Economic uncertainty and natural language processing; the case of Russia

Wojciech Charemza — Vistula University, Poland & University of Leicester, UK
Akademia Finansów i Biznesu Vistula, Stokłosy 3
02-787 Warsaw, Poland
wch@le.ac.uk (corresponding author)

Svetlana Makarova — University College London, UK & Vistula University, Poland
University College London, SSEES
16 Taviton St.
London WC1H 0BW, UK

Krzysztof Rybiński — Vistula University, Poland & EMLYON Business School, France
Akademia Finansów i Biznesu Vistula, Stokłosy 3
02-787 Warsaw, Poland

This research has been financed by the National Science Centre in Poland Opus 15 project 2018/29/B/HS4/01462, "*Construction and testing of country-specific measures of economic uncertainty*"

Abstract

The paper proposes a method of constructing text-based country-specific measures for economic policy uncertainty. To avoid problems of translation and human validation costs, we apply natural language processing and sentiment analysis to construct such measures for Russia. We compare our measure with that developed earlier using direct translations from English and human validation. In this comparison, our measure does equally well at evaluating the uncertainty related to key events that affected Russia between 1994 and 2018 and performs better at detecting the effects of uncertainty in Russia's industrial production.

Keywords: Economic policy uncertainty, natural language processing, sentiment analysis, Russian economy.

Acknowledgement

We are very grateful to the Editor and the Referees of Economic Analysis and Policy for their valuable comments and suggestions. We also thank Ben Noble for stimulating discussions that led us to correct the earlier version of the paper and to Robin Hazlehurst for helpful comments and suggestions. Needless to say, they are not responsible for any deficiencies in this paper.

Economic uncertainty and natural language processing; the case of Russia

1. Introduction

We propose a method of constructing a country-specific text-based measure of economic and policy uncertainty that can account for linguistic and sentiment particularities. We apply this methodology to Russian data. The essence of text-based measures of uncertainty lies in counting the frequency of terms in press articles that indicate uncertainty in the context in focus. Such measures, usually presented in the form of time-varying indices, first became available in 2012 based on press articles written in English and are described in Baker et al. (2016), where the methodology for constructing economic policy uncertainty (EPU) indicators was proposed. It has been found that this type of uncertainty is an important factor for affecting various economic and related phenomena, like economic cyclicality (Bloom, 2014; Bloom et al., 2018; Gu et al., 2021), investment climate (e.g. Liu and Zhang, 2015; Peng et al., 2018; Yang and Hamori, 2021), money demand (Hossain and Arwatchanakarn, 2020; Bahmani-Oskooee et al., 2016; Bahmani-Oskooee and Maki Nayeri, 2021), housing (Bahmani-Oskooee Ghodsi, 2017), investment fraud (Hou et al., 2021), and carbon emissions (Adams et al., 2020; Atsu and Adams, 2021). For a comprehensive review of other related findings, see Al-Thaqeb and Algharabali (2019).

These text-based economic indicators use articles in the press or other media as their source of data and are often built using large data sets. The EPU indices are created by first identifying the topic of an article that has at least some economic context and then by searching for words and phrases in it that can be categorised as having meanings related to *economic*, *policy*, and *uncertainty*. We later call such words *descriptors* and denote them as {economic}, {policy} and {uncertainty} accordingly.

Since the first EPU index based on English-language media sources appeared, interest has grown in constructing and applying similar indices for different countries and regions. Two approaches have so far been applied for evaluating uncertainty in non-English speaking countries. The first is to search English-language newspapers for articles concerning a given country or region and look for EPU descriptors in them. There is already some evidence that using English-language newspapers in country-focused uncertainty studies for non-English language countries might cause bias and fail to identify some topics. Huang and Luk (2018), for instance, present the disadvantages of using English-language indices to analyse uncertainty in China, and similar results have been obtained for Turkey (see Kılıç, 2021; however, Jirasavetakul and Spilimbergo, 2018, provide some arguments to the contrary).

The second approach is to translate the English language descriptors into a local language and search articles published in that language. This approach has been used by Baker, Bloom and Davis among others, and they publish EPU indices at https://www.policyuncertainty.com/ for many countries using searches of the local language newspapers; they currently cover the USA, Australia, Brazil, Canada, Chile, China, Colombia, Croatia, France, Germany, Greece, Hong Kong, India, Ireland, Italy, Japan, South Korea, Mexico, Netherlands, Russia, Singapore, Spain, Sweden and the UK. This method might also be not fully adequate because differences in journalistic styles, conventions and writing patterns, and also in the perception and expectations of readers, might cause the essence of uncertainty to be distorted (for a comparison of the reporting systems of journals in Eastern and Western European countries, see e.g. Wessler et al., 2008; for a similar analysis for Western European countries only see Esser and Umbricht, 2013). In an extreme case, the reporting differences might result in the descriptors {economic}, {policy} and {uncertainty} becoming uninformative as they might be crowded out by a different word or phrase that carries a similar message but has not been identified in advance as belonging to a corresponding set of descriptors. Such competing sets might come from expressions of negative sentiment that appear in articles on economic matters and that might in the local language carry a message that suggests economic and political uncertainty. Consequently the original EPU methodology requires human validation of newspaper articles for non-English speaking countries in order to avoid bias and inefficiency. The human validators must have relevant linguistic and economic knowledge, and must be aware that in some languages, the sense of uncertainty often comes from the sentiment expressed by sentences. This possible problem of loss in translation has been noticed before and is usually solved by enhancing and adjusting the sets of descriptors (see e.g. Ghirelli et al., 2019, 2021 for construction of the economic policy uncertainty index for Spain).

To minimise, or even avoid, these difficulties, we apply a pragmatic and cost-saving approach based on natural language processing (NLP). We do this by (i) applying word embedding methods to create vocabularies of words that describe economic, policy, and uncertainty in a way that is specific to the local language; (ii) using topic modelling to detect the leading topics of press articles so that we can eliminate articles whose main themes are not related to economics; and (iii) weighting the selected text articles by their sentiment weights, which represent the intensity of positive and negative words in the articles selected. We compare our measure with that developed earlier by direct translations from English and human validation.

These concerns might be particularly relevant for the case of Russia and the Russian language. The differences between Russian and English are large but not extreme (see, e.g. Chiswick and Miller, 2005; or Isphording and Otten, 2013, for a quantitative evaluation of this distance). In this way, our results may provide a convenient starting point for studies of uncertainty in different countries that use different languages. While the linguistic distance might not be great, several studies show there to be substantial contextual and emotional differences between Russian and other languages, including English. Especially striking are the results of Jha et al. (2021), who analyse sentiment towards finance in books published in Chinese, English (UK and US separately), French, German, Italian, Spanish and Russian in the period 1870-2009. Of all these languages, the one in which the sentiment expressed towards finance was clearly the lowest was Russian. Furthermore, linguistic difficulties might arise because there may be sociological, political or even psychological reasons why uncertainty might not be evident in articles merely from counting the appearances of words related to it (for a discussion of the specifics of Russian media models, see Vartanova, 2012; and to an extent, Petersson and Persson, 2011). It is also possible that self-censorship and external pressure factors might restrain journalists in Russia from using particular expressions and phrases, which might also distort the assessment (see Schimpfössl and Yablokov, 2020). Such complexity might push the costs of proper human validation prohibitively high if the language subtleties, media specifics and cultural differences are to be accounted for without bias being caused in the understanding of the undertones.

We use data for four countrywide Russian newspapers that are available electronically. Our data series begins in 1992 with the first electronic availability of the newspaper *Kommersant* (the other three newspapers were added subsequently), and it ends in 2018. This means we cover a turbulent period of Russia's history that takes in the Yeltsin era, Putin's presidencies, Medvedev's interim period, the first and second Chechen wars, the Russo-Georgian armed conflict, the annexation of Crimea, and the Donbas conflict. However, we do not include the Covid-19 pandemic as uncertainty has to be measured differently for this period (see, e.g. Altig et al., 2020; Charemza et al., 2020).

The plan of the paper is as follows. After the introduction given in Section 1, Section 2 briefly outlines the methodology applied and describes how our uncertainty measures are constructed. Section 3 discusses our uncertainty index for Russia in more detail. Section 4 evaluates our index in comparison to the EPU index of Baker et al. (2016) constructed for Russia by checking whether the periods of increased uncertainty identified by the indices match particular events in Russia that might cause uncertainty to increase. In this comparison, both indices do equally well, as they match 90%-92% of the events that have the potential to raise uncertainty. Section 5 compares the estimated effects of the uncertainty measured by both indices on industrial production in Russia. Section 6 discusses various robustness results. Section 7 concludes and discusses ways the research could be developed further. The paper contains two Appendixes. Appendix A contains some descriptive statistical measures for comparing the indices, and in Appendix B, we give a list of a priori identified events considered in Section 4.

2. Methodology and the related literature

The original EPU methodology by Baker et al. (2016) consists of the following four steps:

Step 1: Define the descriptors {economic}, {policy} and {uncertainty} by translating the descriptors from English to a local language.

Step 2: Search for all the articles in the database to identify those containing at least one word from each set of descriptors.

Step 3: Do a human validation of a sample of articles selected from the entire dataset of articles, often called the corpus, to confirm that the selection does not contain articles on unrelated subjects.

Step 4: By aggregating the count of the selected articles linearly, or summing it up, across the newspapers and across time and scaling it, build an index reflecting the frequency of the newspaper articles that concern economic policy uncertainty within the total number of articles.

To deal with the problems caused by linguistic differences and to minimise the need for human validation, we modify Steps 1, 3 and 4 of this methodology. Step 1, where the descriptors are defined by translating the corresponding words from English into a local language, might result in clarity and context being lost, and in consequence, it may cause bias. Instead of translating, we feel that using natural language processing rather than dictionary comparison to define sets of words that are similar to the English language descriptors would help identify the economic policy uncertainty context more precisely. Human validation in Step 3 can in practice, only be done on a small sample of articles, and the selection depends on the quality of the work of the validators, and to an extent, on their subjective judgements. This might create heterogeneity and some bias in the selection of articles. This is why we apply machine learning to identify the leading topics of the articles and then eliminate those that are not relevant. In Step 4, linear aggregation of the selected articles using equal weights ignores the feelings and emotions conveyed in them, but these are clearly important for evaluating the strength and magnitude of the uncertainty. By applying the sentiment-related weights we diversify the articles by the power and clarity of the messages of uncertainty carried by particular articles.

