



WP/21/192

IMF Working Paper

A Sentiment-Enhanced Corruption Perception Index

by Yongquan Cao, Yingjie Fan, Sandile Hlatshwayo, Monica Petrescu, and Zaijin Zhan

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Statistic Department

A Sentiment-Enhanced Corruption Perception Index

Prepared by

Yongquan Cao, Yingjie Fan, Sandile Hlatshwayo, Monica Petrescu, and Zaijin Zhan¹

Authorized for distribution by Florina Tanase

July 2021

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

Direct measurement of corruption is difficult due to its hidden nature, and measuring the perceptions of corruption via survey-based methods is often used as an alternative. This paper constructs a new non-survey based perceptions index for 111 countries by applying sentiment analysis to Financial Times articles over 2005–18. This sentiment-enhanced corruption perception index (SECPI) captures not only the frequency of corruption related articles, but also the articles' sentiment towards corruption. This index, while correlated with existing corruption perception indexes, offers some distinct advantages, including heightened sensitivity to current events (e.g., corruption investigations and elections), availability at a higher frequency, and lower costs to update. The SECPI is negatively correlated with business environment and institutional quality. Increases in the perceived incidence or scope of corruption influences economic agents' behaviors, and thus economic dynamics. We found that when the SECPI is at least one standard deviation above the mean, the growth per capita falls by 0.65 percentage point on average, with more pronounced impacts for emerging market and low income countries.

JEL Classification Numbers: D73, E02, C38

Keywords: corruption, text mining, sentiment analysis

Author's E-Mail Address: YCao2@imf.org; Yingjie@princeton.edu; SHlatshwayo@imf.org; MPetrescu@imf.org; ZZhan@imf.org

¹ We thank Yunhui Zhao, Tatiana G. Goriainova, Mamoon Saeed, and IMF library for assistance with access to FT news archive. We would also like to thank Pheabe Morris and other participants at the IMF seminar for their valuable comments.

Contents	Page
Abstract.....	2
I. Introduction	4
II. Literature Review	5
III. A Sentiment-Enhanced Corruption Perception Index	7
A. Data Sources.....	7
B. Sentiment Analysis Methods.....	8
C. Construction of the Sentiment-Enhanced Corruption Perception Index.....	8
D. Robustness of the Index	9
E. Comparison with the Existing Indexes.....	11
F. Corruption, Elections, Short and Long-term Trends	13
IV. Macro-Relevance Analysis	16
A. Simple Correlation.....	16
B. Local projections.....	17
V. Conclusions	20
References.....	22
 Figures	
Figure 1. Three Generations of Corruption-Related Measurements	6
Figure 2. Comparison with the CoC and the CPI	12
Figure 3. Comparison with NIC.....	13
Figure 4. Corruption Perceptions, Elections, and Investigations: Case Studies	15
Figure 5. Correlations Between Negative Corruption Sentiment & Structural Indicators.....	16
Figure 6. Effect on GDP Per Capita Growth, Full sample	18
Figure 7. Macroeconomic Impacts of Negative Corruption Sentiment, Advanced Economies	19
Figure 8. Macroeconomic Impacts of Negative Corruption Sentiment, Emerging Market & Low-Income Countries	20
 Tables	
Table 1. Summary Statistics	7
Table 2. Correlation with the SECPI.....	10
Table 3. Summary of Human Audits in Quantifying Sentiment	10
 Boxes	
Box 1. Refinements in Extracting Country Information from FT articles	11
Box 2. Case Studies.....	14
 Annex	
Annex I. List of Variables for Local Projections	26

I. INTRODUCTION

Corruption is difficult to measure directly, and measuring the perceptions of corruption via survey-based methods is often used as an alternative. Recent years have seen a growing demand for more robust measurement of corruption from international investors, politicians, citizens, and international organizations.² Nevertheless, measuring corruption directly based on the prevalence of realized corruption events remains challenging, in part because of a wide range of concealed behaviors that constitute corruption.³ Against this backdrop, corruption is more frequently measured indirectly by gauging the perception of the extent of corruption in a specific country, based on subjective opinions of citizens, experts, or stakeholders. With recent technological advances allowing for text mining, newspapers offer a new source for measuring public perception of corruption, providing a welcome alternative to survey-based methods.

This paper constructs a novel sentiment-enhanced corruption perception index (SECPI) for 111 countries from 2005-18.⁴ In the process, the SECPI considers not only the number of corruption-related articles from the Financial Times (FT) over 2005–18, but also corruption perception by conducting in-depth sentiment analysis of these articles, based on the Loughran-McDonald financial dictionary. The sentiment analysis allows us to better gauge the perception of corruption and helps this index to achieve an improved performance over other news-based indexes that only quantify the frequency of mentions of corruption. With hundreds of thousands of FT news articles in the last two decades, covering major political, economic, and social events around the globe, we have an abundant number of corruption-related articles upon which to draw in order to measure corruption perceptions across countries and over time when we construct the SECPI.

The SECPI, while correlated with existing corruption perception indexes and business environment and institutional quality indicators, offers distinct advantages. When comparing the SECPI with two widely-used measures of corruption perception, the Corruption Perception Index (CPI) from Transparency International and the Control of Corruption (CoC) by Kaufmann, et al. (2010), the correlation coefficients⁵ between SECPI and the other two indexes are around -0.9.⁶ Furthermore, we find that countries with less negative corruption perception as measured by the SECPI have better governance and business environments, and less poverty. The SECPI has three advantages over existing corruption perception indexes. First, the SECPI reacts quickly to current events. For example, the SECPI can quickly pick up shifts in corruption perceptions caused by a corruption investigation. In addition, the SECPI tends to spike during election years when corruption is sometimes a major campaign or policy topic. Second, the SECPI can be updated at higher

² In 2018, the IMF adopted a new [framework](#) for enhanced engagement on governance, highlighting some of the limitations of existing measures and calling for a holistic approach to measurement that draws on multiple complementary measures of corruption.

³ See, for example, the UNODC's [module](#) on measuring corruption.

⁴ The SECPI is not an IMF-approved or recommended measure and is presented only for the purposes of this research.

⁵ The correlation coefficients are calculated in two steps. First, we group countries into advanced economies, emerging market and low-income countries and calculate average values over time across different groups. Second, we pool them together and calculate correlation coefficients between two indexes over time.

⁶ The negative sign of the correlation coefficients is due to the index construction that, as the corruption perception deteriorates, the SECPI increases but the CoC and the CPI decreases.

frequency, as FT news are published daily, while most existing survey-based corruption perception indexes are produced on an annual basis. Third, the SECPI is much less costly to construct and to update, relative to survey-based corruption perception indexes.

