of the coronavirus crisis targeted transfers are urgently needed to stay above the poverty line."

This approach involves a "human-assisted" training of grammar. We review the results of an initial trial and then change the “grammar” accordingly. Thus we develop four “grammars” : *causal_forward*, *policy_or_covid*, *list_of_factors*, and *effect* which are linked to one another (hence the use of the term *sub-grammar* below when one of these is invoked by another). We use the frequently-encountered terms that were selected by RAKE in developing these grammars. As shown in Figure 2, to find a causal or forward relationship in a text, we introduce the sub-grammar called *causal_forward* (shown in Figure 5 in the appendix). The part of the text that matches the causal or forward relationship is delimited by XML[13] tags <excerpt>. The <excerpt> tag indicates the relevant part of the text to be extracted from the newspaper article. Since the grammar engine executes an *exact matching* on the entire article’s full text, a “null_insert” is inserted to capture all of the text in the article *before* and *after* the causal_forward segment, to match the article’s beginning and end.

[10] Their approach was one of the winners among 500 participants at the related Kaggle competition. See (Sastre et al., 2020).

[11] Based on (Gross, 1997), Unitex/Grammlab is an open-source, cross-platform, multilingual, lexicon- and grammar-based corpus processing tool. It can be downloaded from https://unitexgramlab.org/

[12] This illustrative example of how to use GrapeNLP can be found in this link https://www.kaggle.com/javiersastre/grapenlp-grammar-engine-in-a-kaggle-notebook

[13] XML which stands for the extensible markup language, is a text-based markup language set of codes, or tags, that describes the text, which is both human and machine-readable. XML tags are case-sensitive. Words starting with upper or lower-case letters are treated differently.

Figure 2: Grammar axiom: Policy or COVID-19

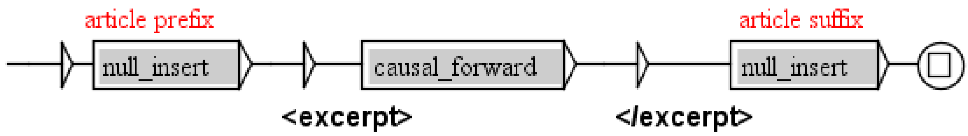

A lexical token denoted by <TOKEN> matches any word, symbol or digit. Hence the matches occur at the token level (not at the character level, so that one could not break up the word go-ing into separate tokens go and -ing. It is therefore necessary to enter "go" and "going" and other forms of the verb as separate lexical tokens). A penalizing_insert is similar to a null_insert. It allows the causal_forward expression to contain unknown token inserts, that is, words that are not related to the search terms (denoted by *blah blah*), but ensures a minimal occurrence of these by penalizing them. (See Figure 3)

Figure 3: Grammars null_insert (left) and penalizing_insert (right)

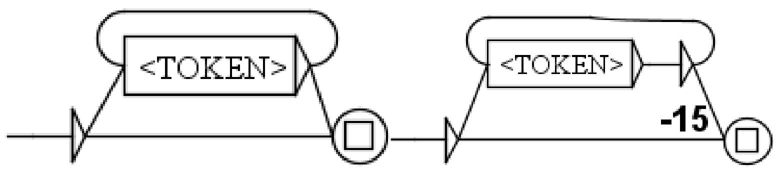

The grammar *list_of_factors* in Figure 4 contains a list of one or more relevant factors. The matched list must start and end with a known factor. The matched list of factors is delimited by the tag <factors> indicating the part to be extracted. It is dependent on the sub-grammar *policy_or_covid* which is given in Figure 6 in the appendix. The latter covers nine channels that take into account the different impacts that either a government policy or COVID-19 may have. The <E> on top of a box makes that box optional. That is, the words in the boxes indicated with an <E> may or may not appear. However, when a box does not have an <E> then the box is compulsory in the sense that at least one of the words in that box must be present. For example, by looking at the bottom of Figure 6 , one possible path is "economic uncertainty caused by the coronavirus lockdown". The terms "economic uncertainty" and "coronavirus" are compulsory. These words must be present in order to select the articles, while "caused by" and "lockdown" are optional. These need not be present in the text. The grammar *effects* given in Figure 7 is the sub grammar for Figure 5. It considers different expressions capturing the effects of the COVID crisis such as economic policy uncertainty, economic uncertainty, economic devastation, etc. We use a similar approach for the US. Only slight changes are made in the *policy_or_covid* and *effects* grammars. These are shown in Figure 8 and Figure 9, respectively.

Figure 4: List of factors

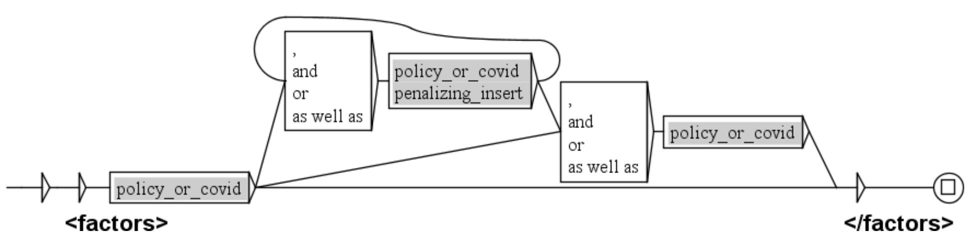

3.4 Calculating the EPU-NLP

In the final step of our procedure, we follow (Baker et al., 2016, p.1599)'s method to calculate the economic policy uncertainty index (EPU).

1. Let $c_{i,t}$ $i = 1, \ldots, N$, $t = 1, \ldots, T$ denote the raw count of articles found to be uncertainty-related in newspaper i in month t and let $Total_{it}$ be the *total* number of articles in newspaper i in month t. The scaled count is then given by C^*_{it}

$$C^*_{it} = \frac{c_{it}}{Total_{it}}, \;\; i = 1, \ldots, N; \;\; t = 1, \ldots, T. \tag{1}$$

2. Compute the variance of scaled counts as ${\sigma_i}^2 = (1/T)\Sigma_{t=1}^{T}(C^*_{it} - \overline{C}_i)^2$ where $\overline{C}_i = (1/T)\Sigma_{i=1}^{T} C^*_{it}$ is the average over the entire period of the scaled counts of articles in newspaper i.

3. Divide the scaled counts by their standard deviation : $Y_{it} = C^*_{it}/\sigma_i, i = 1, ..., N; t = 1, ...T$. (Thus, for example, $Y_{it} = 2$ would indicate that this scaled count obtained in month t is two standard deviations above the mean for newspaper i for the entire time period, and would point to a period of higher uncertainty.)

4. Compute $Z_t = \frac{1}{N}\sum\limits_{i=1}^{N} Y_{it}$, where Z_t is the average of the scaled standardized counts over all newspapers for month t.

5. Calculate $M = \frac{1}{T}\sum\limits_{t=1}^{T} Z_t$where M is the average of the standardized scaled s over all newspapers and for all months in the data set

6. Calculate the normalized EPU time-series index as $EPU_t^{NLP} = \left(\frac{100}{M}\right) Z_t$. With this normalization, the EPU-NLP has a mean of 100.

4 Testing the model

We adopt a structural vector autoregression approach (SVAR) to test our EPU-NLP index and our EPU-Boolean index for both Canada and the US using data from January 2015 to October 2020. As in (Baker et al., 2020) and (Altig et al., 2020), we detrend all the variables using the Hamilton filter (Hamilton, 2018) and then take the first difference of the log of these variables.