We prepare the data by applying the 'bags of words' approach, where we treat each text as a set of words, ignoring the order of them. We also remove punctuation, white spaces, special characters, stop words and digits. Next, we convert full words in the 'bags of words' to their stems. In English, the differences between stem-based and non-stem-based natural language processing can easily be identified and the results compared, but this is practically impossible for Slavic languages, including Russian, at the current stage of development of machine searches. These languages have complex declinations and conjugations so that words with the same meaning may have many different endings, which requires that lemmatisation or stemming is performed during the text pre-processing. For an introduction to the algorithmic approach to stemming, which is applied here, see, e.g. Manning et al. (2008). The data preparation or text pre-processing is conducted using the `quanteda` library in R (see https://quanteda.io/). For a description and evaluation of `quanteda` see Benoit et al. (2018). A possible disadvantage of using stems is overstemming, where two different words have an identical stem. This may, to an extent, affect the quality of the search. The possible effects of stemming in our analysis are further discussed in Section 6.

The four steps of our approach are then as follows:

Step 1: As in the original EPU methodology, we define the sets of stems for search as the descriptors for {economic}, {policy} and {uncertainty}. Instead of translating the descriptors from English, we apply one of the simplest word embedding methods, *Word2vec* (see Mikolov et al., 2013). The descriptors that we obtain as a result are those that fit best in their cosine similarity to the words 'uncertainty', 'economic' and 'policy' translated into Russian[1]. Word embedding is an encoded representation of a word and its context as a vector. Cosine similarity is a measure that is similar to the correlation coefficient used for evaluating the relation between two vectors (for a description and comparison with other methods, see, e.g. Sidorov et al., 2014). There are alternatives to Word2vec that each has its own advantages and disadvantages (see, e.g. Naili et al., 2017, for a comparison with other word embedding methods), and there are also generalisations of Word2vec, most notably BERT (see Devlin et al., 2018). We choose Word2vec for its relative simplicity and also because the Russian language dictionary can be accessed much better by Word2vec than by the alternatives. We download the Russian Word2vec data from https://fasttext.cc/docs/en/crawl-vectors.html and create descriptors with the stems that have the highest similarity to the terms 'uncertainty', 'economics' and 'policy'. The calculations were conducted in Python with the help of the `genism.models` library. We use the `ru_model=KeyedVectors.load_word2vec_format()` function to load the embeddings and the `ru_model.most_similar()` method to find the most similar

[1] Some examples of stemmed descriptors obtained by Word2vec methods, translated into English, are: economic: {ceconomic, macroeconomic, business, inflatio, monet, demograph, industr, branch, financ, unemploy, coniunctur, oil}; policy: {politic, geopolitic, dictator, ideolog, protectionist, reform, propagand, anticorrupt, populist, concept}; uncertainty: {uncertain, unclear, unstabl, unforcastab, ambigu, blurr, contradict, mess, tens}.

words. Word embedding and topic modelling have already been used in other applications to construct uncertainty measures, see Aromí (2017) and Davis et al. (2019).

Step 2: After defining the descriptors using Word2vec, we conduct dictionary-based text mining of the corpus with the goal of finding all the articles that contain at least one stem from each set of descriptors. This is just the same as Step 2 in the original EPU methodology.

Step 3: We apply natural language processing to identify the leading topics of all the articles. As a result, articles that have leading topics like 'sport' or 'fashion' that are clearly irrelevant for evaluating economic policy uncertainty are eliminated. The need to eliminate irrelevant articles arises because stems in the descriptors might have ambiguous meanings and so be used in different contexts. For instance, the sentence 'This car is not economic in fuel consumption, and the producer's patchy delivery policy makes its availability uncertain' contains all the words needed to identify it as coming under the economic policy uncertainty measure, but its topic is clearly far away from economics. Human validation would detect this, but it is not practically viable for the human eye to scrutinise hundreds of thousands of articles. We use one of the more popular topic modelling methods, the Latent Dirichlet Allocation, LDA; see, e.g. Blei et al. (2003) and Blei (2012). The LDA, in its essence, is a method of cluster analysis that identifies, by learning, (i) the latent topics, (ii) the sets of stems that define each latent topic and (iii) the probabilities that an article discusses a latent topic, for each article and each topic. Each article is assigned to a topic with the highest probability in the per-document distribution of topics estimated by the LDA model. We then give the latent topics meaningful names. Hence, the human input is substantially smaller than that required for human coding of the text corpus. For topic modelling, we use the `LDA(k=20)` and `LDA(k=30)` functions from the `topicmodels` library in R with its default parameters. For a description of this and a comparison with the alternative methods, see, for example, Grün and Hornik (2011). The Word2vec calculations were done in Python. The outcome is that we can identify articles which topics might be irrelevant for our purpose, like ‘sport’ or fashion’.

LDA and similar methods have been applied for analysing uncertainty before. Calvo-González et al. (2018) notably analysed policy volatility by applying LDA to identify the topics of presidential speeches in a number of countries, and Larsen (2021) used this method to construct a series of uncertainty measures for Norway. Tobback et al. (2018) used the support vector machine (SVM) technique, an alternative to LDA, to develop an uncertainty index for Belgium (for a comparison of LDA and SVM, see Luo and Li, 2014). Azqueta-Gavaldón et al. (2020) applied both Word2vec and LDA to evaluate economic policy uncertainty for the euro area.

Step 4: Instead of aggregating the selected articles linearly, we apply weights that represent the sentiments conveyed in each of the articles. This means we assume that the manifestation of sympathy, compassion and empathy might either emphasise or soften the expression of economic policy uncertainty conveyed in the article. We set the aggregation weights by conducting sentiment analysis, which calculates the number of stems with positive and negative feelings using the Russian sentiment lexicon. The general methodology that is used here is a slightly simplified version of the lexicon-based approach proposed by Taboada et al. (2011) that was originally proposed for English-language media. A sentiment lexicon is a set containing words or phrases in a given language that are evaluated as being associated with a sentiment. The choice and evaluation of the lexicon are always controversial (for a discussion, see Algaba et al., 2019). From the various sentiment lexicons available for the

Russian language (for a review and a methodological proposition for creating a unified lexicon see Kotelnikov et al., 2018), we have decided to use RuSentiLex by Loukachevitch and Levchik (2016), available at https://www.labinform.ru/pub/rusentilex/, which contains more than 12,000 words and expressions. We do not use intensifiers, so we do not consider the intensity of the positive and negative feelings. The proper application of intensifiers in Russian would require additional study. To conduct sentiment analysis, we convert the stemmed corpus into a `quanteda` document-feature matrix (DFM), convert the sentiment lexicon to an R dictionary and apply the `dfm` function from the `quanteda` library `dfm(data_dfm, dictionary=sentiment_dictionary)`. Next, we compute for each article the fraction of the stems with sentiment in the entire number of stems in the article and rescale these fractions on the interval from 0 to 1, defined by the 0.15, 0.5, 0.75, and 0.9 quantiles (for the methodology, see Thelwall et al., 2010). We denote these rescaled fractions of the positive sentiments as $SP(i)$ and negative sentiments as $SN(i)$, where i is the number of the article in the corpus. Finally, we compute the uncertainty weights, $W(i)$, using the balance of sentiments so that $W(i) = 1-(SP(i) - SN(i))$. The intuition here is that if negative sentiment prevails over positive sentiment, the uncertainty increases in proportion to the balance of sentiments, and similarly, if the balance of sentiments is positive, uncertainty decreases. If the numbers of positive and negative expressions are similar, the weight of each article is one. In the original EPU methodology, the index is computed on the basis of the frequency of appearance of the articles selected in Step 2. Instead, we sum up the sentiment weights of each article selected and divide this sum by the total number of articles published in each period. That is, both indices coincide if for all selected articles $W(i) = 1$.

The sentiment-weighted aggregation might also be done for English-language indices, but it seems to be particularly relevant for Russian. Gladkova (2010) gives examples of several possible misrepresentations of sympathy, compassion and empathy between English and Russian, all of which might be relevant when feelings related to uncertainty are expressed. Moreover, Jha et al. (2021) show that the level and dynamics of the sentiment in expressions related to finance are substantially different in Russian to what they are in other languages.

Our methodology described in the four steps above avoids the difficult task of proper human validation of the entire set of articles, and to some extent, accounts for the particular linguistic features of the Russian language. There is still a degree of arbitrariness here, though, as the sets of topics identified by the LDA still have to be named. As each topic contains a limited number of stems, with only 20 in our case, this is a relatively easy task.

3. Uncertainty index for Russia

We construct our index data from four newspapers available electronically, which are:

1. *Kommersant*	(Oct 1992 – Feb 2018),	579 997 articles
2. *Vedomosti*	(Dec 2003 – Feb 2018),	342 309 articles
3. *Novaya Gazeta*	(Feb 2004 – Feb 2018),	63 884 articles
4. *Moskovsky Komsomolets*	(Jan 2005 – Feb 2018),	143 758 articles

These newspapers represent a good spectrum of the newspapers available for different categories of readers. *Kommersant* is a newspaper that is primarily but loosely associated with information and news on business and commerce for a wide group of readers. According to the *Kommersant* website https://www.kommersant.ru/ about/kommersant, on 23 January 2020, its daily circulation is around 100,000—110,000 copies. *Moskovsky Komsomolets* is a popular newspaper addressed to a general audience with a print circulation of around

700,000 copies, according to https://ria.ru/20091211/198562973.html. *Vedomosti* is a business daily aimed at students and professionals. According to the Russian Wikipedia page https://ru.wikipedia.org/wiki/ведомости, its daily circulation is 75,000 copies. *Novaya Gazeta* is regarded as relatively independent and sometimes critical towards the Russian government. It is not a proper daily as it is published irregularly. Its reported circulation in August 2009 was 104,700 (see https://web.archive.org/web/20090822153334/http://www.pressaudit.ru/registry).

In Step 1 of our algorithm, we set the number of cosine-similar stems for each set of descriptors to 50. In Step 2, we set the number of LDA topics to 30 for all the newspapers except for *Kommersant*, where we reduce the number of topics to 20 because of the heavy computational burden. Articles with leading topics that we identify as 'sport', 'fashion', 'arts', 'crime', are eliminated. We extract the data by scraping using the `RSelenium` library, which is an adaptation of `Selenium WebDriver` in R (see Gojarea et al., 2015).