Consistent with much of the existing literature, we find a negative relationship between corruption perceptions, as measured by the SECPI, and growth. Increases in the perceived incidence or scope of corruption influence economic agents' behaviors, and thus economic dynamics. A more negative perception of corruption is associated with lower growth, especially in emerging market and low-income countries; during periods when corruption perception is at least one standard deviation above the mean, GDP per capita growth is on average 0.65 percentage point lower cumulatively by the second year.⁷ Splitting the sample based on income levels shows that this effect is only significant in emerging and low-income countries, where household consumption and private investment in percent of GDP also fall.

The rest of the paper is organized as follows. Section II introduces the related literature and highlights our contributions. Section III demonstrates the construction of the index and offers some stylized facts. Section IV shows the macro-relevance analysis, and Section V concludes.

II. LITERATURE REVIEW

There are broadly three generations of measures of corruption, with the SECPI forming part of the third generation of measures, which are unified by their use of big data (Figure 1). The first generation of corruption measures, which includes the well-known Corruption Perceptions Index, relied on surveys that measure experts' and/or citizens' perceptions of the prevalence and nature of corruption. The second generation of measures moved to more direct measures of people's and firms' experiences with corruption using victimization surveys (e.g., asking for direct evidence of how often bribes were requested and of what amounts) and bureaucratic quality indicators (e.g., tax collection efficiency or fiscal transparency). Examples of the second generation indicators include the Global Corruption Barometer. Finally, over the past decade, with increasingly improved access to large databases and tools to scrape and code data using algorithms, the third generation of corruption measures has emerged. Some notable examples include India's IPaidABribe.com, procurement analyses that identify contractual outliers that might uncover corrupt practice (e.g., Fazekas, Mihály, and Gábor Kocsis 2020), and news-flow measures of corruption by Hlatshwayo, et al. 2018. This generation of measures improves on the earlier approaches by adopting reproducible and less subjective approaches (e.g., without reliance on the often sticky perceptions of citizens or assessments of a select group of experts). Our index contributes to improving the third generation of corruption perception indexes by measuring the sentiment of corruption-related articles using sentiment analysis.

⁷ The immediate impact is 0.5%.

Figure 1. Three Generations of Corruption-Related Measurements



Source: Hlatshwayo, et al., 2018.

Sentiment analysis has been applied in the field of economics and social sciences extensively to better inform public opinion, media tone, and market expectations (Young and Soroka 2012; Cambria et al., 2013; Oliveira et al., 2017). Sentiment analysis is normally carried out by two methods: dictionary-based and machine learning-based approaches. Dictionary-based approaches use sentiment lexicons, which are collections of words that convey feelings, to match such feelings with text in the documents and calculate the polarity of a body of text (Eshbaugh-Soha 2010; Liu 2015). A sentiment lexicon is a list of words that are either dichotomously classified as positive or negative or contain more continuous measures of their content. Machine learning-based approaches apply statistical classification techniques, such as support vector machines or neural networks (Manning & Schutze 1999; Hopkins & King 2010), to classify text. Depending on whether human coders classify a representative proportion of the text as the basis, supervised and unsupervised learning applies algorithms to sort the rest of the documents into categories.

The paper adopts a widely-used dictionary-based approach with discretionary caution. Scores in the sentiment lexicon should be reflective of the context. Sentiment lexicons are usually developed for a wide array of applications. While applying them in a specific area, the measurement of the sentiment could be inaccurate and imprecise (Krosnick 1999; Grimmer and Stewart 2013). Other than off-the-shelf dictionaries, there are dictionaries developed for a specific field, such as Loughran-McDonald sentiment lexicon, which is created for use primarily with financial documents (Loughran et al., 2011). It rates commonly used words, which might have different meanings in earning reports, in the field of finance. Therefore, this paper uses Loughran-McDonald sentiment lexicon as the benchmark, because it is considered to be relatively accurate to capture the sentiment in the field of economics and finance, which fits well for FT news articles.

On the impact of corruption, a rich body of literature has long documented its negative effects on public finances, monetary and financial stability, confidence and trust, and, in turn, growth. Perceptions of extensive corruption can lower tax compliance and increase evasive practices (Aghion et al., 2016). Corruption can also lead to leakages within public spending and/or poorer quality of executed projects. Similarly, corruption can increase fiscal deficits and lead to substantial debt accumulation (Tanzi and Davoodi, 1998; Kaufmann, 2010; Achury et al., 2015). Corruption is sometimes coupled with fiscal dominance, leading to inflation and eroding the independence of monetary policy (Huang and Wei, 2006; Cavoli and Wilson, 2015). Furthermore, perception of extensive corruption can lead to increased

borrowing costs via higher risk premia and higher default risk (e.g., Akitoby and Stratmann, 2010). A recent literature has also highlighted the role corruption can play in lowering households' and firms' trust in institutions and processes, with adverse impact on formal financial market participation and investment (Guiso, Sapienza, and Zingales, 2009; Sapienza & Zingales, 2012). Together with muted productivity due to misallocated resources and brain drain (Gupta et al., 2002; Rajkumar and Swaroop, 2008), the above channels can lower growth and harm competitiveness (Campos, et al., 2010; Hlatshwayo, et al., 2018). This paper confirms the claims in the existing literature that the corruption perceptions have negative impacts on both structural variables and growth.

III. A SENTIMENT-ENHANCED CORRUPTION PERCEPTION INDEX

A. Data Sources

We constructed the SECPI using the FT online archive over 2005–18, which includes over 800,000 articles. We selected the corruption related FT articles based on keywords starting with “corrupt” and “brib”, which includes “corruption”, “corrupted” and “bribery”, etc.⁸ The selected database consists of 27,935 articles covering 182 countries, of which 47 countries have at least one selected article every year. Table 1 shows the summary statistics of advanced economy, emerging market, and low-income countries.^{9,10} The emerging market has the greatest number of corruption related articles. However, its ratio of corruption-related articles over the total number of articles mentioning these countries is not the highest. The issue of FT’s uneven coverage across countries will be addressed in Section C.

Table 1. Summary Statistics¹¹

	# of articles	# of corruption articles	Ratio (%)
Advanced economy	1,420,344	10,431	0.73
Emerging market	477,953	22,028	4.61
Low income countries	53,182	4,731	8.90

There are some limitations of using the FT as the sole source for the SECPI. First, the sentiment based entirely on FT articles may reflect the opinions of the FT editors on corruption across countries. It does not necessarily reflect the opinions of a wider range of people, such as the citizens of the country or the international communities. Second, FT news focuses primarily on big events that may have sizable macroeconomic and/or political implications. Therefore, some local corruption news (e.g., covering petty corruption), which are more likely captured in local newspapers, could be neglected by the FT. Nevertheless, it is worth noting that the potential selection bias goes beyond the size of the events, and it is often

⁸ This keywords-based selection strategy performed well in identifying corruption-related articles. The type I and type II errors are very limited based on human reading.

⁹ The categorization of countries follows the following link: <http://www.imf.org/external/datamapper/FMEconGroup.xlsx>.

¹⁰ The methodology to identify the country information is detailed in Box 1.