Even though VAR models may not be used to establish causality, as (Altig et al., 2020) note, they can indicate whether uncertainty shocks are precursors to a slowdown in economic activity, such as a fall in GDP and employment. Using a vector autoregression analysis, they find significant contractions in economic variables during COVID-19 for both the US and the UK. (Baker et al., 2020) use stock market volatility, newspaper-based economic uncertainty, and subjective uncertainty in business expectation surveys to measure the COVID-19 induced uncertainty for the US. As mentioned earlier, (Moran et al., 2020) constructed an uncertainty measure for Canada by applying the method of (Jurado et al., 2015) and assessed the impact of COVID-induced uncertainty on economic variables using a SVAR analysis for Canada. To indicate the advantages of a SVAR, let us start with a standard VAR model which in our case may be described as follows:

$$\boldsymbol{y}_t = \boldsymbol{A}_1 \boldsymbol{y}_{t-1} + \boldsymbol{A}_2 \boldsymbol{y}_{t-2} + \cdots + \boldsymbol{A}_p \boldsymbol{y}_{t-p} + \boldsymbol{\varepsilon}_t \tag{2}$$

where y_t is a 5×1vector of 4 macroeconomic variables and one uncertainty index, p denotes the number of lags. The $\boldsymbol{A}_i$, $i = 1, ...p$ are 5×5 matrices of parameters, while $\boldsymbol{\varepsilon}_t$ is a 5×1 vector of innovations with $\boldsymbol{\varepsilon}_t \sim N(0, \boldsymbol{\Sigma})$ and $E(\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}'_s) = \boldsymbol{0}$ for all $s \neq t$. Equation (1) is a standard VAR which describes a reduced form model and which does not allow contemporaneous effects of the endogenous variables on each other. It also has the underlying counterfactual assumption that the innovations of the different equations are mutually uncorrelated. Since the innovations are in fact correlated in our model (as an analysis of the covariance matrix $\boldsymbol{\Sigma}$ reveals), a shock to one variable will have an impact on the innovation of another variable, precluding a clear interpretation of the impulse-responses. We therefore adopt a SVAR approach and use a Cholesky decomposition in order to identify the shocks. Our equation may be written as:

$$\boldsymbol{A}(\boldsymbol{I}_5 - \boldsymbol{A}_1 L - \boldsymbol{A}_2 L^2 - ... - \boldsymbol{A}_p L^p)\boldsymbol{y}_t = \boldsymbol{B}\boldsymbol{e}_t \tag{3}$$

where $\boldsymbol{A}$ is a lower triangular matrix with ones in the diagonal, $\boldsymbol{B}$ is a diagonal matrix and $\boldsymbol{e}_t$ is a 5×1 vector of orthogonalized innovations with $\boldsymbol{e}_t \sim N(\mathbf{0}, \boldsymbol{I}_5)$ and $E(\boldsymbol{e}_t \boldsymbol{e}_s') = 0$ for all $s \neq t$ such that $\boldsymbol{B}\boldsymbol{e}_t \equiv \boldsymbol{A}\boldsymbol{\varepsilon}_t$. The structure of the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ is set in accordance with the order of the variables in the VAR. The order of the variables matters for the results. The variable that is listed first is assumed to have a contemporaneous impact on the rest of the variables, while none of the other variables has a contemporaneous impact on the first. Similarly, each of the variables will have a contemporaneous impact on the rest of the variables that are listed after it, but will not be affected by them contemporaneously. In our case, since we wish to analyze the impact of a EPU shock on the macroeconomic variables and since our uncertainty measures are news-based, it seems reasonable to consider them to be exogenous and to order them first. Consequently, we adopt the following order of the variables for Canada : EPU-NLP (or EPU-Boolean), TSX, employment, industrial production, and GDP. For the US, the order of the variables is: EPU-NLP (or EPU Boolean), S&P 500, employment, industrial production, and consumption.[14] [15]

In accordance with the optimal lag criteria SBIC and HQIC we choose three lags for Canada and one lag for the US.

Before presenting the impulse response functions, we first present a comparison of the EPU-NLP and EPU-Boolean. As mentioned earlier, to highlight the benefits of using the NLP approach, we also constructed an index on the basis of the same data-set, using a strictly Boolean approach. As shown in Figure 12, EPU-NLP better captures the large increase in uncertainty in March and April 2020 due to COVID-19 compared to EPU-Boolean for both Canada and the US. We then estimate the response of the economic variables to an uncertainty shock as captured by a one-standard-deviation (SD) innovation in EPU-NLP and EPU-Boolean respectively.

The impulse response functions for Canada for the period including COVID-19 (Jan. 2015-Oct 2020) with the EPU-NLP and EPU-Boolean indices, are shown in Figure 13. The Canada SVAR results indicate that a one-SD policy uncertainty innovation with EPU-NLP leads to declines of 1.04 % in real GDP, 1% in industrial production, 0.95 % in employment, and 1.08 % in TSX respectively. By contrast, a one SD policy uncertainty innovation with EPU-Boolean results in declines of only 0.42% in real GDP, 0.33 % in industrial production, 0.41% in employment, and 1.02% in TSX.

Similarly for the US, we find that a one-SD EPU-NLP shock has a more pronounced effect than a one-SD EPU-Boolean shock on most of the economic variables. The impulse response functions for the US for the period including COVID-19 with EPU-NLP and EPU-Boolean are shown in Figure 14. One-SD shock to uncertainty with EPU-NLP results in a 0.90% drop in industrial production, 0.70% in real consumption, 0.83% in employment, and 2.1% in S&P 500. On the other hand, a one-SD shock to uncertainty with EPU- Boolean provokes only a 0.19 % fall in industrial production, 0.11 % in real consumption, 0.16 % in employment, and 0.60 % in S&P 500. Hence, for both the US and Canada, we observe a less pronounced response to a one-SD shock when we use the EPU-Boolean instead of the EPU-NLP.

We also follow (Altig et al., 2020), and compare the impact of a one-SD innovation with EPU-NLP for the period *including* COVID-19 (Jan. 2015-Oct-2020) with the *pre-COVID-19* period (Jan. 2015-Dec. 2019) for Canada and the US. The impulse response functions are shown in Figure 15 for Canada and Figure 16 for the US. These results are striking. For Canada, for the period including COVID-19, a one-SD innovation with EPU-NLP leads to declines of 1.04 % in real GDP, 1% in industrial production, 0.95 % in employment, and 1.08 % in TSX respectively. By contrast, for the pre-COVID-19 period, a one SD innovation with EPU-NLP results in no change in real GDP, 0.17 % in industrial production, 0.025% in employment, and 0.6% in TSX.

The results are similar for the US. For the period including COVID-19, one SD shock to uncertainty (with EPU-NLP) results in a drop of 0.90% in industrial production, 0.70% in real consumption, 0.83% in employment, and 2.1% in S&P 500. By contrast, for the pre-COVID-19 period, one SD shocks to uncertainty (with EPU- NLP) provokes a fall of only 0.06 % in industrial production, 0.05% in real consumption, 0.015 % in employment, and 0.34 % in S&P 500.

In other words, for both the US and Canada, we observe a less pronounced response to a one-SD shock to uncertainty (with EPU-NLP) for *the pre-COVID-19 period.* Hence, the EPU-NLP index is able to capture

[14] We use industrial production at the monthly frequency as a proxy for real GDP in the case of the US since real GDP is not available on a monthly basis for the US.

[15] Since our primary purpose is not to make a comparison of the EPU-NLP^{Canada} and EPU-NLP^{US} indices, using different variables is not consequential.

the COVID-19-induced uncertainty and its severe negative impact on economic variables in the both Canada and the US.

We also present a graph of the *weekly* EPU-NLP for Canada in Figure 10. Interestingly, we observe peaks and troughs in the index corresponding to all major events that created uncertainty, such as investor fears of recession, the 2015 recession, Trudeau's stimulus package in January 2016, the slump in the energy sector, Trudeau's visit to the US in February 2017, the job loss in Alberta, the G-7 meeting in Canada, as well as COVID-19, and other events. We present a similar graph of the weekly index for the US in Figure 11 and again observe how key uncertainty-generating events (such as elections, the US-China trade-war and COVID-19 among others) are captured by the index.

We also compared EPU-NLP to other uncertainty measures such as VIX as well as to BBD-EPU. As is shown in Figure 17 , EPU-NLP and VIX closely follow each other, with matching peaks and troughs during the COVID-19 period for Canada and the US. The results for the correlation are shown in Table 5. We observe a correlation of 0.85 between $EPU-NLP^{Canada}$ and VIX, and of 0.80 between $EPU-NLP^{US}$ and VIX. (Baker et al., 2016) found a correlation of 0.58 between the EPU and VIX. As they explain, their EPU is more specialized on policy uncertainty as opposed to financial uncertainty captured by the VIX. They therefore developed an EMV (equity market volatility) index that better captures financial uncertainty. In our case, in view of our search words, the EPU-NLP may be more attune to generalized uncertainty, and COVID-19 uncertainty in particular. This may be an explanation of its closer correlation we obtain between the EPU-NLP and the VIX. We find a correlation of 0.79 between BBD's EMV index and our $EPU-NLP^{Canada}$ and of 0.70 between the EMV index and $EPU-NLP^{US}$. In addition, the correlation between the $EPU-NLP^{US}$ index and BBD-EPU^{US} is 0.85 and that between $EPU-NLP^{Canada}$ and EPU-NLP^{Canada} is 0.72.

5 Conclusion

This paper described a new approach, based on text mining and natural language processing techniques, towards constructing an EPU index based on newspaper articles. For this purpose, we use the RAKE, RoBERTa/SBERT, and GrapeNLP algorithms. RAKE is used to determine the popularity of various words or phrases that are to be used to filter articles and develop search queries and grammars. We use the RoBERTa algorithm which is pre-trained on large a large news dataset (CC-News). We further fine-tune it on our own newspaper data. However, since RoBERTa is not well suited for semantic searches, we combine RoBERTa with the Sentence-BERT (SBERT) algorithm. Finally, we use the GrapeNLP grammar engine to select the final EPU-related articles on the basis of which we calculate our EPU-NLP index.