Table 1 gives a list of topics identified by the 20-topics LDA for *Kommersant* and named by us. Articles for which the leading topic is marked with an asterisk (*) have been disregarded. We understand the leading topic to be the one with the highest frequency of appearances in an article in comparison to other topics.

Table 1: List of topics with their arbitrary names for *Kommersant*, 20-topics LDA
Articles on topics that are marked with an asterisk (*) have been eliminated

1.	Commodities & markets	11.	(*)Traffic & accidents
2.	Policy & other matters	12.	(*)Culture
3.	(*)Culture	13.	(*)Europe & international (tourism & geography)
4.	Crime & economics	14.	War & international & US
5.	(*)Crime	15.	Air & trade & industry
6.	(*)Popular culture & former Soviet republics	16.	Markets
7.	(*)Sport	17.	Policy
8.	Former Soviet Union & war (history)	18.	Finance & banking
9.	Public finances and administration	19.	Mobile communication & industry
10.	Markets	20.	Commodities

The distributions of words in topics is illustrated in Figure 1, which presents word clouds for four selected topics, where the original stems are translated into English. This figure shows the most relevant or most frequent stems for each topic, and their size is proportional to the weights given to them by the Dirichlet distribution, which in turn represent the frequencies of their appearances. As we substitute English expressions for the original Russian language stems, the proportions might be slightly distorted. There might be an overlap between the topics, as some words might fit under more than one topic, depending on the context.

It may be noted that the size of the words in the topic cloud for 'culture' is smaller than that in the economics-oriented topics. This indicates that the words for culture are distributed more evenly, as the vocabulary used to describe cultural matters is richer than that for economic matters.

Figure 2 shows the heatmap of the frequencies of the monthly appearance of the articles with {uncertainty} descriptors, as defined by the 50-word Word2vec in *Kommersant*, split across the topics identified by LDA. For the clarity of the graph, we use 15-topic LDA rather than 20-topic LDA. We plot 12 of the 15 topics; otherwise, the plot would not be clearly readable. For each topic, the scale of colours represents the frequency of the appearance of articles where stems from {uncertainty} appear. The dark blue colour indicates that there were no articles in a given month on the topic that contained the {uncertainty} stems. The deep brown colour indicates that such stems are found in all the articles published that month. The figure shows that words from {uncertainty} did not frequently appear in the topics identified until 2011, with the exception of the topic ‘crime’.

Figure 1: Distributions of words (stems) in sample topics

Commodities & markets	**Finance & banking**
	
Culture	**Commodities**
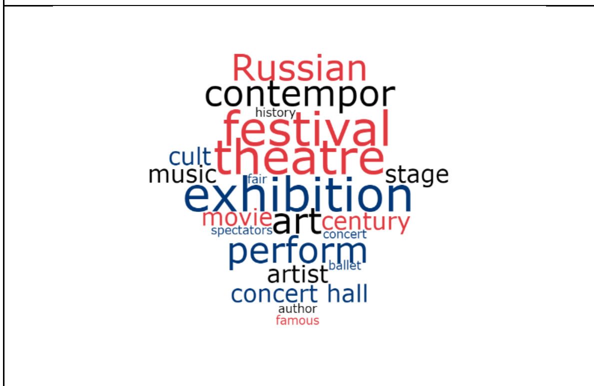	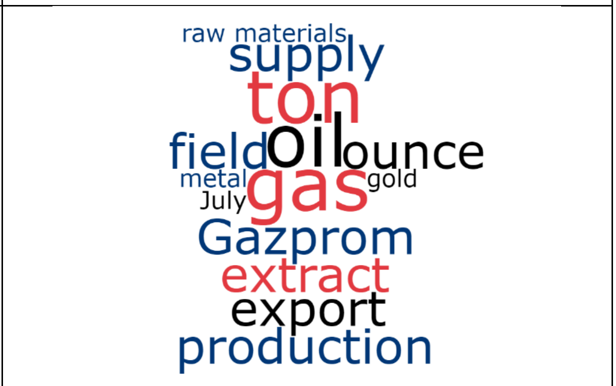

Figure 3 compares the time series of a single-paper version of the index, where only the data for *Kommersant* are used, with the fuller index based on all four newspapers. Both indices are weighted by the negative sentiments, that is, are computed for the monthly sums of negative sentiment weights of the articles selected and, after standardisation scaling, the longer index, that is, that for the *Kommersant* only, is normalised by its first observation and multiplied by 100. The differences between the means and variances of the two series, computed for data after *Vedomosti* was added to the index in December 2003 as the second journal, are strongly statistically significant. To test the difference in means, we use the t-test. To test the difference in variances, we apply the Pitman-Morgan-McCulloch test is robust to the non-normality of the data (see McCulloch, 1987). For the detailed results, see Appendix A. This shows that both the mean and the variance of the *Kommersant*-only series are greater than the corresponding characteristics of the fuller index. That is, heterogeneity in the style of newspapers affects the computed uncertainty measures markedly and confirms the rationale for using measures based on a wider selection of journals.

Figure 2: Monthly percentages of the appearance of {uncertainty} descriptors in *Kommersant,* Oct 1992 – Feb 2018

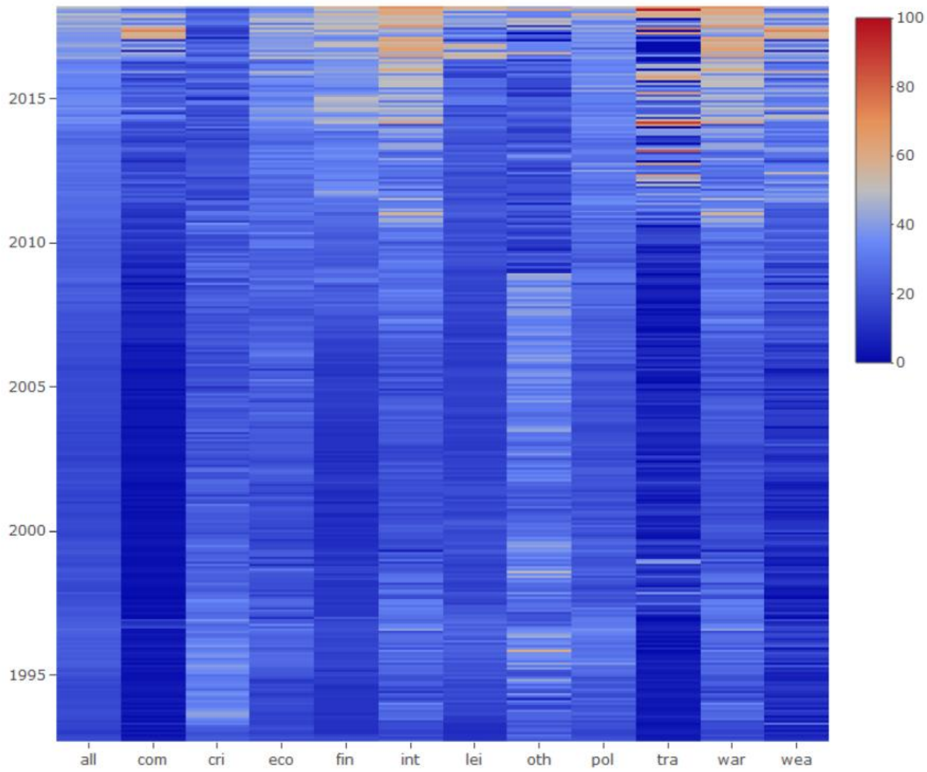

Legend: abbreviations for the topics names: `all`: all articles; `com`: commodities; `cri`: crime; `eco`: economics; `fin`: finance; `int`: international affairs; `lei`: leisure (mainly sport); `pol`: politics; `tra`: trade; `war`: war; `wea`: weather; `oth`: other.

Figure 3: Comparison of the uncertainty indices for Russia computed for *Kommersant* and all four newspapers[2]

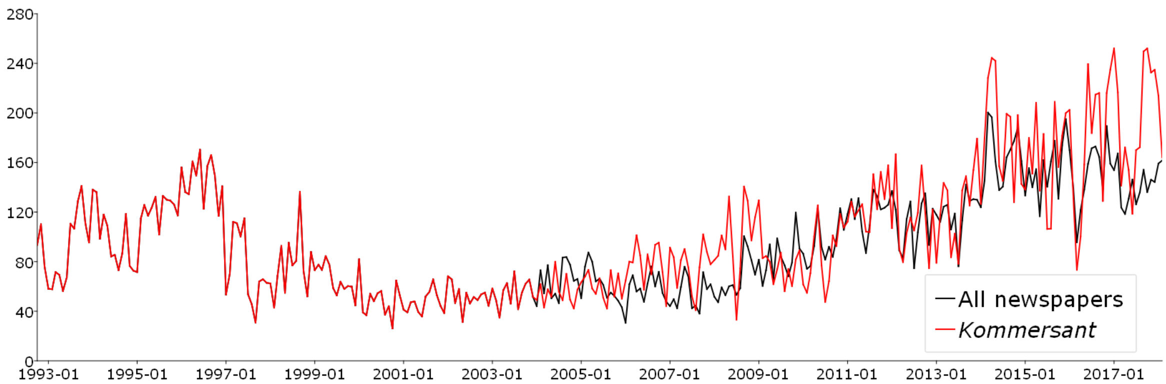

Legend: All newspapers: index computed using Steps 1-4 of the new methodology explained in Section 2 above with data from all 4 newspapers

Kommersant: as above, but with data from one newspaper (*Kommersant*) only

[2] Data are available upon request.

4. Tracing events that generate uncertainty

One of the simplest ways of evaluating the quality of uncertainty indices is to find out whether they reflect rises in uncertainty that are associated with events that may be expected to have generated it. These include unexpected incidents that attract large-scale publicity with unclear economic and political repercussions, such as terrorist attacks or threats, natural or human-made disasters, politically motivated arrests and killings, sudden economic policy changes like a devaluation, or international incidents. They also include pre-announced events that have an unclear outcome, like closely contested elections, meetings in which changes in policy might be announced, or major international sporting and cultural events.