¹¹ The number of articles on a specific country is calculated by summing up all articles mentioning the country. An article can mention more than one country, thus counted multiple times during the calculation. Table 1 shows the number of articles, the number of corruption-related articles, and the ratio between the two numbers by grouping countries in the three categories. It is worth noting that Table 1 does not imply the ranking of corruption perception.

affected by FT's own preference. Third, the country coverage of the FT is uneven. For some low-income countries, there are less than ten corruption related articles during 2005–18. Fourth, the database only contains English articles, which prevents us from analyzing the corruption perception in news in other languages. Section C describes how we address some of these limitations.

B. Sentiment Analysis Methods

We conduct in-depth sentiment analysis on the corruption related articles using a sentiment lexicon. Sentiment analysis is essentially a natural language processing mechanism to extract information to determine the attitude of authors on some chosen subjects from a large number of documents (Turney, 2002; Pang and Lee, 2004; Hu and Liu, 2004; Kim and Hovy, 2004; Wilson et al., 2005; Agarwal et al., 2009). A sentiment lexicon is a list of words scored according to their emotion or opinion content. Sentiment lexicons, developed by crowdsourcing hand-coded sentiment of an individual word, may be binary (e.g., positive or negative), or have more specific categorization (e.g., anger, joy, and sadness). The number of words from a sentiment lexicon is limited, as not all English words contain sentiment. Sentiment scores are calculated by matching a selected sentiment lexicon with the body of text of articles and aggregating the assigned value of matched words. The final score provides an aggregated measure of a document's tone.

The Loughran-McDonald Financial dictionary¹² was selected as our benchmark sentiment lexicon. Relative to other popular lexicons, this lexicon includes more economic and financial terminology, which would help study the corruption perception index's macroeconomic relevance. The dictionary-based approach cannot be applied without caution and validation. We use other popular sentiment lexicons, such as the Harvard-IV dictionary, Henry's Financial dictionary, and Qdap Dictionaries, to construct alternative indexes to check the benchmark index's robustness. Furthermore, we validate the lexicon through human audits (see Section D).

C. Construction of the Sentiment-Enhanced Corruption Perception Index

The SECPI is constructed in the following four steps:

Step One: We select the corruption related FT articles as described in Section A. We further restrict our benchmark database to the countries with at least one article per year in at least seven of the 14 years' sample period. Without this restriction, we would have included 182 countries, one sixth of which have articles in less than two years out of 14 years. There are two reasons as to why some countries do not have enough corruption-related articles. First, FT does not have much coverage for those countries (potentially because their readerships in those countries are not large enough or because FT chooses not to cover events in those countries based on its own preference). Second, the country does not have significant corruption related issues. Adding countries without enough coverage will create noise in our index. On the other hand, if we restrict our benchmark database to the countries with articles every year, we are left with only 47 countries. The restriction would exclude many emerging market and low-income countries, which do not have FT coverage every year but valuable information can still be gleaned from FT. Therefore, we exclude countries with less than

¹² Details about this dictionary can be found at <https://drive.google.com/file/d/15UPaF2xJLSVz8DYuphierz67trCxFLcl/view>

seven years of coverage in order to strike a balance between accuracy and country coverage. We end up having 111 countries in the sample.

Step Two: For each selected article, we use the negative word list from Loughran-McDonald financial dictionary to calculate a sentiment score. The benchmark score of article i , $sentiment_i$, is calculated by aggregating the sentiment of individual words in the article, which have a sentiment score. As a robustness check, we construct alternative indexes using both negative and positive words in the Loughran-McDonald financial dictionary as well as using other popular dictionaries (see Section D).¹³

Step Three: For country j at time t , we aggregate all corruption-related articles' sentiment scores. A simple sum would be biased since some countries have more extensive FT coverage than others, according to Table 1. Thus, we weight the sum for each country by the total number of FT articles mentioning the country j . The scores for country j at time t are calculated using the following formula:

$$scores_{t,j} = \frac{\sum_{\{i|time=t, country=j\}} sentiment_i}{\sum_i 1\{country = j\}}$$

Where $sentiment_i$ is the sentiment score of article i . The $1\{country = j\}$ is an indicator function, which is equal to 1 when the article is about country j and 0 otherwise.

Step Four: We construct the sentiment-enhanced corruption perception index based on the ranking of the scores. $Scores_{t,j}$ by construction have a very wide range, thus are hard to interpret. Against this backdrop, we calculate $Index_{t,j}$ as the percentile of the $scores_{t,j}$ among all $scores$. The index is bounded within zero and one. One means the most negative corruption perception and zero means the least negative corruption perception across countries and time.

$$Index_{t,j} = Prob(scores < scores_{t,j})$$

As the percentile is calculated using the scores in all 14 years, the index can capture both the cross-sectional movements and changes over time.¹⁴

D. Robustness of the Index

The SECPI is robust when different sentiment lexicons are applied. We constructed alternative indexes by replacing the negative word list of Loughran-McDonald dictionary with seven different dictionaries, including the Harvard-IV dictionary (SentimentHE), negative word list of the Harvard-IV dictionary (NegativityHE), the Henry's Financial dictionary (SentimentGI), negative word list of the Henry's Financial dictionary (NegativityGI), the Qdap dictionaries (SentimentQDAP), negative word list of the Qdap dictionaries (NegativityQDAP) and the Loughran-McDonald dictionary (SentimentLM). The correlation

¹³ See <https://cran.r-project.org/web/packages/SentimentAnalysis/SentimentAnalysis.pdf> for details. R and Python are used in constructing the index.

¹⁴ One drawback of the index is that the value will slightly change once new articles are included as the same article will have slightly different percentile in the new dataset.

coefficients between the benchmark index and other alternative indexes across countries and time are generally high, except for SentimentQDAP¹⁵ (Table 2).

Table 2. Correlation with the SECPI

Sentiment Lexicon	Correlation
NegativityQDAP	0.99
NegativityHE	0.90
NegativityGI	0.99
SentimentQDAP	0.29
SentimentLM	0.98
SentimentHE	0.70
SentimentGI	0.65

The SECPI provides consistent results with human audits in quantifying the sentiment on corruption-related articles. As people may have different views on the same article, we designed specific guidance to rank the corruption-related articles. We group all articles in the following five categories, from the most negative to the least negative:

1. The articles on confirmed corruption with big economic or political impacts.
2. The articles on corruption under investigation with big economic or political impacts
3. The main topic is not about corruption, but corruption is an important factor. The overall sentiment of the article is negative.
4. The main topic is not about corruption, but anticorruption is an important factor. The overall sentiment of the article is positive.
5. The main topic is about anticorruption and its good economic or political impacts.

We audited 147 articles and grouped them into five categories based on the negativity of the corruption identified therein, from category 1 (most negative) to category 5 (least negative); the average sentiment scores decline across the categories (Table 3), which means the sentiment scores are broadly consistent with our expectation.