We compare the EPU-NLP index with another EPU which we construct on the basis of the same dataset using a strictly Boolean approach. We observe that EPU-NLP better captures the COVID-19 induced uncertainty than the EPU-Boolean. We also compare the EPU-NLP with other leading uncertainty indices and find that it is closely correlated with several of them (BBD-EPU, Equity Market Volatility (EMV), the CBOE Volatility Index (VIX)). We further assessed the impact of EPU-NLP and EPU-Boolean using a SVAR model with Canadian and US economic variables. We found that EPU-NLP created a greater dip in economic variables compared to EPU-Boolean. Lastly, we compare the impact of a one-SD EPU-NLP shock on pre-COVID-19 data (Jan. 2015-Dec. 2019) with its impact on the entire data range including COVID-19 data (Jan. 2015-October 2020). The SVAR results showed once more that EPU-NLP generated a greater dip in economic variables for Canada and the US for the period including COVID-19 than for the pre-COVID-19 period.

References

Altig, D., S. Baker, J. M. Barrero, N. Bloom, P. Bunn, S. Chen, S. J. Davis, J. Leather, B. Meyer, E. Mihaylov, et al. (2020). Economic uncertainty before and during the covid-19 pandemic. *Journal of Public Economics 191*, 104274. 2, 9, 10

Altug, S., F. S. Demers, and M. Demers (2009). The investment tax credit and irreversible investment. *Journal of Macroeconomics 31*(4), 509–522. 2

Baker, S., N. Bloom, S. Davis, and S. Terry (2020). Covid-induced economic uncertainty (no. w26983). 9

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The quarterly journal of economics 131*(4), 1593–1636. 2, 3, 4, 8, 11

Barz, B. and J. Denzler (2020). Deep learning on small datasets without pre-training using cosine loss. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1371–1380. 6

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The quarterly journal of economics 98*(1), 85–106. 2

Caballero, R. J. and E. M. Engel (1999). Explaining investment dynamics in us manufacturing: a generalized (s, s) approach. *Econometrica 67*(4), 783–826. 2

Demers, M. (1991). Investment under uncertainty, irreversibility and the arrival of information over time. *The Review of Economic Studies 58*(2), 333–350. 2

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*. 2, 5

Dixit, A. and R. Pindyck (1994). Investment under uncertainty. princeton university press, princeton, nj. 2

Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature 57*(3), 535–74. 2

Gross, M. (1997). *The Construction of Local Grammars.* in E. Roche and Y. Shabes, eds., Finite-state language processing, Cambridge, Mass.: MIT Press, pp.329-354. 7

Hamilton, J. D. (2018). Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics 100*(5), 831–843. 9

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216. 2, 9

Liu, Y., M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*. 2, 6

Moran, K., D. Stevanović, and A. K. Touré (2020). *Macroeconomic uncertainty and the covid-19 pandemic: Measure and impacts on the canadian economy.* CIRANO. 2, 9

Paumier, S., T. Nakamura, and S. Voyatzi (2009). Unitex, a corpus processing system with multi-lingual linguistic resources. *eLEX2009 173*. 2, 7

Reimers, N. and I. Gurevych (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*. 2, 6

Sastre, J. (2011). *Efficient finite-state algorithms of application of local grammars.* Ph. D. thesis. 2, 4, 7

Sastre, J., A. H. Vahid, C. McDonagh, and P. Walsh (2020). A text mining approach to discovering covid-19 relevant factors. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 486–490. IEEE. 4, 7

6 Appendix

Table 1: Canadian newspapers

Montreal Gazette	National Post	Toronto Star	Edmonton Journal
Vancouver Sun	Globe and Mail	Wind Star	The Star Phoenix
Ottawa Citizen	CBC News	Calgary Herald	Leader Post
Financial Post	National Observer	Hill Times	

Table 2: US newspapers

USA Today	Los Angeles Times	Wall Street Journal	New York Times
The Dallas Morning News	Star Tribune	San Francisco Chronicle	

Table 3: List of words from RAKE used for the Lucene Index search

economy OR economic OR economies OR uncertain OR uncertainty OR uncertainties OR policy OR policies OR regulation OR regulations OR tax OR taxes OR "central bank" OR deficit OR budget OR spending OR coronavirus OR virus OR covid19 OR challenging OR "trade war" OR "trade wars" OR election OR recession OR monetary OR fiscal OR unclear OR unpredictable OR unemployment OR damage OR devastation OR social OR investment OR production OR investor OR investors OR inflation OR consumption OR hiring OR afloat OR firms OR companies OR tariff OR growth OR "labor demand" OR lockdown OR lockdowns OR upheaval OR "economic damage" OR concern OR concerns OR "aid package" OR stimulus OR transfers OR "tax relief" OR lingering OR heightened OR political OR deepened OR sow OR pandemic OR prolonged OR national OR nationwide OR strictest OR full OR increased OR "financial crisis" OR "economic crisis" OR "job losses " OR "economic recovery" OR crisis OR unprecedented OR debt OR vaccine OR "white house" OR "federal reserve" OR congress OR "stimulus package" OR election OR "relief plan" OR elections OR "presidential elections" OR "political polarization" OR polarized'

Table 4: Queries for the sementic search with the RoBERTa sentence transformer

queries = ['there is economic policy uncertainty due to coronavirus crisis', 'government creating and fuelling uncertainty with economic policies in response to virus crisis', 'uncertain situation creating measures of economic distress and economic crisis', 'unclear measures distressing households and investors', 'economic stimulus package lacks clear implementation creating uncertain' , 'unprecedented levels of rise in trade policy uncertainty', 'many do not know whether they will get unemployment benefit next month', 'tax relief is not sufficient to decrease the impact of crisis' ,'social protection and social assistance needed to make ends meet', 'health policy causing uncertainty' , 'real risk of higher taxes in the immediate future is greater concern of household and business', 'whether taxes might be raised by how much', 'deficits created uncertainty about future taxes', 'coronavirus induced recession', 'government s indecisive handling of the shut down canada', 'they came amid a backdrop of global trade policy uncertainty and concerns about china s outlook', 'the outlook for the global economy is worrying, and policy-makers should be concerned', 'the future direction of trade policies and global demand conditions remains highly uncertain', 'investors are more likely to postpone, amid fears of job losses and lacking in clarity policy response', 'many fear due to greater economic uncertainty and financial crisis in canada', 'The growing virus fears in these unprecedented times']

Table 5:

	Indices		
	VIX	EMV_{US}	$BBD-EPU_{US}$
$EPU-NLP_{US}$	0.80	0.70	0.85
	VIX	EMV_{US}	$BBD-EPU_{CAN}$
$EPU-NLP_{CAN}$	0.85	0.79	0.72