We evaluate our index in comparison with the EPU index for Russia, which is available at https://www.policyuncertainty.com/russia_monthly.html and is based solely on the *Kommersant* newspaper. Both indices are plotted in Figure 4, with our index denoted by U. For the sake of comparison, we divide the indices by their standard deviations. After this scaling, the remaining main sources of difference between the indices are that (1) Word2vec cosine similarity principle is used to select descriptors for {economic}, {policy} and {uncertainty} in U, while they are set by human choice in EPU; (2) topic modelling is used in U and articles found to be irrelevant by word embedding are discarded, while human validation of articles is used for EPU; and (3) sentiment-related aggregation weights are applied rather than unity weights that are identical for all articles; (4) the data for EPU begin in January 1994, and those for U begin in October 1992. In Figure 4, we mark with grey bars the periods of military tensions that directly or indirectly involved Russia, which were the first Chechen war in December 1994 to August 1996 and the most intensive period of the second Chechen war in August 1999 to May 2000, the Georgian military conflict in August 2008, and the period of the most intensive conflict in the Donbas and Crimea areas in February 2014 to July 2015. We note that the first Chechen war overlaps to a large extent with the period of turmoil related to Russian privatisation. The Georgian conflict was followed very closely by the beginning of a period of strong depreciation of the Russian currency, which also blurs the picture. We also mark the dates of presidential elections, marking the transfer of power from Yeltsin to Putin as 30 December 1999 rather than the official election in March 2000.

Comparing the two series shows that the periods of temporary increases in uncertainty shown by the two indices usually coincide. Both series indicate an increase in the general level of uncertainty during times of military conflict. It is interesting that statistically significant differences, where the standardised differences between the corresponding values of the two indices are in the 5% tail of the standard normal distribution, mostly occur in January. The uncertainty shown in the EPU is significantly greater than that shown in U in the Januarys of 1996, 1997, 2015, and 2017, but in two other cases, it is the other way around, as the values of U for January 2016 and January 2018 are significantly greater than those of the EPU. It is likely that these anomalies happen because of the different content and vocabulary of articles in the Russian press after the New Year. This is often a rather uneventful month of the year, with a lot of days off work and school holidays for the New Year celebrations under both the Julian and Gregorian calendars. In January, the press often publishes summaries of the previous year's events and forecasts for the current year, which have little relation to the current situation. Hence the picture of uncertainty measured by text-based indices for January is likely to be distorted.

For the period from January 1994, which is the earliest time covered by both indices, to February 2018, we identify 49 uncertainty-generating events. The list of them is given in Appendix B. We then compute the peaks of the indices, defining a peak as an increase over the previous month's value of more than 7.5%. Using this definition, we match 46 peaks in the U series and 45 peaks in the EPU series with the pre-defined uncertainty-generating events (see the last column in the table in Appendix B). This gives a rather impressive 92% accuracy for the match of peaks with events for U and 90% for EPU.

Figure 4: Comparison of the U and EPU uncertainty indices
(divided by their standard deviations)

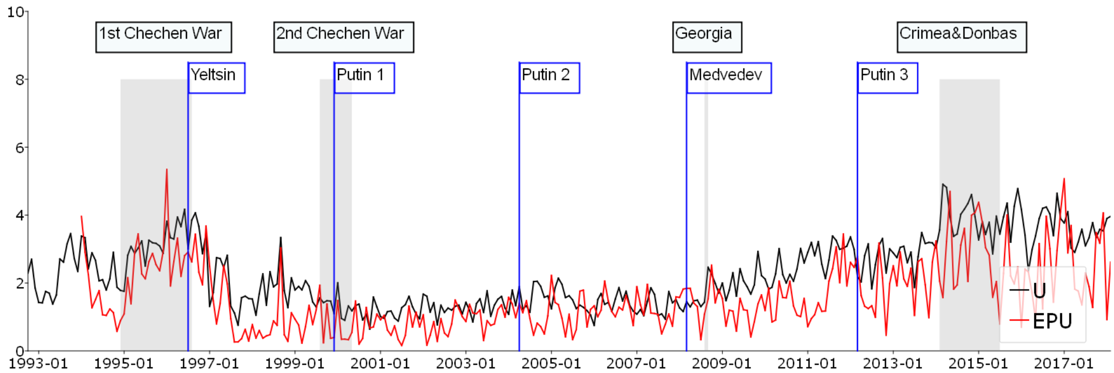

Source of data for EPU: https://www.policyuncertainty.com/russia_monthly.html

Realising that our selection of events and the definition of peaks might to some extent be arbitrary, we repeat the matching using different selections of events and different thresholds for the peaks. We also apply scaling to the series by moving the standard deviation to allow for long-run changes in uncertainty and substitute the analysis of peaks with analysis of substantive increases over the long-run tendency in the form of the 'steps' analysis. The results are similar to those presented here as the accuracy of both indices in matching the uncertainty-generating events is always above 80%, with both indices performing with similar accuracy.

5. Real effects of uncertainty

For further evaluation and comparison, we apply both indices to analyse the effects of uncertainty on Russia's industrial production. That the uncertainty shocks have a negative effect on growth has been documented by Bloom et al. (2018), and this long-run effect is usually attributed to investors taking a wait-and-see stance (the real options effect; see Bloom, 2009, 2014, and the references in Section 1 of this paper). There is also some evidence to the contrary, as technology news is often correlated with uncertainty, and so an increase in uncertainty might, in some cases, cause a positive real effect (Segal et al., 2015, Cascaldi-Garcia and Galvao, 2020).

However, the picture for studies of individual countries is often more blurred. Most papers analysing countries with open economies using various methodologies have found some confirmation of the real options effect, like for example for Chile (Cerda et al., 2016), China (Huang and Luk, 2018; Gu et al., 2021), Japan (Arbatli et al., 2019) and Switzerland (Dibiasi et al., 2018). The story is different though when a distinction is made between global and idiosyncratic effects. Ozturk and Sheng (2018) show that common uncertainty shocks do indeed provoke an adverse response in real economic activity. In the case of Russia, there may also be a short-run 'rush to complete' effect. In this case, the real sphere may have an initial positive reaction to an uncertainty shock that might reflect the specific nature of the Russian labour market. The Russian labour market has both stable employment and high wage flexibility at the same time (see, e.g. Gimpelson, 2019). The chance to create instant financial motivation for workers together with the flexible interrelation between the official and informal labour markets (see, e.g. Kapelyushnikov et al., 2012) might stimulate this effect.

We assess the possible real effects by computing and testing the impulse responses from the (local) linear projections (LP-IRs) from the uncertainty indices to industrial production in Russia (for the methodology, see Jordà, 2005, and Jordà and Marcellino, 2010). This requires us to orthogonalise least-squares projections. To do this, we first estimate the vector autoregressive (VAR) model of the Russian economy and compute the Cholesky decomposition of the covariance matrix of its residuals. We next apply the least-squares projections of the uncertainty into industrial production and orthogonalise using the decomposition matrix from VAR. The local projection is an attractive alternative to the widely used orthogonal impulse responses of the linear VAR model. Its simplicity and good asymptotic properties allow for testing using well-developed traditional statistical methods. On the practical side, the linear projections usually give impulse responses that might be interpreted more sensibly than those obtained directly from VAR. The usual lag selection procedures like the Akaike Information Criterion or similar often suggest VAR's with short lags as optimal, which in turn leads to trivial or not fully interpretable impulse responses. On the other hand, the linear projection produces more complex and interpretable impulse responses even if the VAR lags are short.

Our VAR model describing the Russian economy consists of five variables:

1. The measure of uncertainty being tested, either EPU or U.
2. The Brent oil price, deflated by the US retail price index.
3. The interbank annual interest rate deflated by annual inflation measured monthly.
4. The MOEX index of stock market prices on the Moscow Exchange in logs, deflated by the consumer price index for Russia.
5. The production output gap, measured by the log deviations from the Hodrick-Prescott cycle-trend.

Data for all the variables except the uncertainty indices and MOEX are available at https://fred.stlouisfed.org/. The MOEX data can be retrieved from various sources such as https://www.moex.com/. We obtain consistent and comparable data for all the variables from October 1997 to the end of our uncertainty data in February 2018.

The construction of the model and the selection of variables is typical for a VAR model of a natural resource economy (for Russia see, e.g. Ito, 2008; Perifanis and Dagoumas, 2017; Kholodilin and Netšunajev, 2019; and, in a more general form, Oloko et al., 2021). The variables enter the model after being transformed into stationary data. They are tested for stationarity using several optimal point unit root tests that allow for up to three structural breaks (see Carrion-i-Silvestre et al., 2009). All the variables except for oil prices are trend-stationary so that a linear trend is removed. Oil prices are stationary in the first differences of the trend-removed data. The variables are ordered from the uncertainty variable to industrial production, and changing the ordering does not affect the results. The order for the VAR is decided by the corrected Akaike Information Criterion. The model is estimated using the maximum likelihood method.

Figure 5 gives the estimates of the LP-IRs of industrial production to a one-standard-deviation uncertainty shock, where the uncertainty variable is defined either by EPU or by our uncertainty index U. The inside bands give the 90% individual region and the outside, wider, bands are the simultaneous Bonferroni 90% bands.

The pictures of real effects shown by the two uncertainty indices are substantially different. The impulse responses of industrial production to uncertainty shocks measured by EPU show some signs of the positive rush to complete effect, with a culmination about five months after the shock. There is also an indication of a more delayed negative response occurring 11-13 months after the shock. Our index U indicates a negative response 1-2 months after the shock and some more delayed positive and negative responses after about one year. This means the results obtained with U confirm the existence of the real options effect, which is in line with the global findings of Bloom et al. (2018).

An obvious question that arises is which of the two indices is to be believed? The statistical support for the effect of U is stronger than that for EPU. When our index measures uncertainty, the negative real effect is clearly statistically significant, as both individual and simultaneous confidence bands are lower than zero for time horizons 1 and 2. They also show a reasonably strong further negative response for about a year.