Table 3. Summary of Human Audits in Quantifying Sentiment

Category	Average Scores	Count
1	0.140	2
2	0.109	15
3	0.066	109
4	0.066	17
5	0.057	4

¹⁵ SentimentQDAP uses the dictionary built by Hu and Liu (2004) based on customer reviews, which contains a different positive word list than others.

Box 1. Refinements in Extracting Country Information from FT articles

Identifying a country in FT news articles is not a straightforward task, given that the database does not provide country information. There are several difficulties. First, sometimes the country is not mentioned, but the major city in the country is mentioned. For example, China is not mentioned but Beijing is mentioned. Second, cities from different countries can share the same name, for example, London, United Kingdom and London, Canada. Third, several countries are mentioned in the same article, such as a corruption case related to a multinational company.

To address these difficulties, we build a dataset consisting of the countries and their major locations and adopt the following two refinement measures. First, to minimize errors, we use the most popular location and country pair when the city names are shared by many countries. Second, we allow only one country linked to one article. In most FT articles, more than one country is mentioned. However, around 80 percent of all human audited articles focus on the issues in only one country. As part of our approach, for example, if we have an FT article reporting a corruption case of an important political figure in country X, who maintains a close tie with country Y, only country X should be linked to the article instead of both country X and country Y. In practice, we link the most frequently mentioned country to the article as we assume the country mentioned most frequently is the focus of the article. Also, Hlatshwayo, et al. 2018 found only a small number of articles that point to supply-side corruption involving multiple countries.

The articles' main location focus is mostly correctly identified using our method, according to the human audits. Among 1,070 articles in the human audits, 830 articles are linked to the correct countries, 22 are wrongly linked and others can be linked to not just one country.

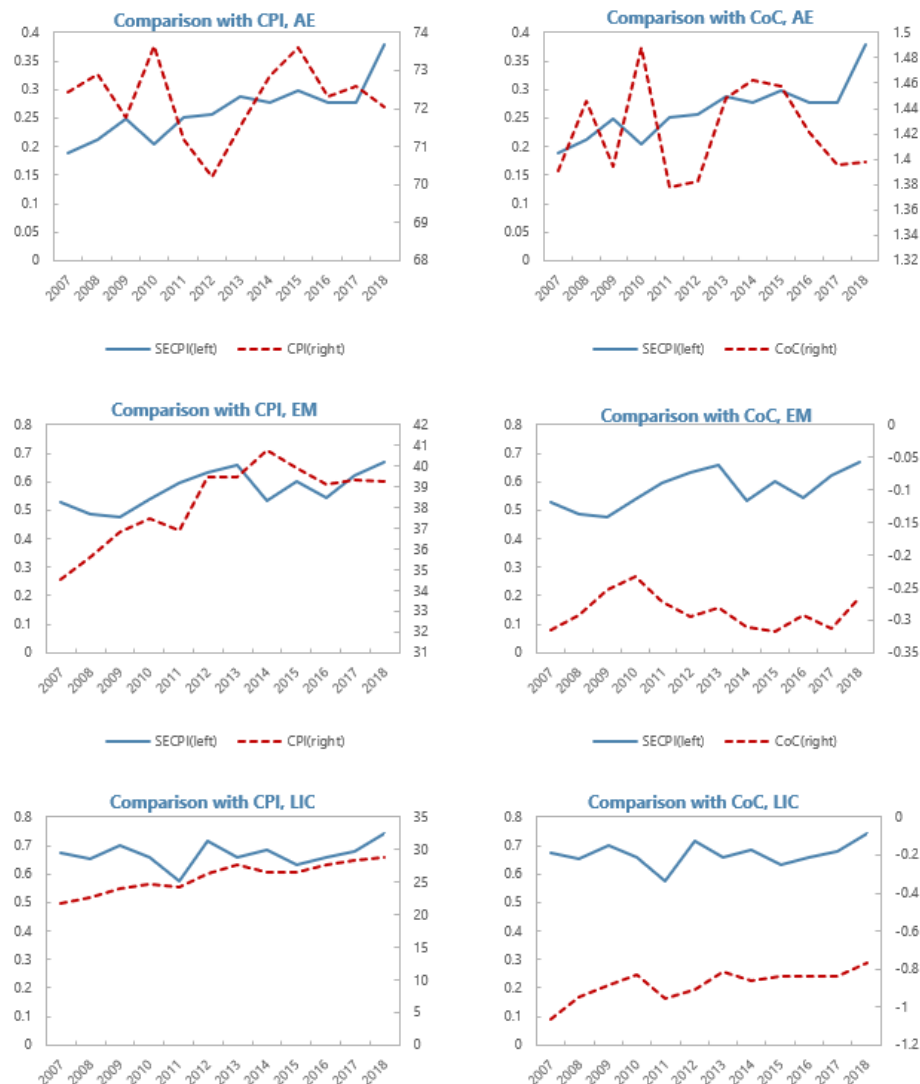
E. Comparison with the Existing Indexes

The SECPI is well correlated with the existing popular corruption perception indexes.

We compare the average values of SECPI among different country groups with the CoC and the CPI. The SECPI negatively correlates with the CoC and the CPI across country groups (Figure 2). The correlation coefficients are around -0.9 for each pair.¹⁶ All three indexes share the similarity that the average corruption perception is less negative in advanced economies (AEs) than in emerging market (EMs) and low-income countries (LICs). We also compare the SECPI with the NIC developed in Hlatshwayo, et al. 2018, which is close to our index in terms of methodology. These two indexes follow the same trend across different country groups, with a correlation coefficient around 0.65 (Figure 3).

¹⁶ See footnote 5.

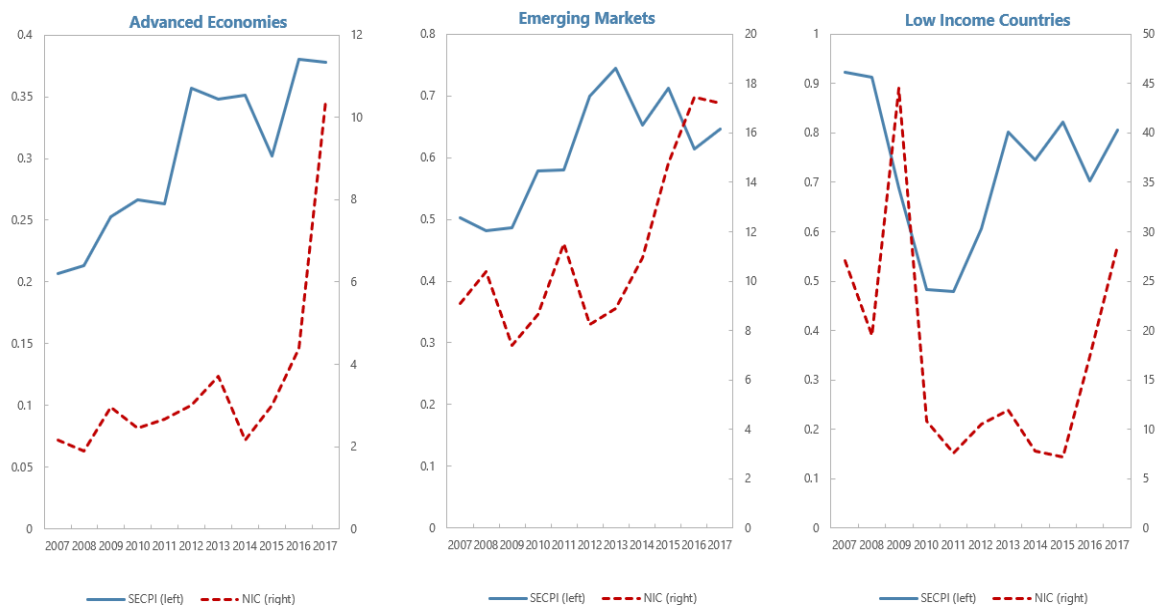
Figure 2. Comparison with the CoC and the CPI