Figure 5: Causal or Forward for Canada

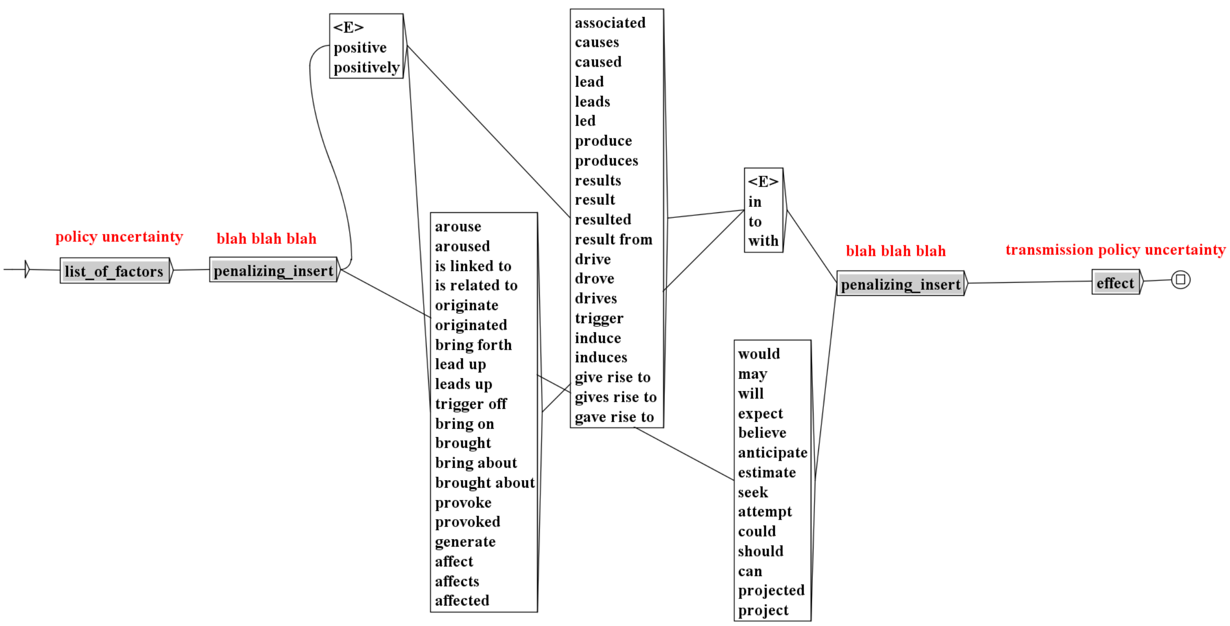

Figure 6: Policy or Covid for Canada

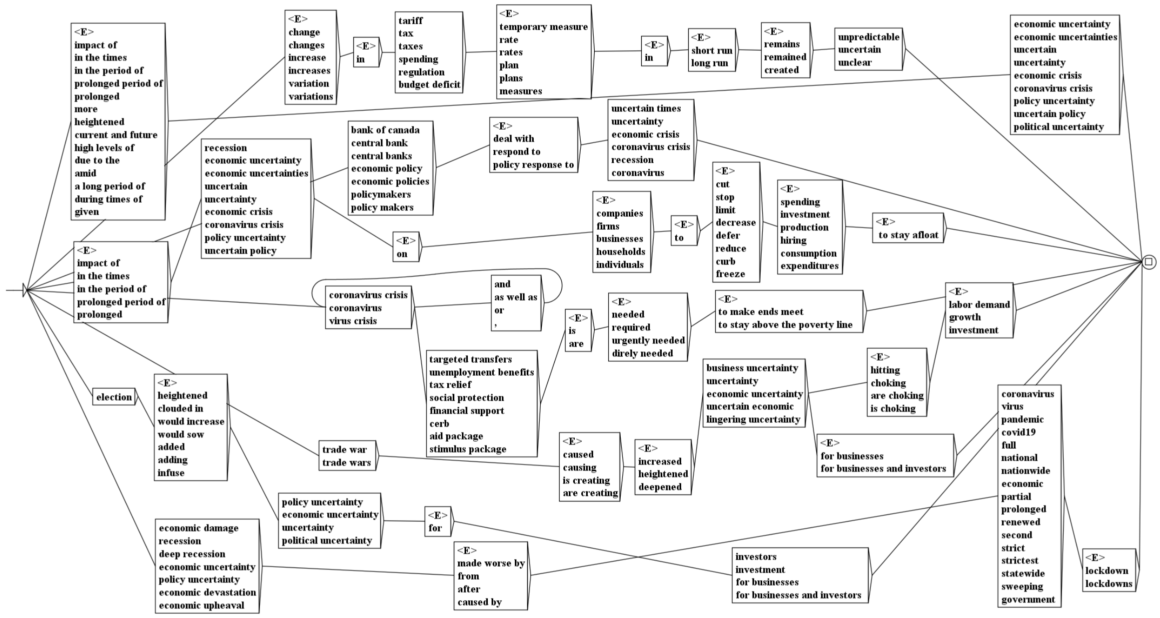

Figure 7: Effects for Canada

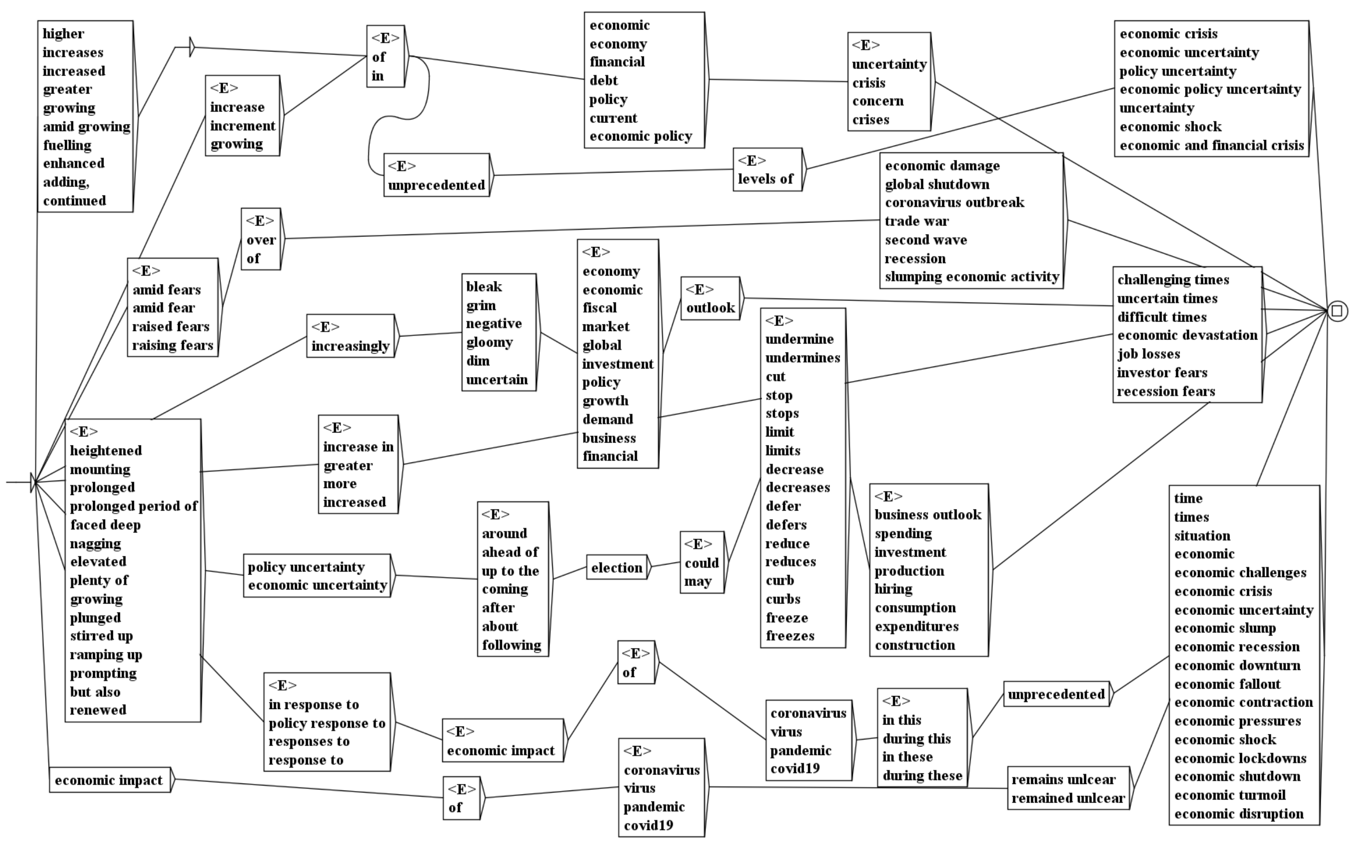