Figure 5: Linear projection impulse responses of industrial production in Russia to a one standard deviation uncertainty shock

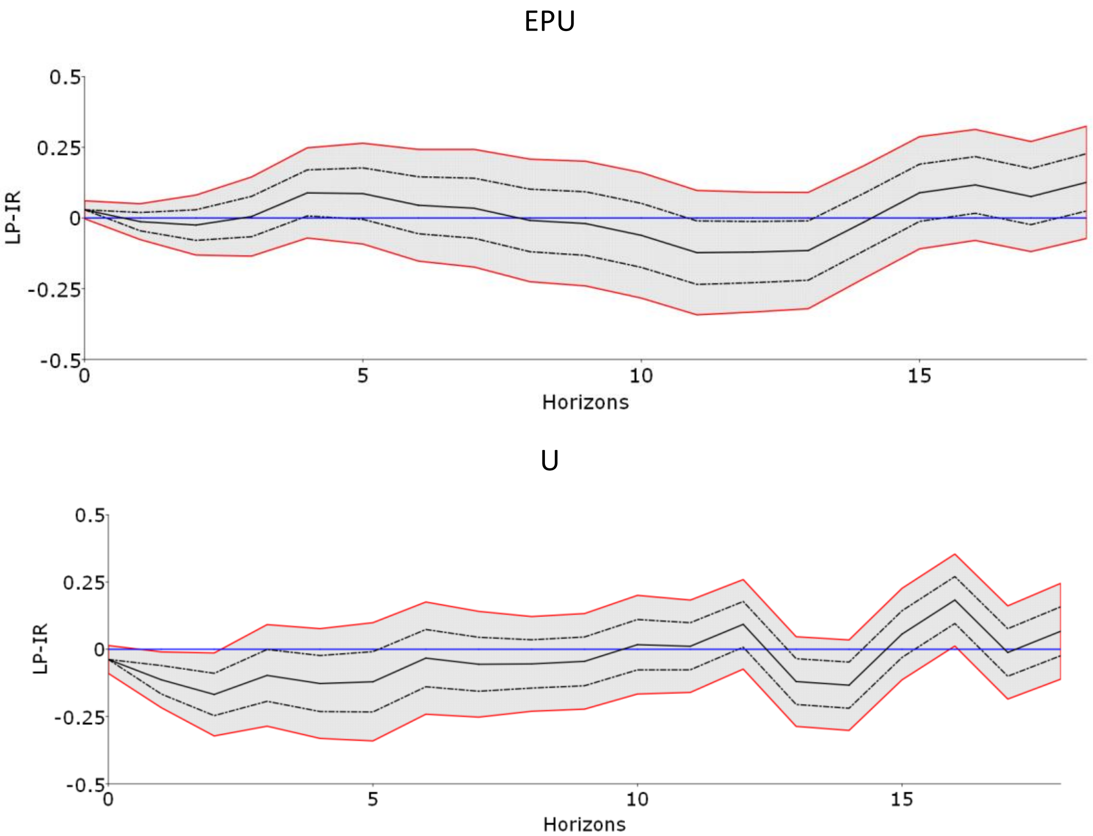

Figure 6 sheds some light on the role the particular components of the U index play in discovering this real effect. It shows the local projection impulse responses (upper panel) of industrial production to a one standard deviation increase in uncertainty and the upper bands of their 90% confidence intervals (lower panel) when uncertainty is measured alternatively by five indices: the original EPU index (cyan line); our reconstructed EPU indices obtained by using all four journals (red line) and *Kommersant* only (blue line) without any human validation in either case; and two variants of the U index computed with topic modelling (LDA) and word embedding (Word2vec), one marked in green with the sentiment weighting and one without in black. For the sake of clarity, we show the impulse responses for the first eight forecast horizons only and smoothed by the Bezier smoother. This means the black line in the upper panel of Figure 6 is the first eight horizons of the Bezier-smoothed impulse response of U in the lower panel of Figure 5. In the lower panel of Figure 6, the black line corresponds to the upper dotted line in the lower panel of Figure 5. Similarly, the light green lines in Figure 6 representing the impulse response of EPU correspond to the impulse response from EPU (the black line in the upper panel of Figure 5) and its 90% upper band (the upper dotted line in the upper panel of Figure 5).

Figure 6: Comparison of the linear projection impulse responses from different uncertainty indices and their upper 90% bands

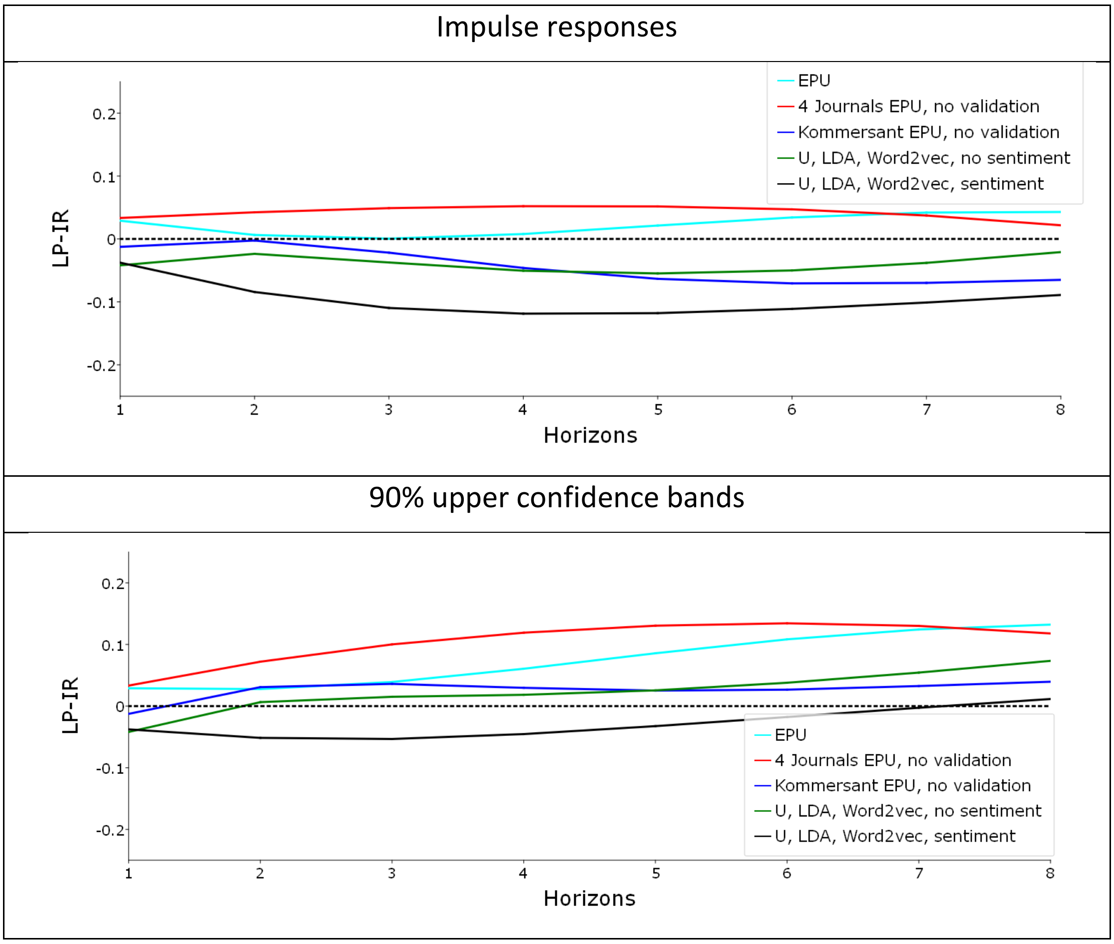

If the upper confidence band is below the zero line, it indicates that the real options effect is significant. The results show that the worst index at being able to confirm this effect is our replication of the EPU index based on data from four newspapers without human validation, topic modelling, word embedding or sentiment weighting. Comparing this with our one-newspaper reconstruction of EPU indicates that adding newspapers actually weakens the significance of this effect. This might arise because the volatility of the index decreases with the increase in the number of newspapers, as shown by Figure 3. By construction of the impulse response, lower volatility implies a smaller absolute magnitude for the one-standard-deviation shock, which might dampen the impulse response function. It might also reflect changes in the structure through the context and vocabulary of newspapers added at later dates to the index that was originally built on the *Kommersant* data alone.

The comparison of impulse responses from the original EPU index based only on the *Kommersant* data and the reconstruction of it without human validation does not show much difference for forecast horizons 2 and 3. For further horizons, the differences become more visible, though both impulse responses remain insignificant. This shows, in our opinion, that human validation plays a somewhat limited role here.

The U index obtained without applying sentiment weighting but based on the LDA-selected articles and using descriptors obtained by word embedding is slightly better at confirming the real options effect than the original EPU is. For the forecast horizons from 1 to 8, its impulse responses are consistently below zero, though they are not significant, while the impulse responses for EPU are positive or very close to zero. We actually expected these two results to be close to each other, but there are two possible reasons for the difference. Firstly, substituting human validation of a sample of articles by topic identification for all the articles allows a substantial part of nuisance articles to be eliminated, giving a clearer picture of the real options effect. Secondly, the word embedding process gives, as a result, richer and, more specifically, country-oriented sets of descriptors, particularly for {economic} which, in combination with {uncertainty}, reflects the Russian economic policy uncertainty more accurately.

The only uncertainty index that convincingly and significantly confirms the real options effect is our index constructed by applying topic modelling and word embedding, but only if it is weighted with the sentiment weights. If these weights are not applied, the impulse responses from such an index are insignificant except for at forecast horizon 1.

6. Robustness checks and discussion[3]

To evaluate how much our results depend on the methodological assumptions we apply, we run several robustness checks, changing the methodology used for computing the indices and testing. The changes we made are:

(1) We experiment with larger and smaller VAR models, including the interest rate as an extra variable in the model and dropping variables with less explanatory power. We also use different deflators for the Brent oil price variable. These do not change the results markedly.

(2) In addition to the local projection impulse responses, we compute the naïve orthogonal impulse responses for a VAR model and also use the 'bootstrap after bootstrap' approach, which gives responses with a bootstrap-mean correction (see Kilian, 1998; Kilian and Kim, 2011) and bootstrapped bands. Here the results were sometimes slightly different,

[3] The detailed results of the robustness analysis are available upon request.

suggesting the existence of the positive rush to complete effect alongside the negative real options effect.