Sources: Transparency International, the World Bank, and authors' calculation.

Note: The Corruption Perceptions Index (CPI) is a composite indicator used to measure perceptions of corruption in the public sector in different countries around the world, drawing upon several sources. Control of corruption captures perceptions of the extent to which public power is exercised for private gain, including both petty and grand forms of corruption, as well as "capture" of the state by elites and private interests. See <https://transparency.hu/wp-content/uploads/2018/02/CPI-2017-Technical-Methodology-Note-English.pdf> for the methodology of CPI and <http://info.worldbank.org/governance/wgi/#doc-methodology> for methodology of Control of Corruption.

Figure 3. Comparison with NIC



Source: Hlatshwayo, et al. (2018) and author's calculation.

Note: NIC is the big data and cross-country news flow indices of corruption by running country-specific search algorithms over more than 665 million international news articles. See <https://www.imf.org/en/Publications/WP/Issues/2018/08/31/The-Measurement-and-Macro-Relevance-of-Corruption-A-Big-Data-Approach-46157> for details.

F. Corruption, Elections, Short and Long-term Trends

A closer look at within-country trends reveals that the SECPI identifies long-term trends in corruption perception similar to other indexes but is more responsive to individual events in the near term. It performs similar to other indexes in detecting long-term trends in corruption, which are at times associated with long-term changes in growth, stock market returns, or FDI inflows (see Box 2). However, in cases where a shift in corruption perceptions is associated with a singular event (e.g., start of a corruption investigation or a change in political leadership), the SECPI reflects most of the change during the year of the event. The start of new high-profile corruption investigations, new allegations, or changes in political regimes can lead to sharp increases in the SECPI during the year these events occur, while these events are likely reflected in other indexes by more gradual shifts and/or with a lag. Given its ability to capture perception shifts rapidly, the SECPI may allow for a more accurate analysis of the relationship between corruption sentiment and fast-moving economic indicators.

Unlike other corruption perception indexes, the SECPI often spikes in election years.¹⁷ The average value of the index during a legislative or executive election year is 0.05 higher

¹⁷ This feature also highlights that the corruption perception indicators, while a good proxy of corruption, may be affected by special events, such as elections.

than outside election years (controlling for country-specific effects). Spikes in election years appear to occur when corruption is a key campaign issue and thus receives additional news coverage (see Box 2 for examples). To the extent the campaign process offers an opportunity to take stock of policy priorities, a spike in the SECPI during election years suggests a high importance of corruption among policy issues. There is no empirical evidence that the size of the spike depends on the overall level of corruption (an increased focus on corruption in election years is not a unique feature of countries with higher levels of corruption). However, there is some evidence that election-related spikes are higher (lower) when countries are experiencing lower (higher) growth than recent historical trends; this is consistent with the view that corruption is more likely to emerge as a popular concern during economic downturns.

Box 2. Case Studies

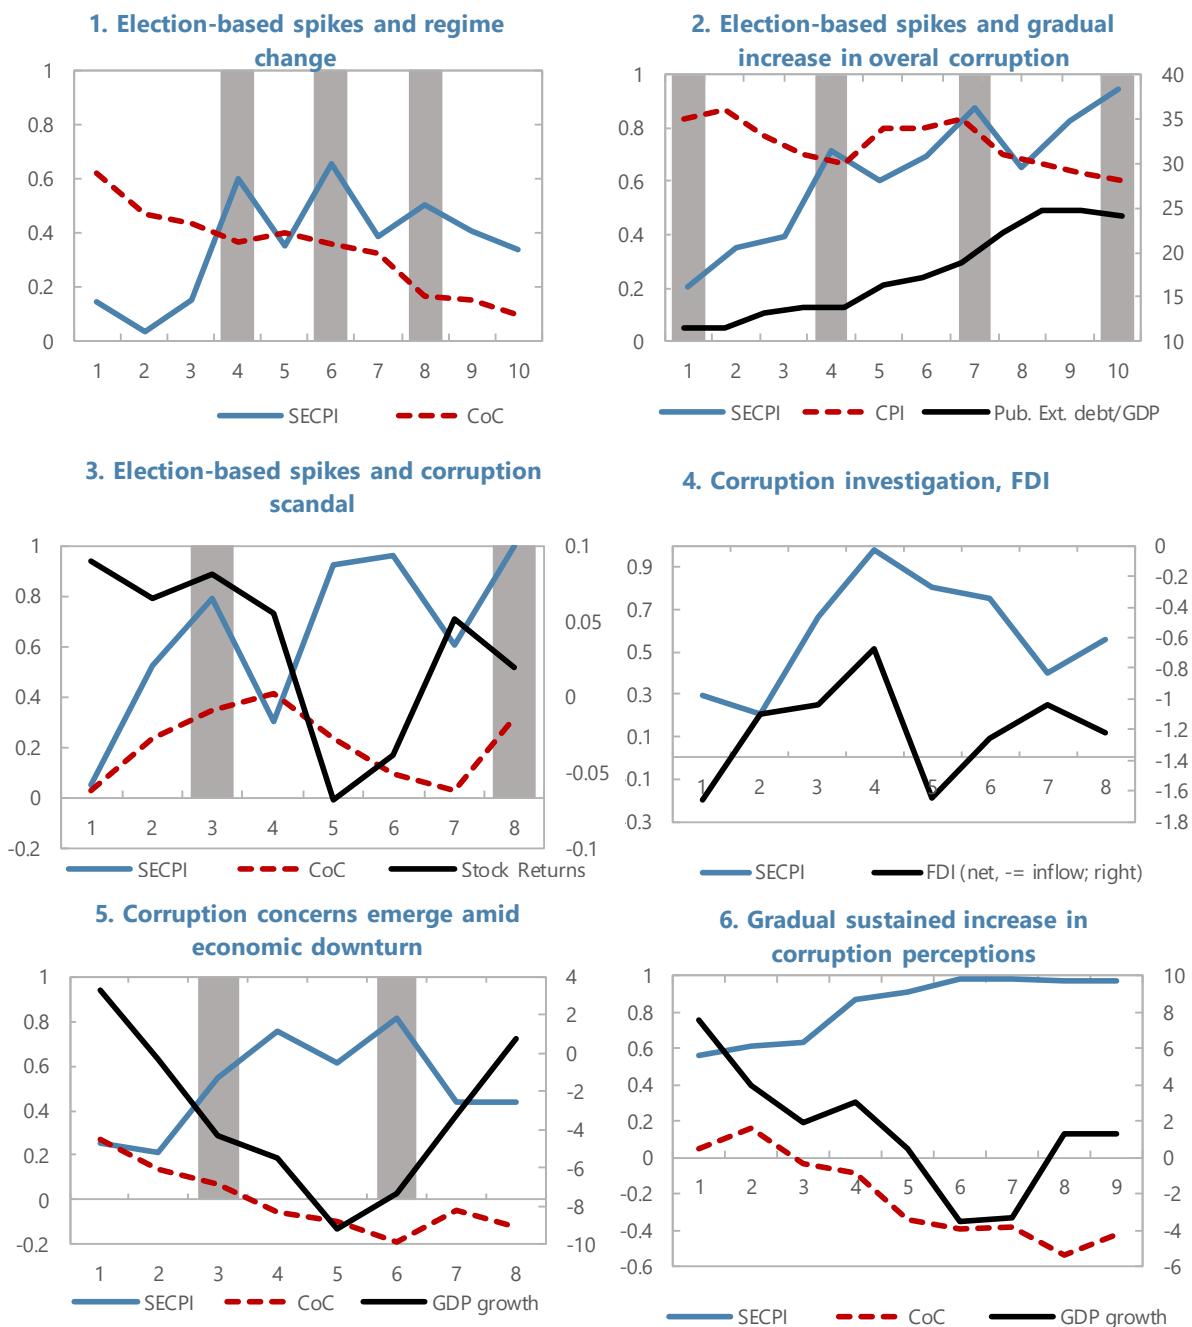
This box considers six case studies to illustrate the unique evolution of the SECPI in election years, and the reaction of the index to both short-term and long-term developments related to corruption (see Figure 4). The analysis is based on an assessment of the relevant articles for each case.