Figure 8: Policy or Covid for US

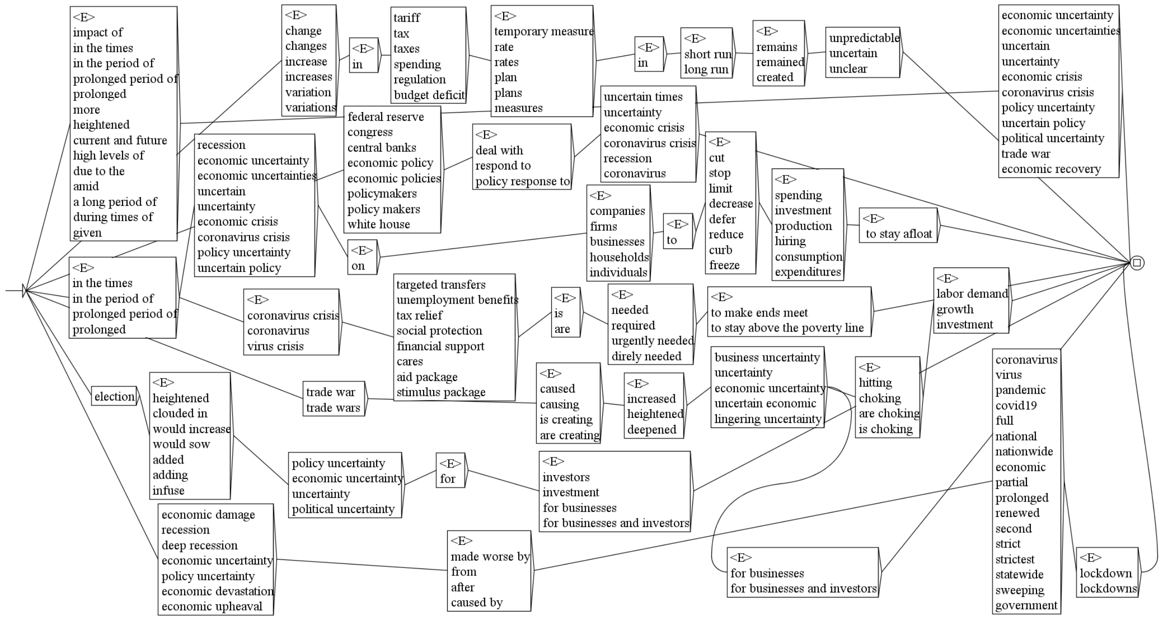

Figure 9: Effects for the US

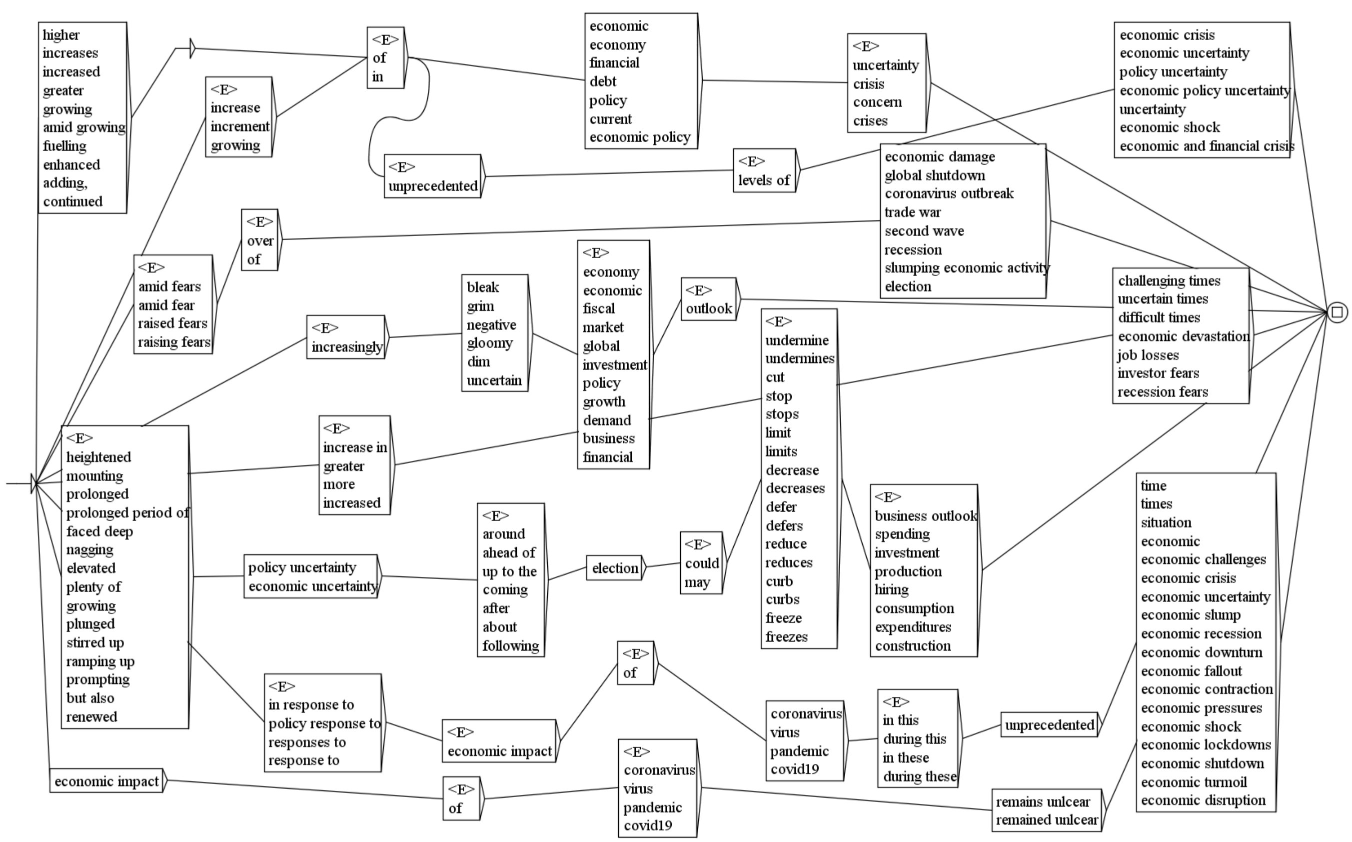

Figure 10:

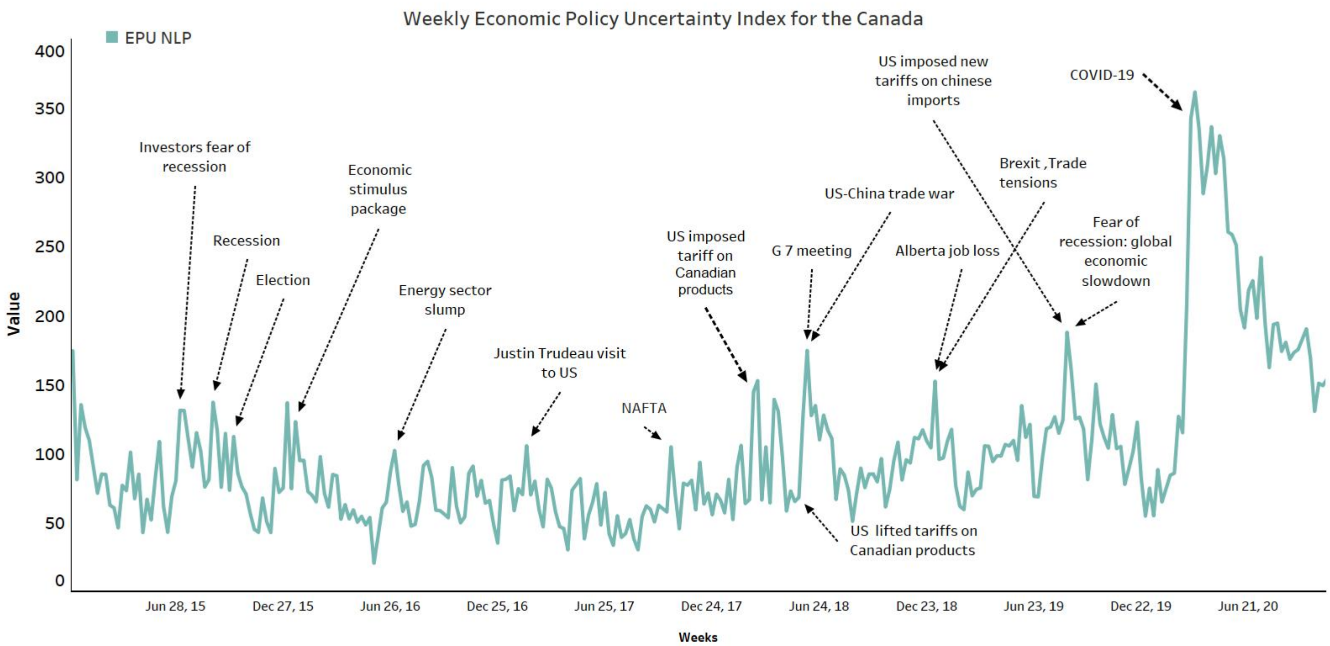

Figure 11:

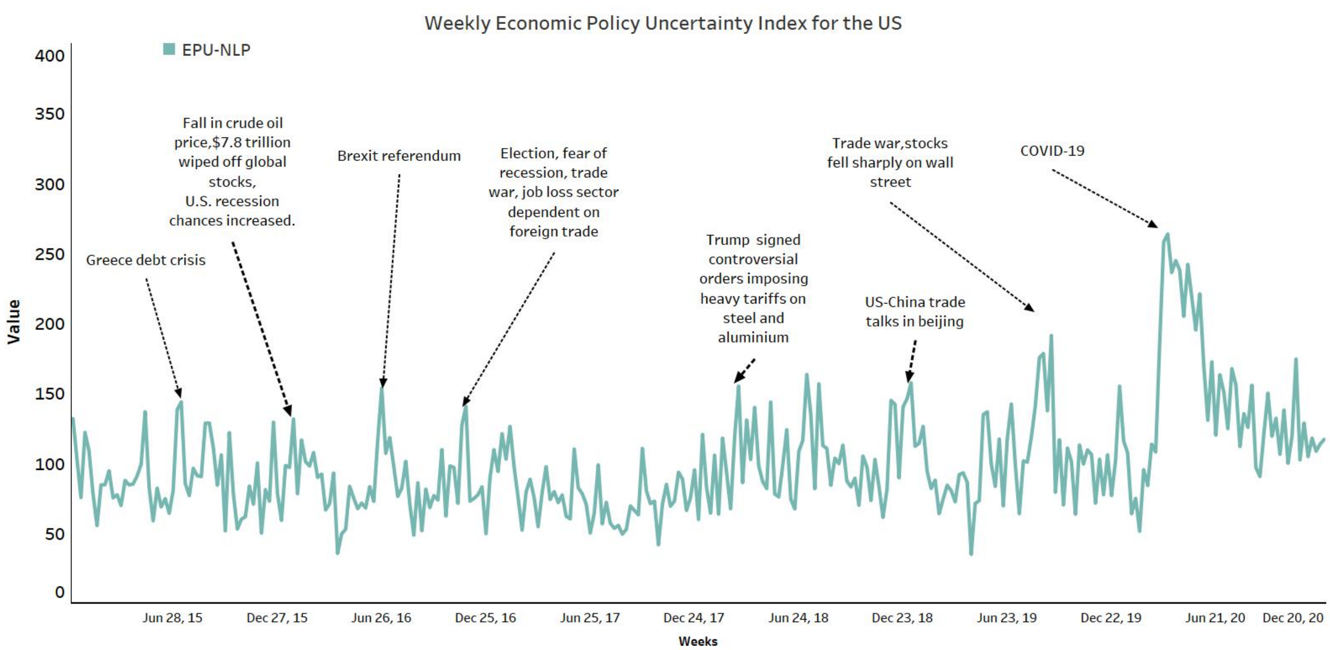

Figure 12:

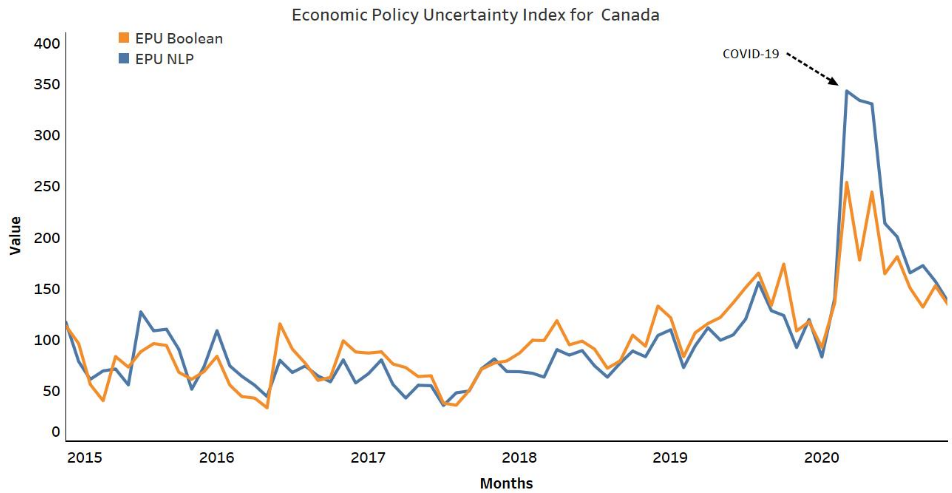
Economic Policy Uncertainty Index for Canada
EPU Boolean
EPU NLP
COVID-19
Value
400
350
300
250
200
150
100
50
0
2015
2016
2017
2018
2019
2020
Months

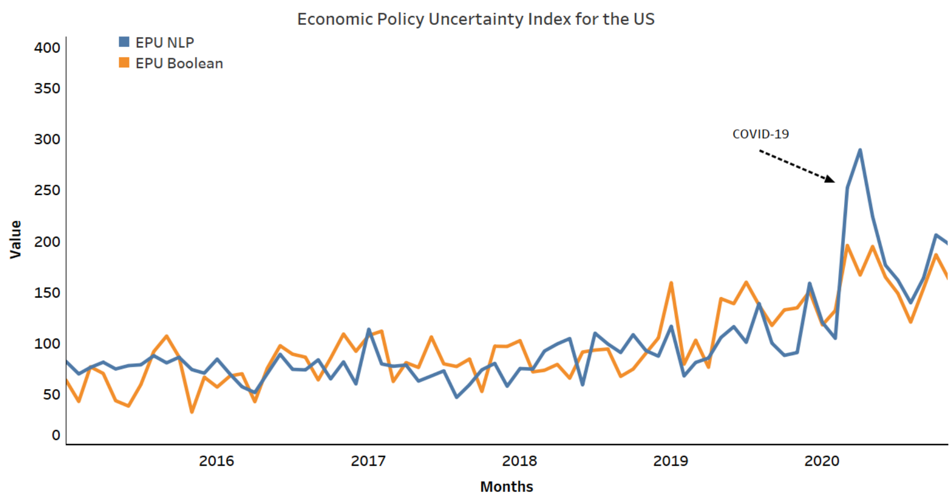
Economic Policy Uncertainty Index for the US
EPU NLP
EPU Boolean
COVID-19
Value
400
350
300
250
200
150
100
50
0
2016
2017
2018
2019
2020
Months