(3) We use a different Russian language lexicon, available from the Kaggle website (https://www.kaggle.com/rtatman/sentiment-lexicons-for-81-languages/metadata). It contains only about 3000 words for positive and negative sentiments. Again, the differences from the results presented here are negligible. As documented by the descriptive statistics in Appendix A, moments of data obtained by using different sentiment lexicons are similar for series in levels (Table A1) and first differences (Table A2), and the correlation is very high and significant. However, the tests for comparing variances in dependent samples (the Pitman-Morgan and Pitman-Morgan McCulloch tests) suggest that there might be some differences in the variance of the series if their dependence is accounted for.

(4) We experiment with different formulae for the sentiment weights. We (i) use the fractions of positive and negative words in the articles, (ii) truncate and rescale the ratios in order to increase the variability of the indices, and (iii) use quantiles of the exponential distribution as weights. Again, the descriptive statistics in Appendix A, where the comparison between the indices scaled by the balance of sentiments and negative sentiment are shown, indicate a very high correlation between them and no significant difference between their means and variances. Such correlation holds for both levels (Table A1) and first differences (Table A2) of the indices.

(5) The sentiment scaling described here is based on the assumption of homogeneity of journalistic styles in expressing sentiments in newspapers. That is, we scale the sentiments using the 0.15, 0.5, 0.75, and 0.9 quantiles division jointly for the data from all the newspapers. We also compute and apply the sentiment scaling under the assumption of heterogeneity of styles, as we compute the quantile divisions separately for each newspaper. That is, we adjust the index to reflect that some newspapers might use sentiment-related words more frequently than others do. The resulting difference between the heterogeneous and homogeneous indices is minimal, as the tests for means and variances differences (the t-test and Pitman-Morgan-McCulloch test) are insignificant. (see Appendix A).

(6) We replicate the results using a different number of stems selected as descriptors by Word2vec and a different number of LDA topics. Increasing the number of stems does not affect the results. Increasing the number of topics, however, blurred the results and made them difficult to interpret.

(7) To evaluate the possible effect of overstemming, that is, reducing words of different meanings to the same stem, we prepare sets of descriptors for {economic}, {policy} and {uncertainty} using, in place of stems, full words in all declinations. We also create the ‘bags of words’ using full words rather than their stems. In such a case, natural language processing is not practically possible. Consequently, we compute two rudimentary indices for the stemmed and full words cases without topic modelling and human validation; that is, for all articles in the corpus. The results in Appendix A, Tables A1 and A2, show that possible overstemming might indeed affect the results slightly, as the variance difference tests indicate some significance. Nevertheless, the correlation between the indices is very high; which indicates that the distortion might not be substantial.

We can therefore conclude that our results are reasonably robust. They are consistent for different settings in machine learning techniques, for matching events with uncertainty, and for the econometric techniques applied to evaluate the real effects.

We realise that the uncertainty measure we propose here can still be improved. This is particularly so for the sentiment measures, which are based on assigning sentiments to individual words without analysing the context. If more accurate methods were to be applied, further progress would be likely. Such methods could mean using more accurate word embedding techniques or context-dependent methods of word embedding, which look at the context of sentences instead of the similarity of words, and, above all, more sophisticated methods of sentiment analysis, which is usually based on a sentence embedding approach. However, even the relatively simple methods we apply here appear to have done their job well.

7. Conclusions

The method we propose here for constructing the economic policy uncertainty index overcomes some problems that language diversity creates here. We apply our approach for Russia, but its straightforward methodology and recent advances in the natural language processing methods suggest it should be relatively easy to apply it for other countries and different languages. In the simplest approach, the sets of cosine similar words describing economic policy uncertainty can be created for a selected language and applied in search of newspapers published in this language, where the relevance of the articles is defined by examining their leading topics. These leading topics could be identified by Latent Dirichlet Allocation or a similar method. It would also be relatively straightforward to separate the idiosyncratic or country-specific uncertainty from global uncertainty by including terms for global uncertainty in the search. This is left for further studies. As sentiment lexicons become widely available for hundreds of languages and computational costs diminish, applying the natural language-based method of analysing uncertainty might lead to uncertainty measures being produced that can track changes in idiosyncratic and global uncertainty quickly and regularly. Also, in future, it might be of advantage to apply lemmatisation rather than stemming. In the case of languages with complex declinations, like Russian, lemmatisation might improve the quality of the constructed indices. This, however, has to wait until sufficient progress in lemmatisation of various national languages is reached.

This research also reveals that translating English-language uncertainty-related descriptors into a different language might lead to a loss of clarity and may consequently blur the uncertainty measures because the sentiments and undertones of the two languages are different. The index we construct, which accounts for sentiment, works well at measuring the level of uncertainty in Russia. Its general tendency confirms the existence of two phases of heightened uncertainty in Russia before the pandemic, with one that arose during the first Chechen war, which largely coincided with the turmoil related to privatisation, and another during the Crimea and Donbas crises. It allows the uncertainty generated by particular events to be identified and measured.

In a statistically significant way, our results show that the predominant impact on Russia's industrial production of an increase in economic policy uncertainty is negative, as it points to a decrease in industrial production for at least two months after the uncertainty rises. This is surely a mix of positive and negative effects, but the negative effect prevails on average. It is possible to account for it using data from 1998-2018, despite changes that happened in the editorial strategies in the press in this period and some self-censorship and direct and indirect pressure imposed upon some journals and journalists.

References

Adams, S., F. Adedoyin, E. Olaniran, and F. V. Bekun (2020), 'Energy consumption, economic policy uncertainty and carbon emissions; causality evidence from resource rich economies', *Economic Analysis and Policy* **68**, 179–190.

Atsu, F. and S. Adams (2021), 'Energy consumption, finance, and climate change: Does policy uncertainty matter?', *Economic Analysis and Policy* **70**, 490–501.

Algaba, A., Ardia, D., Bluteau, K., S. Borms and K. Boudt (2019), 'Econometrics meets sentiment: An overview of methodology and applications", presented at 13th International Conference on Computational and Financial Econometrics (CFE 2019), London.

Al-Thaqeb, S.A. and B.G. Algharabali (2019), 'Economic policy uncertainty: a literature review, *Journal of Economic Asymmetries* **20**, e00133.

Altig, D., S. Baker, J.M. Barrero, N. Bloom, P. Bunn, S. Chen, S.J. Davis, J. Leather, B. Meyer, E. Mihaylov, P. Mizen, N. Parker, T. Renault, P. Smietanka and G. Thwaites (2020), 'Economic uncertainty before and during the Covid-19 pandemic', *Journal of Public Economics* **191** 104274.

Arbatli, E.C., S. J. Davis, A. Ito, and N. Miake (2019), 'Policy Uncertainty in Japan', Working Paper 23411, National Bureau of Economic Research.

Aromí, J.D., (2017), 'Measuring uncertainty through word vector representations', *Económica,* La Plata **63**, 135-156.

Bahmani-Oskooee, M. and H. Ghodsi (2017), 'Policy uncertainty and house prices in the United States of America", *Journal of Real Estate Portfolio Management* **23**, 73-85.

Bahmani-Oskooee, M., A. Kutan, and A. Kones (2016), 'Policy uncertainty and the demand for money in the United States', *Applied Economics Quarterly* 62, 37-49.

Bahmani-Oskooee, M. and M. Maki Nayeri (2021), 'Policy uncertainty and the demand for money in the United Kingdom: Are the effects asymmetric?, *Economic Analysis and Policy* **66**, 76–84.

Baker, S.R., N. Bloom and S.J. Davis (2016), 'Measuring economic policy uncertainty', *Quarterly Journal of Economics* **131**, 1593–1636.

Benoit, K., K. Watanabe, H. Wang, P. Nulty, A. Obeng, S. Müller and A. Matsuo (2018), 'quanteda: an R package for the quantitative analysis of textual data', *Journal of Open Source Software* **3**, 774, 1-4.

Blei, D. (2012), 'Probabilistic topic models', *Communications of the ACM* **55**, 77-84.

Blei, D.M., A.Y. Ng and M.I. Jordan (2003), 'Latent Dirichlet allocation', *Journal of Machine Learning Research* **3**, 993-1022.

Bloom, N. (2009), 'The impact of uncertainty shocks', *Econometrica* **77**, 623-685.

Bloom, N. (2014), 'Fluctuations in uncertainty', *Journal of Economic Perspectives* **28**, 153-175.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten and S.J. Terry (2018), 'Really uncertain business cycles', *Econometrica,* **86***,* 1031–1065.

Calvo-González, O., A. Eizmendi and G. Reyes (2018), 'Winners never quit, quitters never grow. Using text mining to measure policy volatility and its link with long-term growth in Latin America', The World Bank Group, Policy Research Working Paper 8310.

Carrion-i-Silvestre, J.L, Kim D. and Perron, P. (2009), 'GLS-based unit root tests with multiple structural breaks both under the null and the alternative hypotheses', *Econometric Theory* **25**, 1754-1792.

Cascaldi-Garcia, D., and Galvao, A. B. (2020), 'News and uncertainty shocks', *Journal of Money, Credit and Banking*, in print

Cerda, R., Silva, A. and T. Valente (2016), 'Economic policy uncertainty indices for Chile', Latin American Center for Economic and Social Policies (CLAPES-UC) Chile.

Charemza, W. S. Makarova and K. Rybiński (2020), 'Economic uncertainty and anti-pandemic government policy: the cross-country perspective', 14th International Conference *Computational and Financial Econometrics*, London.

Chiswick, BR and P.W. Miller (2005), 'Linguistic distance: a quantitative measure of the distance between English and other languages', *Journal of Multilingual and Multicultural Development* **26**, 1-11.

Davis, S.J., D. Liu and X, S. Sheng (2019), 'Economic policy uncertainty in China since 1949: The View from Mainland Newspapers', Fourth Annual IMF-Atlanta Fed Research Workshop on China's Economy, Atlanta.

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

Dibiasi, A., K. Abbergera, M. Siegenthaler and J-E. Sturma (2018), 'The effects of policy uncertainty on investment: Evidence from the unexpected acceptance of a far-reaching referendum in Switzerland', *European Economic Review*, **104**, 38-67.

Esser, F. and Umbricht, A. (2013), 'Competing models of journalism? Political affairs coverage in US, British, German, Swiss, French and Italian newspapers', *Journalism* **14,** 989–100.

Ghirelli, C., J. J. Pérez, and A. Urtasun (2019). 'A new economic policy uncertainty index for Spain', *Economics Letters* **182** 64-67.