Elections. Case 1 (top left chart) is one example in which the SECPI spiked during legislative and executive election years, and a regime change led to a permanent shift in corruption perceptions. During the first election depicted (year 4), corruption linked to the ruling party became a significant campaign topic, and the SECPI spiked; the election resulted in a regime change. In the years that follow, the SECPI remained above the pre-election levels, as anti-corruption institutions are reported to lose their independence under the new ruling party, and news articles focused on worsening corruption outcomes reflected in various reports and surveys. The CoC index similarly suggests an increase in corruption over this period – though the change is more gradual. Corruption remained a key topic in subsequent elections, leading to further spikes in the SECPI. Case 2 (top right chart) is another example where the index spiked during election years, as corruption was identified as one of the most important areas for policy reform in multiple political campaigns. At the same time, there was an upward long-term trend in the SECPI driven by the emergence of multiple high-profile corruption scandals involving corporations and government institutions; a similar long-term trend was identified by the CPI. The positive correlation between corruption and debt, linked by fiscal deficits and low growth, observed in the literature, was also present.

Corruption investigations. The next two case studies consider how the SECPI performs when single events – such as large-scale investigations – drive corruption perceptions. In case 3, in addition to spikes in election years, with corruption and cronyism a central issue in elections, the index spiked when an investigation involving large sums of money and high-level officials dominated the news (year 5-6). Unlike the CoC, which shows a more gradual shift, the SECPI increased sharply in the year that the investigation became a central topic in global news; stock market returns dropped sharply in the same year. In case 4, the SECPI also spiked as the news focused on a corruption investigation, which revealed incriminating evidence against high-level officials involving a large sum of money and government contracts (year 3-4). The news also identified that the high-level scandal unsettled markets and concerned investors, consistent with lower FDI inflows. The SECPI declines gradually after the incident but remains elevated as the implicated officials remained in power.

Macroeconomic trends. The next two case studies consider examples of direct links between economic conditions and corruption perceptions. In case 5, the SECPI rose due to an increased focus on corruption as a primary cause of low tax revenues when fiscal balances worsened rapidly during an economic crisis, marked by large negative growth rates. Concerns over corruption persisted as the crisis deepened, but after political and economic stability was restored, the focus on corruption dissipated, and the benchmark index reverted. In case 6, a sequence of corruption investigations involving high-level officials led to a gradual increase in the SECPI and the CoC. In parallel, growth declined dramatically and turned negative, which was partly linked to political paralysis resulting from the corruption scandals.