Figure 13: SVAR results, Canada: Comparing EPU-NLP and EPU-Boolean

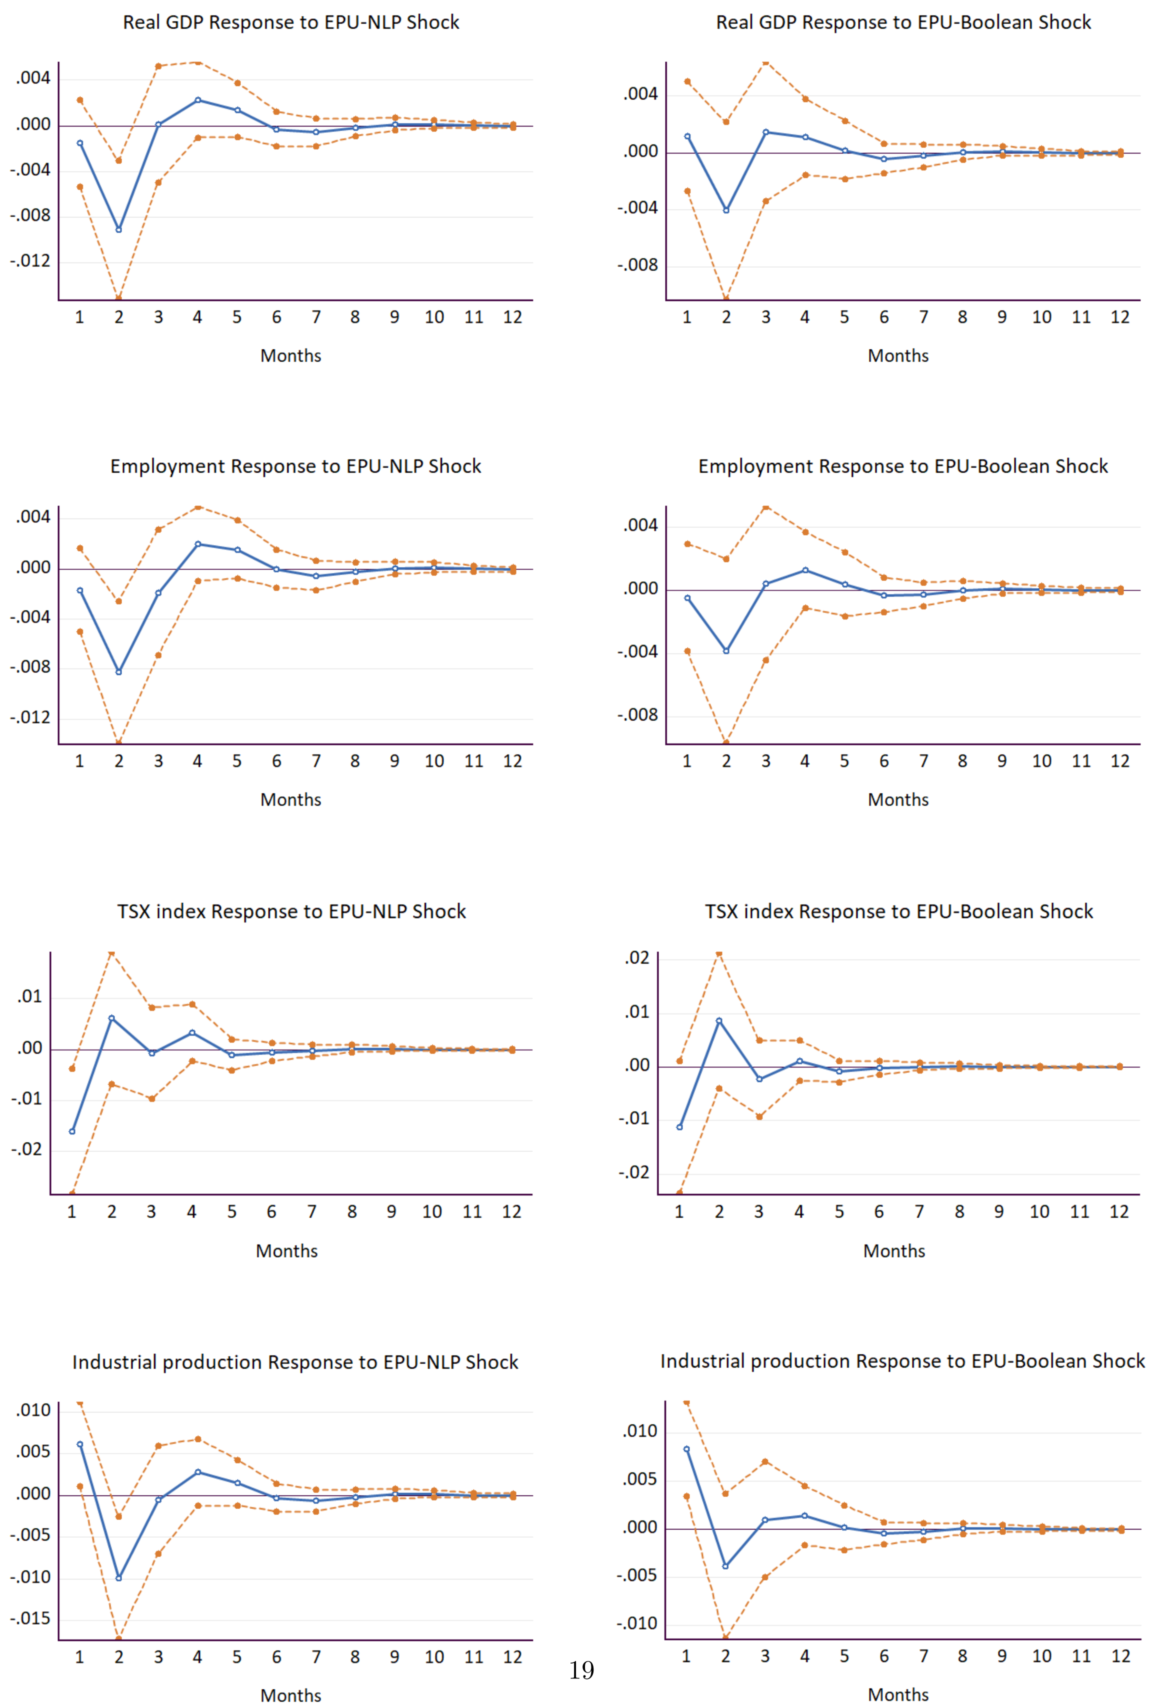

Figure 14: SVAR results, US: Comparing EPU-NLP and EPU-Boolean

Figure 15: EPU-NLP shock: period including COVID-19 vs pre-COVID-19 for Canada

Period including COVID-19

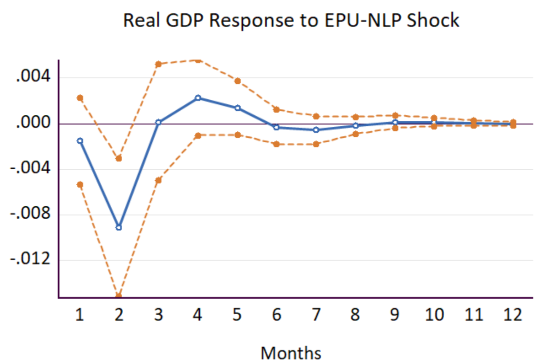

pre-COVID-19

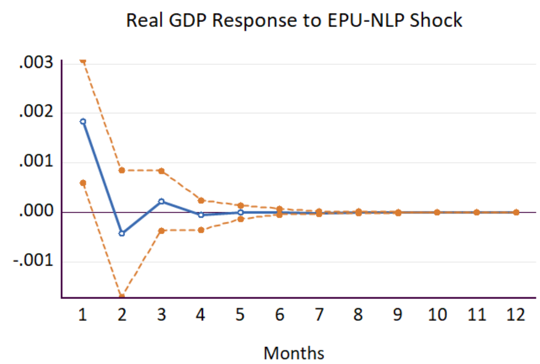

Period including COVID-19

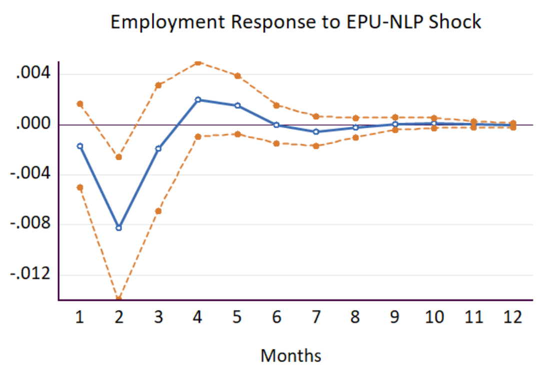

pre-COVID-19

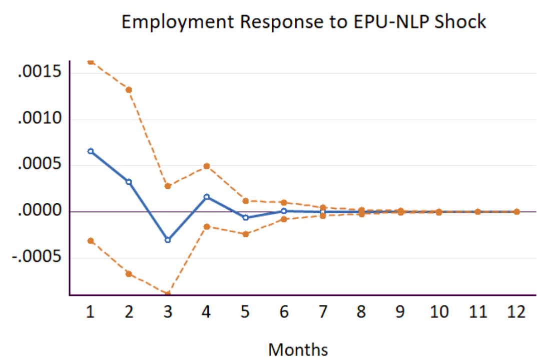

Period including COVID-19

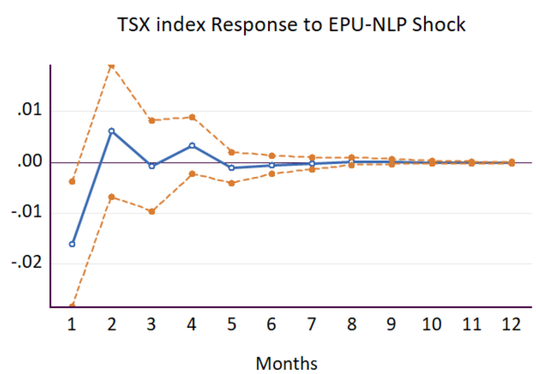

Pre-COVID-19

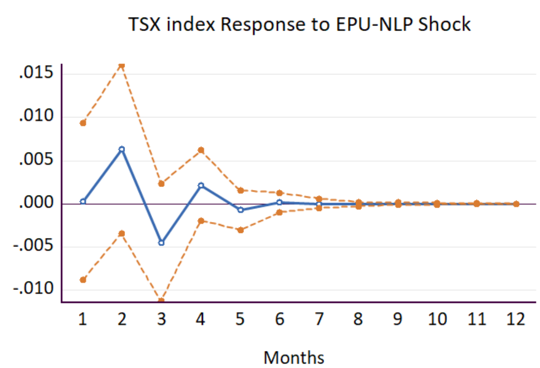

Period including COVID-19

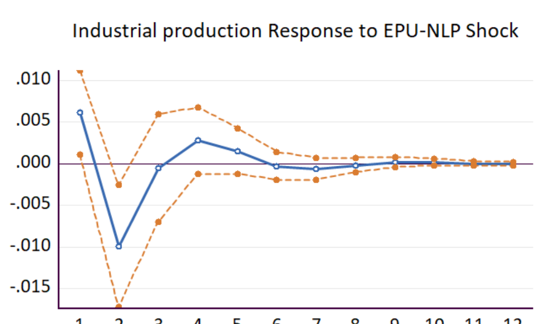

Pre-COVID-19