Ghirelli, C., M. Gil, J. J. Pérez and A. Urtasun (2021), 'Measuring economic and economic policy uncertainty and theirmacroeconomic effects: the case of Spain', *Empirical Economics* **60**, 869–892

Gimpelson, V. (2019), 'The labor market in Russia, 2000–2017', IZA World of Labor 466, https://wol.iza.org/uploads/articles/501/pdfs/the-labor-market-in-russia.pdf?v=1.

Gladkova, A. (2010), 'Sympathy, compassion, and empathy in English and Russian: a linguistic and cultural analysis', *Culture and Psychology* **16**, 267-285.

Gojarea, S., R. Joshib and D. Gaigaware (2015), 'Analysis and design of Selenium WebDriver automation testing framework', 2nd International Symposium on Big Data and Cloud Computing (ISBCC'15), *Procedia Computer Science* **50**, 341 – 346.

Grün, B. and K. Hornik (2011), 'topicmodels: an R package for fitting topic models', *Journal of Statistical Software* **40**, 1-30.

Gu, X., X. Cheng, Z. Zhu and X. Deng (2021), 'Economic policy uncertainty and China's growth-at-risk', *Economic Analysis and Policy* **70**, 452–467.

Hossain, A. A. and P. Arwatchanakarn (2020), 'The effect of economic uncertainty on narrow money demand and its stability in New Zealand: An empirical investigation', *Economic Analysis and Policy* **68**, 88–100.

Hou, X., T. Wang and C. Ma (2021), 'Economic policy uncertainty and corporate fraud', *Economic Analysis and Policy* **71**, 97–110.

Huang, Y. and P. Luk, (2018), 'Measuring economic policy uncertainty in China', Hong Kong Baptist University.

Isphording, I.E. and S. Otten (2013), 'The costs of Babylon—linguistic distance in applied economics', *Review of International Economics* **21**, 354–369.

Jha, M., H. Liu and A. Manela (2021), 'Does finance benefit society? A language embedding approach', presented at 2nd Workshop 'Big Data and Economic Forecasting', European Commission.

Ito, K. (2008), 'Oil prices and macro-economy in Russia: The co-integrated VAR model approach', *International Applied Economics and Management Letters* **1**, 37-40.

Jirasavetakul, L-B. F., and A. Spilimbergo (2018), 'Economic policy uncertainty in Turkey', IMF Working Paper WP/18/272.

Jordà, Ò. (2005), 'Impulse responses by local projections', *American Economic Review* **95**, 161-182.

Jordà, Ò and M. Marcellino, (2010), 'Path forecast evaluation', *Journal of Applied Econometrics* **25**, 635–662.

Kapelyushnikov R., A. Kuznetsov and O. Kuznetsova (2012), 'The role of the informal sector, flexible working time and pay in the Russian labour market model', *Post-Communist Economies* **24**, 177-190.

Kholodilin, A. and Netšunajev (2019), 'Crimea and punishment: the impact of sanctions on Russian economy and economies of the euro area', *Baltic Journal of Economics* **19**, 39–51.

Kilian, L. (1998) 'Small-sample confidence intervals for impulse response functions', *Review of Economics and Statistics* **80**, 218-230.

Kilian, L. and Y.J. Kim (2011), 'How reliable are local projection estimators of impulse responses?', *Review of Economics and Statistics* **93**, 1460–1466.

Kılıç, İ. (2021), 'Geopolitical uncertainty and its macroeconomic effects: the case of Turkey, 2006-2016', Vistula University, PhD thesis.

Kotelnikov, E., T. Peskisheva, A. Kotelnikova and E. Razova (2018), 'A comparative study of publicly available Russian sentiment lexicons', in: (D. Ustalov, A. Filchenkov, L. Pivovarova and J. Žižka, eds) *Artificial Intelligence and Natural* Language *AINL 2018 Communications in Computer and Information Science* **930** Springer, 139-151.

Larsen, V. H. (2021), 'Components of uncertainty', *International Economic Review*, in press.

Liu, L. and T. Zhang, (2015), 'Economic policy uncertainty and stock market volatility', *Finance Research Letters* **15**, 99-105.

Loukachevitch N., and A. Levchik (2016), 'Creating a general Russian sentiment lexicon', *Proceedings of Language Resources and Evaluation Conference LREC-2016*, 1171-1176.

Luo, L. and L. Li. (2014), 'Defining and evaluating classification algorithm for high-dimensional data based on latent topics', PLoS ONE **9**, e82119., doi:10.1371/journal.pone.0082119

Manning, C.D., P. Raghavan and H. Schütze (2008), *Introduction to information retrieval*, Cambridge UP.

McCulloch, C.E. (1987), 'Tests for equality of variance for paired data', *Communications in Statistics - Theory and Methods* **16**, 1377–1391.

Mikolov, T., K. Chen, G. Corrado and J. Dean (2013), 'Efficient estimation of word representation in vector space', *arXiv*: 1301.3781v3, 1-12.

Naili, M., A.H. Chaibi and H.H. Ben Ghezala (2017), 'Comparative study of word embedding methods in topic segmentation', International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES2017, Marseille.

Oloko, T.F., A.E. Ogbonna, A.A. Adedeji and N. Lakhani (2021), 'Oil price shocks and inflation rate persistence: A Fractional cointegration VAR approach', *Economic Analysis and Policy* **70**, 259–275.

Ozturk, E.O. and Sheng, X.S. (2018), 'Measuring global and country-specific uncertainty', *Journal of International Money and Finance* **88**, 276-295.

Peng, G., Huiming, Z. and Y. Wanhai, (2018), 'Asymmetric dependence between economic policy uncertainty and stock market returns in G7 and BRIC: A quantile regression approach', *Finance Research Letters* **25**, 251-258.

Perifanis, T. and A. Dagoumas, (2017), 'An econometric model for the oil dependence of the Russian Economy', *International Journal of Energy Economics and Policy* **7**, 7-13.

Petersson, B. and E. Persson (2011), 'Coveted, detested and unattainable? Images of the US superpower role and self-images of Russia in Russian print media discourse', *International Journal of Cultural Studies* **14**, 71–89.

Segal, G., Shaliastovich, I., and Yaron, A. (2015), 'Good and bad uncertainty: macroeconomic and financial market implications', *Journal of Financial Economics* **117**, 369-397.

Sidorov, G., A. Gelbukh, H. Gomez-Adorno and D. Pinto (2014), 'Soft similarity and soft cosine measure: similarity of features in vector space model', *Computación y Sistemas* **18**, 491-504.

Schimpfössl, E. and I. Yablokov (2020), 'Post-socialist self-censorship: Russia, Hungary and Latvia', *European Journal of Communication* **35**, 29–45.

Taboada, M. J. Brooke, M. Tofiloski, K. Voll and M. Stede (2011), 'Lexicon-based methods for sentiment analysis', *Computational Linguistics* **37**: 272–274.

Thelwall, M., K. Buckley, G. Paltoglou, D. Cai and A. Kappas (2010), 'Sentiment strength detection in short informal text', *Journal of the American Society for Information science and Technology* **61**, 2544-2558.

Tobback, E., H. Naudts, W. Daeleman, E. J. de Fortuny and D. Martens (2018), 'Belgian economic policy uncertainty index: Improvement through text mining', *International Journal of Forecasting* 34, 355-365.

Vartanova, E. (2012), 'The Russian media model in the context of post-Soviet dynamics', in (D.C. Hallin and P. Mancini, eds.), *Comparing media systems beyond the Western word*, Cambridge UP, chapter 7, 119-142.

Wessler, H., M. Skorek, K. Kleinen-von Königslöw, M. Held, M. Dobreva and M. Adolphsen (2008), 'Comparing media systems and media content: online newspapers in ten Eastern and Western European countries', *Journal of Global Mass Communication* **1**, 165-189.

Yang, L. and S. Hamori (2021), 'Systemic risk and economic policy uncertainty: International evidence from the crude oil market', *Economic Analysis and Policy* **69**, 142–158.

Appendix A

Pairwise comparison of selected indices

Tables in this Appendix contain sets of descriptive statistics comparing pairs of uncertainty indices. The pairs compared are

Pairs of uncertainty indices	symbols of columns
EPU and U computed for all newspapers	EPU ; U
EPU and U computed for *Kommersant* only	EPU ; U(Kom.)
U and U computed for *Kommersant* only	U ; U(Kom.)
U under homogeneity and heterogeneity of journalistic style	U(Hom.) ; U(Het.)
U computed with the use of Loukashevich and Kaggle lexicons	U(Louk.) ; U(Kag.)
U and U weighed by negative sentiments only	U ; U-
U based on stemmed and not stemmed bags of words	U(st.) ; U(nst.)

In the comparison, the length of the longer index has been truncated so that the pairs of indices compared are of equal length. The indices are compared in their levels (Table A1) and first differences (Table A2).

The upper panel of each table contains in rows 2-10 the univariate characteristics of the compared series in the adjacent columns, that is:

row	description (upper panel)	symbol
2	number of observations	nobs.
3	arithmetic means	means
4	standard deviations	st.devs.
5	skewness coefficients	skew.
6	p-value of the skewness coefficient	pval.
7	kurtosis coefficient	kurt.
8	p-value of the kurtosis coefficient	pval.
9	Jarque-Bera normality statistic	J-B
	p-value for the Jarque-Bera statistic	pval.