Figure 4. Corruption Perceptions, Elections, and Investigations: Case Studies

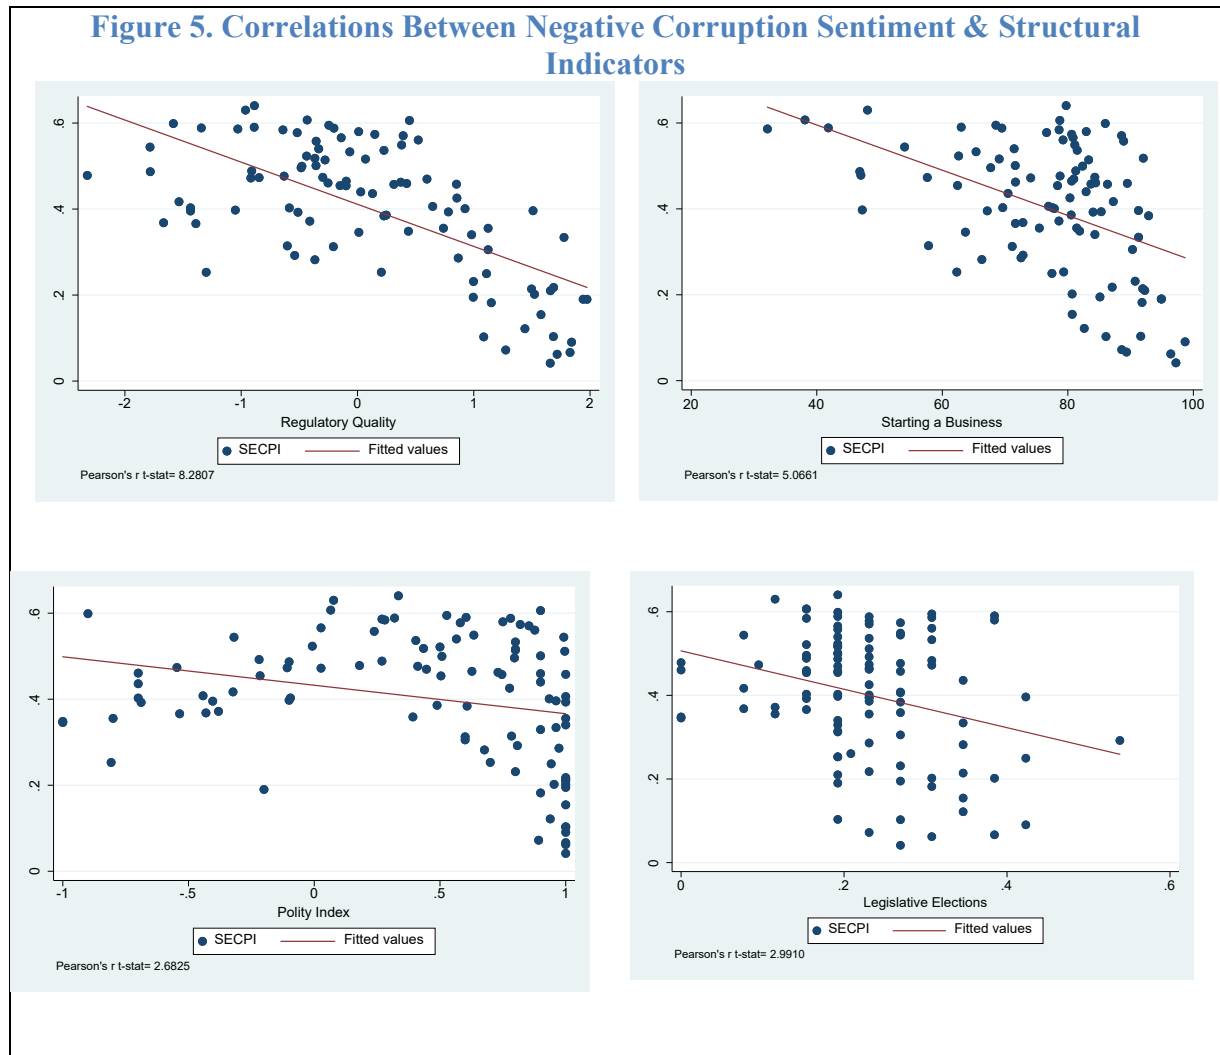


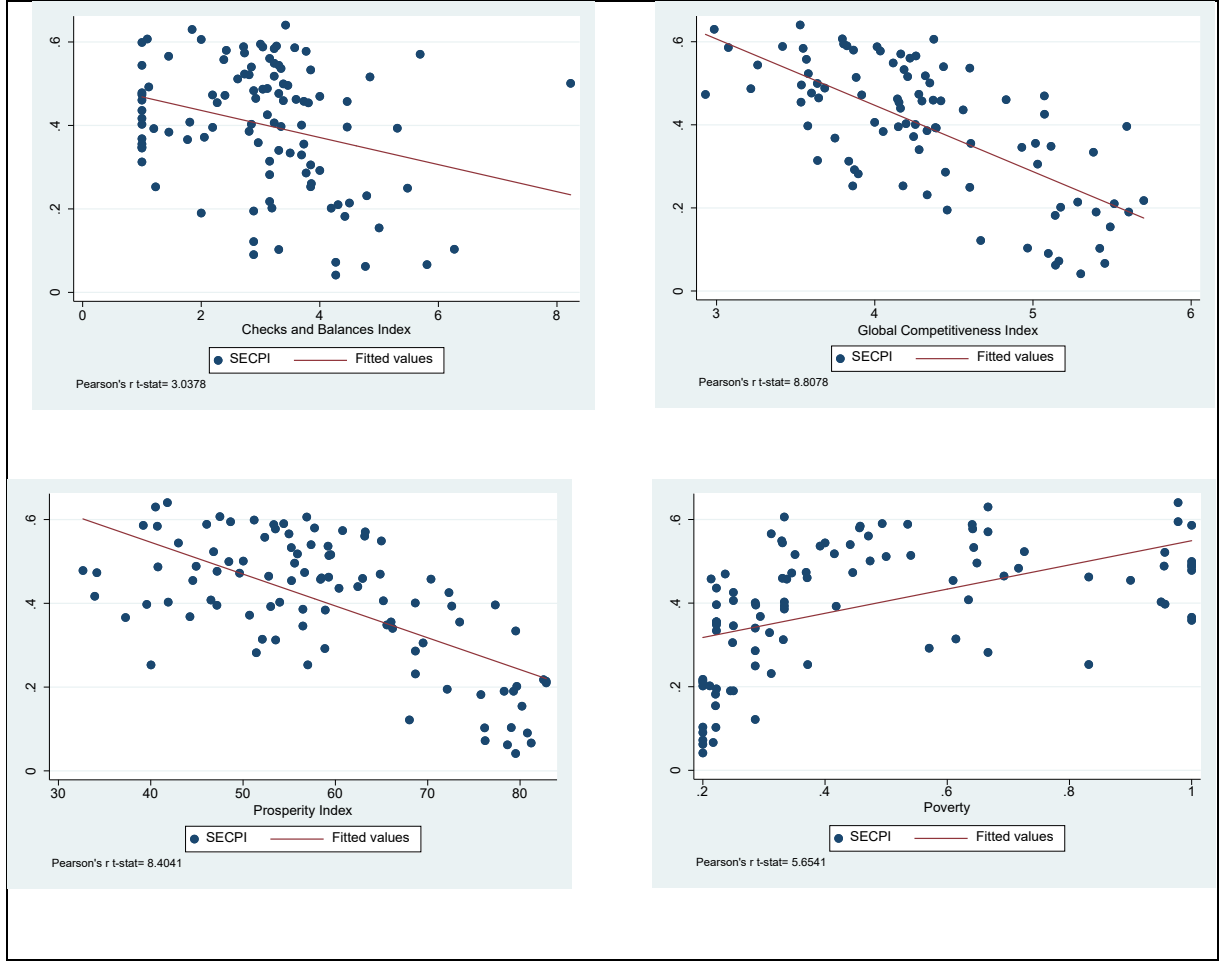
Sources: Transparency International, the World Bank, and authors' calculation.

IV. MACRO-RELEVANCE ANALYSIS

A. Simple Correlation

Negative corruption sentiment is significantly correlated with several structural indicators of the business environment, institutional quality, and poverty. These measures include the Worldwide Governance Indicator of Regulatory Quality; the World Bank's Ease of Doing Business Index; World Economic Forum's Global Competitiveness Index; and the Legatum Institute's Prosperity Index, a multi-pillar measure of prosperity that includes personal freedom, safety and security, environmental, economic, and other factors. It is no surprise that higher levels of negative corruption sentiment are correlated with lower regulatory quality, lower ease of doing business, fewer checks and balances, worse prosperity, and higher level of poverty (see Figure 5). Based on Pearson's correlation test statistics, these relationships are statistically significant. The correlation results for several of these indicators are also significant even after we control for differences in income per capita (e.g., for Starting a Business, Regulatory Quality, Global Competitiveness Index, Polity Index, and the Prosperity Index), suggesting that these structural correlations are not mere reflections of differences in level of development.