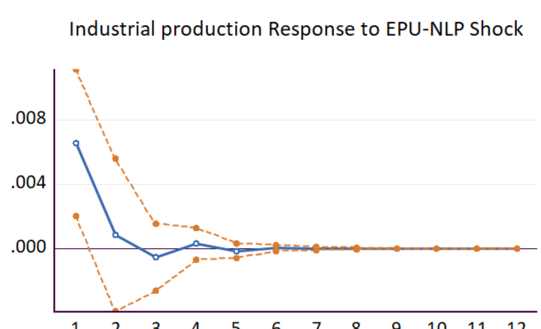

Figure 16: EPU-NLP shock: period including COVID-19 vs pre-COVID-19 for the US

Period including COVID-19

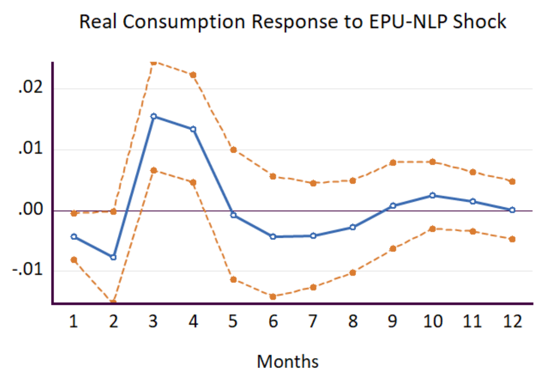

Pre-COVID-19

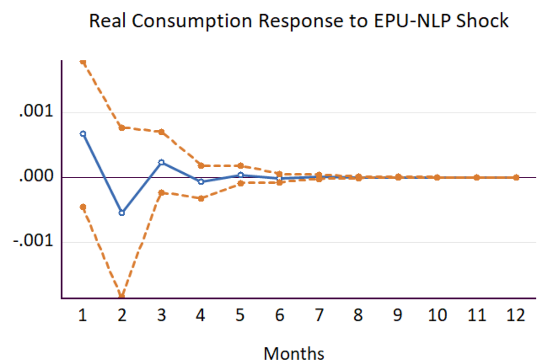

Period including COVID-19

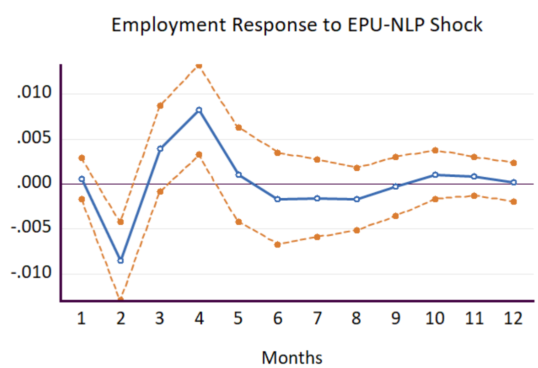

Pre-COVID-19

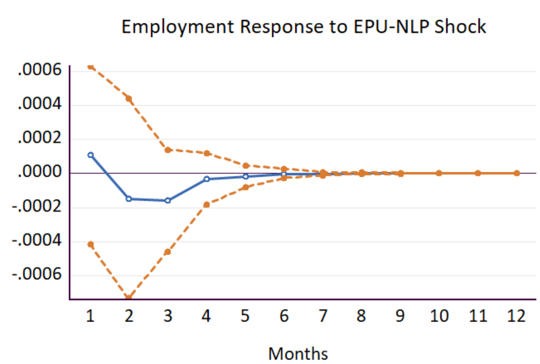

Period including COVID-19

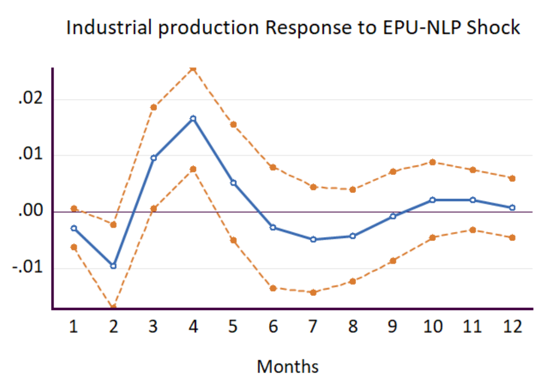

Pre-COVID-19

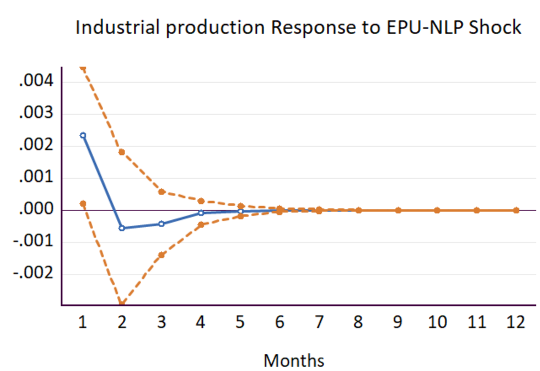

Period including COVID-19

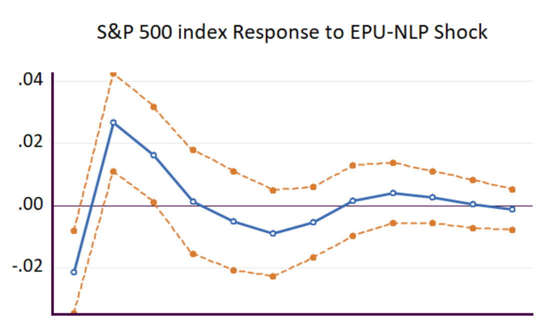

Pre-COVID-19

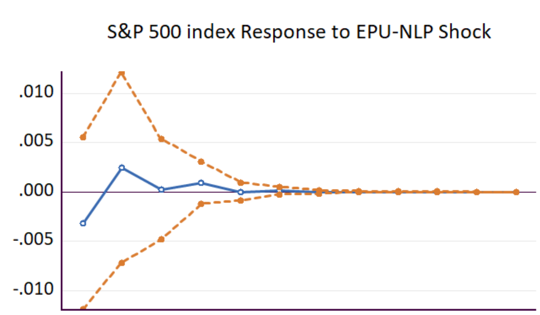

Figure 17: EPU-NLP compared to market-based VIX Index for the US

US

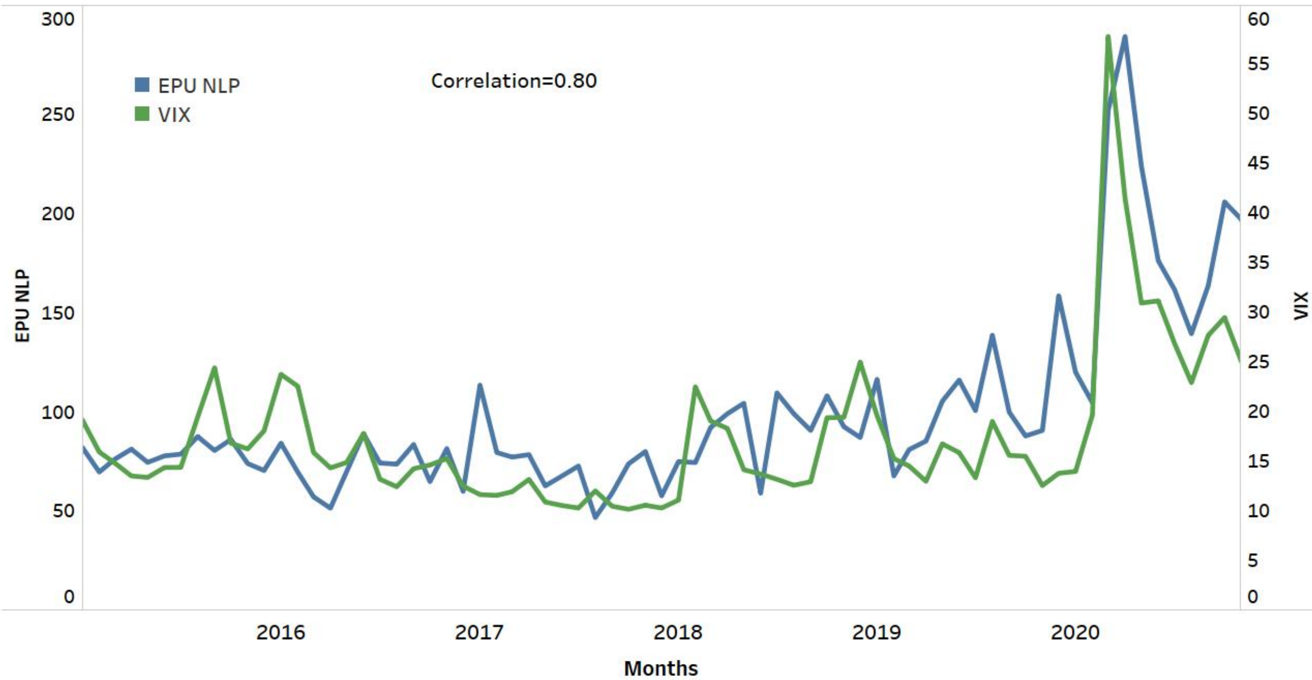

Canada