The lower panel of each table contains various correlation and homogeneity statistics, with their values on column 'Stat' and p-values are given alongside in column 'pval'.

row	description (lower panel)	symbol
2	Pearson correlation coefficient and its bootstrapped p-value	Pearson
3	Spearman correlation coefficient and its bootstrapped p-value	Spearman
4	Kendall correlation coefficient and its bootstrapped p-value	Kendall
5	Li-Li-Tsai quantile correlation coefficient with the first variable truncated at the 0.75 quantile and its bootstrapped p-value	qcorr1
6	Li-Li-Tsai quantile correlation coefficient with the second variable truncated at the 0.75 quantile and its bootstrapped p-value	qcorr2
7	Student t-statistic for testing differences in means and its asymptotic p-value	t-diff.
8	F variance ratio statistic and its asymptotic p-value	F
9	Pitman-Morgan statistic for comparing dependent variances under normality and its asymptotic p-value	P-M
10	Pitman-Morgan-McCulloch statistic for comparing dependent variances under non-normality and its bootstrapped p-values	P-M-M

Table A1: Indices in levels, univariate statistics and pairwise comparison

Univariate statistics

statistics	EPU	U	EPU	U(Kom.)	U	U(Kom.)	U(Hom.)	U(Het.)	U(Louk.)	U(Kag.)	U+-	U-	U(st.)	U(nst.)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
nobs.	201	201	170	170	201	201	305	305	305	305	305	305	278	278
means	130.043	95.877	103.349	115.460	130.043	106.121	93.250	93.255	93.250	92.909	93.250	93.275	97.108	94.325
st.devs.	76.036	43.963	43.608	53.876	76.036	54.329	41.537	41.939	41.537	42.251	41.537	41.322	68.328	71.059
skew.	1.103	0.467	0.231	0.775	1.103	0.924	0.484	0.512	0.484	0.494	0.484	0.474	1.035	1.288
pval.	0.000	0.007	0.218	0.000	0.000	0.000	0.001	0.000	0.001	0.000	0.001	0.001	0.000	0.000
kurt.	0.939	-0.920	-0.998	-0.373	0.939	-0.110	-0.796	-0.762	-0.796	-0.791	-0.796	-0.846	0.962	1.975
pval.	0.007	0.008	0.008	0.321	0.007	0.750	0.005	0.007	0.005	0.005	0.005	0.003	0.001	0.000
J-B	48.143	14.392	8.574	18.011	48.143	28.715	19.984	20.702	19.984	20.342	19.984	20.525	60.340	122.009
pval.	0.000	0.001	0.014	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Correlation and pairwise comparison														
statistics			**EPU & U(Kom.)**		**U & U(Kom.)**		**U(Hom.) & U(Het.)**		**U(Louk.) & U(Kag.)**		**U+- & U-**		**U(st.) & U(nst.)**	
	Stat	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**
Pearson	0.668	0.000	0.847	0.000	0.729	0.000	0.998	0.000	0.996	0.000	0.997	0.000	0.934	0.000
Spearman	0.629	0.000	0.849	0.000	0.693	0.000	0.998	0.000	0.995	0.000	0.996	0.000	0.930	0.000
Kendall	0.454	0.000	0.652	0.000	0.514	0.000	0.966	0.000	0.942	0.000	0.951	0.000	0.790	0.000
qcorr1	0.599	0.000	0.703	0.000	0.608	0.000	0.804	0.000	0.804	0.000	0.803	0.000	0.774	0.000
qcorr2	0.594	0.000	0.697	0.000	0.572	0.000	0.803	0.000	0.802	0.000	0.799	0.000	0.765	0.000
t-diff	-17.038	0.000	8.164	0.000	-12.483	0.000	0.016	0.494	-0.721	0.236	0.057	0.477	-2.249	0.013
F	2.991	0.000	1.526	0.003	1.959	0.000	1.019	0.433	1.035	0.383	1.010	0.464	1.082	0.257
P-M	10.910	0.000	5.197	0.000	7.058	0.000	3.018	0.001	3.253	0.001	1.122	0.131	1.818	0.035
P-M-M	8.381	0.000	-3.898	0.000	6.362	0.000	-1.150	0.121	-2.880	0.002	0.904	0.185	-3.601	0.000

Table A2: Indices in first differences, univariate statistics and pairwise comparison

Univariate statistics

statistics	EPU	U	EPU	U(Kom.)	U	U(Kom.)	U(Hom.)	U(Het.)	U(Louk.)	U(Kag.)	U+-	U-	U(st.)	U(nst.)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
nobs.	200	200	169	169	200	200	304	304	304	304	304	304	277	277
means	0.900	0.639	0.630	0.724	0.900	0.665	0.235	0.259	0.235	0.253	0.235	0.259	0.076	0.138
st.devs.	71.152	23.654	24.819	36.065	71.152	33.723	24.399	24.413	24.399	24.353	24.399	23.943	63.933	67.760
skew.	-0.086	-0.085	-0.082	-0.029	-0.086	-0.026	-0.195	-0.216	-0.195	-0.131	-0.195	-0.243	-0.276	-0.011
pval.	0.621	0.626	0.664	0.878	0.621	0.880	0.166	0.123	0.166	0.352	0.166	0.084	0.061	0.943
kurt.	0.980	0.355	0.211	0.297	0.980	0.658	0.507	0.554	0.507	0.130	0.507	0.447	3.622	5.555
pval.	0.005	0.306	0.575	0.431	0.005	0.058	0.071	0.048	0.071	0.642	0.071	0.112	0.000	0.000
J-B	8.253	1.287	0.503	0.645	8.253	3.627	5.176	6.268	5.176	1.080	5.176	5.512	154.956	356.114
pval.	0.016	0.525	0.778	0.724	0.016	0.163	0.075	0.044	0.075	0.583	0.075	0.064	0.000	0.000

Correlation and pairwise comparison

statistics	EPU & U		EPU & U(Kom.)		U & U(Kom.)		U(Hom.) & U(Het.)		U(Louk.) & U(Kag.)		U+- & U-		U(st.) & U(nst.)	
	Stat	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**	**Stat**	**pval**
Pearson	0.166	0.010	0.521	0.000	0.390	0.000	0.991	0.000	0.979	0.000	0.987	0.000	0.867	0.000
Spearman	0.197	0.010	0.516	0.000	0.325	0.000	0.987	0.000	0.974	0.000	0.985	0.000	0.819	0.000
Kendall	0.130	0.000	0.357	0.000	0.230	0.000	0.912	0.000	0.866	0.000	0.899	0.000	0.652	0.000
qcorr1	0.088	0.128	0.404	0.000	0.220	0.004	0.723	0.000	0.718	0.000	0.715	0.000	0.579	0.000
qcorr2	0.160	0.008	0.368	0.000	0.310	0.000	0.721	0.000	0.714	0.000	0.712	0.000	0.585	0.000
t-diff	-0.116	0.454	0.060	0.476	-0.109	0.457	0.057	0.477	0.034	0.487	0.050	0.480	0.044	0.483
F	9.048	0.000	2.112	0.000	4.452	0.000	1.001	0.496	1.004	0.487	1.038	0.372	1.123	0.167
P-M	19.089	0.000	5.791	0.000	12.498	0.000	0.078	0.469	0.157	0.438	2.004	0.023	1.937	0.027
P-M-M	16.579	0.000	-5.395	0.000	12.403	0.000	0.275	0.392	-0.861	0.190	1.523	0.067	-1.964	0.026

Appendix B

Major events affecting uncertainty in Russia, January 1994 - February 2018

No	Date	Event	Matched by
1	Oct-94	Steep fall of the rouble against the dollar ("Black Tuesday")	U, EPU
2	Dec-94	Beginning of the first Chechen war	EPU
3	Jul-95	The Central Bank introduces a currency corridor	U, EPU
4	Jan-96	Resignations of Chubais and Kozyrev	U, EPU
5	Jun-96	The first round of the presidential elections in Russia	U, EPU
6	Jul-96	Yeltsin elected president for a second term	U, EPU
7	Jun-97	Resolution on privatisation in 1992-1996	U
8	Jan-98	Denomination of the rouble	EPU
9	Aug-98	Default on government short-term bonds	U, EPU
10	Aug-99	Terrorist invasion on the territory of Dagestan	U, EPU
11	Sep-99	Explosions of houses in Buinaksk, Moscow, Volgodonsk	U, EPU
12	Jan-00	Yeltsin's statement on early resignation (31.12.99)	U, EPU
13	Jun-00	Arrest of Gusinsky and the scandal with the Media-Most holding	U, EPU
14	Jul-00	Putin addresses the Federal Assembly for the 1st time	U, EPU
15	Oct-02	Moscow theater hostage crisis (terrorist attack on Dubrovka)	U, EPU
16	Oct-03	Khodorkovsky arrested	U, EPU
17	Feb-04	Terrorist attack in Voronezh	U, EPU
18	Mar-04	Putin election - 2	U, EPU
19	Aug-04	Suicide bombings on two planes from Moscow	U, EPU
20	Sep-04	Seizure of a school in Beslan by terrorists	U, EPU
21	Jan-05	Protests against the monetisation of benefits in a number of cities	U, EPU
22	May-05	First sentence of Khodorkovsky and Lebedev	U, EPU
23	Oct-05	Armed attack on Nalchik, the capital of Kabardino-Balkaria	U, EPU
24	Jul-06	G8 Leaders' Summit (G8) in Strelna (St Petersburg)	U, EPU
25	Oct-06	Anna Politkovskaya killed	U, EPU
26	Jan-07	Putin names his successor as Russian president	U, EPU
27	May-08	Medvedev election	U
28	Aug-08	Georgian conflict	U, EPU
29	Sep-08	U.S. stock market crash	U, EPU
30	May-09	St. Petersburg International Economic Forum (SPIEF'21)	U, EPU
31	Nov-09	The Nevsky Express train crashed (Islamic terrorist attack)	U, EPU
32	Mar-10	Two terrorist attacks in the Moscow metro	U, EPU
33	Apr-10	The Smolensk air disaster	U, EPU
34	Dec-10	Second sentence for Khodorkovsky	U, EPU
35	Jan-11	The Domodedovo International Airport suicide bombing	U, EPU
36	Jul-11	Sinking of the motor ship "Bulgaria", 122 victims	U, EPU
37	Mar-12	Putin election - 3	EPU
38	May-12	March of the Millions on the Bolotnaya Square in Moscow	U
39	Aug-12	Verdict in the case of Pussy Riot	U, EPU
40	Dec-12	Ban on adoption of Russian children by Americans ("Dima Yakovlev's Law")	U, EPU

41	Feb-14	Sochi Olympics	U
42	Mar-14	Referendum in Crimea	U, EPU
43	Jul-14	Malaysian Boeing crash over Ukraine	U, EPU
44	Dec-14	The largest drop in the rouble exchange rate	EPU
45	Feb-15	Assassination of Boris Nemtsov	U
46	Dec-16	Falling oil prices	U, EPU
47	Mar-17	Anti-corruption protest rallies in 95 cities in Russia	U, EPU
48	Apr-17	Terrorist attack in the St Petersburg's underground	U, EPU
49	Sep-17	Large-scale wave of telephone terrorism in Russia	U, EPU