B. Local projections

This work uses local projection approaches combined with machine-learning based variable selection to capture negative corruption sentiment's effects. To empirically examine the dynamic effects of sentiment shocks on macro outcomes, we used Jorda (2005) local projections at an annual frequency with country and year fixed effects, where $p = 2$ lags and $h = 5$ years to account for persistency in the effects of some controls and to consider outcomes over the medium-term horizon (5 years).

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \tau_t^h + \sum_{l=-1}^{L=1} \beta^h SECPI Shock_{i,l} + \sum_{p=0}^P \theta_p^h X_{i,t-p} + \sum_{p=0}^P \gamma_p^h \Delta X_{i,t-p} + \sum_{p=1}^P \omega_p^h y_{i,t-p} + \sum_{p=1}^P \mu_p^h \Delta y_{i,t-p} + \varepsilon_{i,t}$$

Where:

y = outcome variable (e.g., GDP per capita growth)

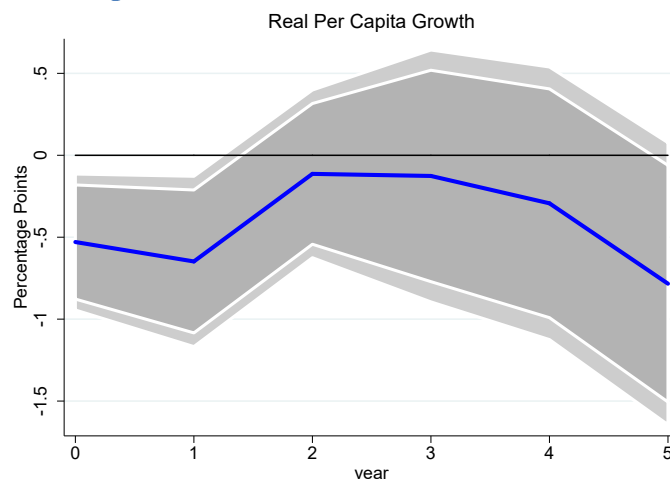
α = country fixed effects

τ = time fixed effects

$SECPI Shock$ = sentiment shock (see below)

X = selected control variable

Figure 6. Cumulative Effect on GDP Per Capita Growth, Full sample

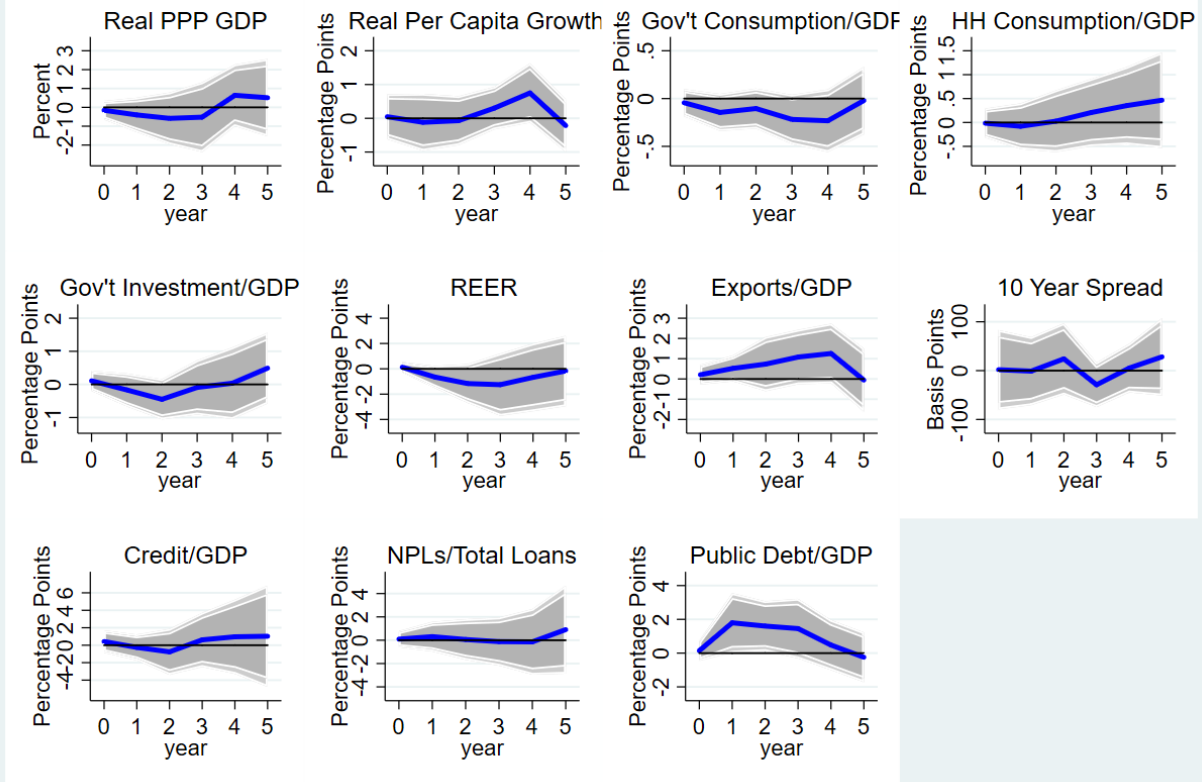


To capture periods with exceptionally high negative corruption sentiment, we create a “shock” dummy, which equals one when corruption sentiment is more than one standard deviation above the mean—where both the standard deviation and means are country-specific. In the context of model uncertainty, we allowed for a battery of possible controls covering prices and exchange rates; volatility; external and domestic demand; and structural variables (e.g., bureaucratic quality, democratic accountability, religious tensions, etc.; see Annex I for a complete list). In addition to more than 25 controls, we also included transformations of the controls (e.g., levels, two lags, first differences, and lagged first differences) and lagged dependent variables. To specify the models, we leveraged Belloni, et al’s (2014) Lasso-based variable selection procedure, which identifies controls whose omission would otherwise generate omitted variable bias, helping address concerns about endogeneity.¹⁸ This procedure, as the authors note, “allows for imperfect selection of the controls and provides confidence intervals that are valid uniformly across a large class of models.” In addition to controlling for executive and legislative elections, both leads and lags of the shocks are added for robustness and standard errors are clustered at the country level. The panel used in our study covers 111 countries and the years of 2007–18.

¹⁸ For real output, the selected variables include growth deviations from 5 year averages; population growth, government interest expenditures/GDP, total debt/GDP; trading partner growth and import demand, and the current account balance/GDP.

Figure 7. Cumulative Macroeconomic Impacts of Negative Corruption Sentiment, Advanced Economies

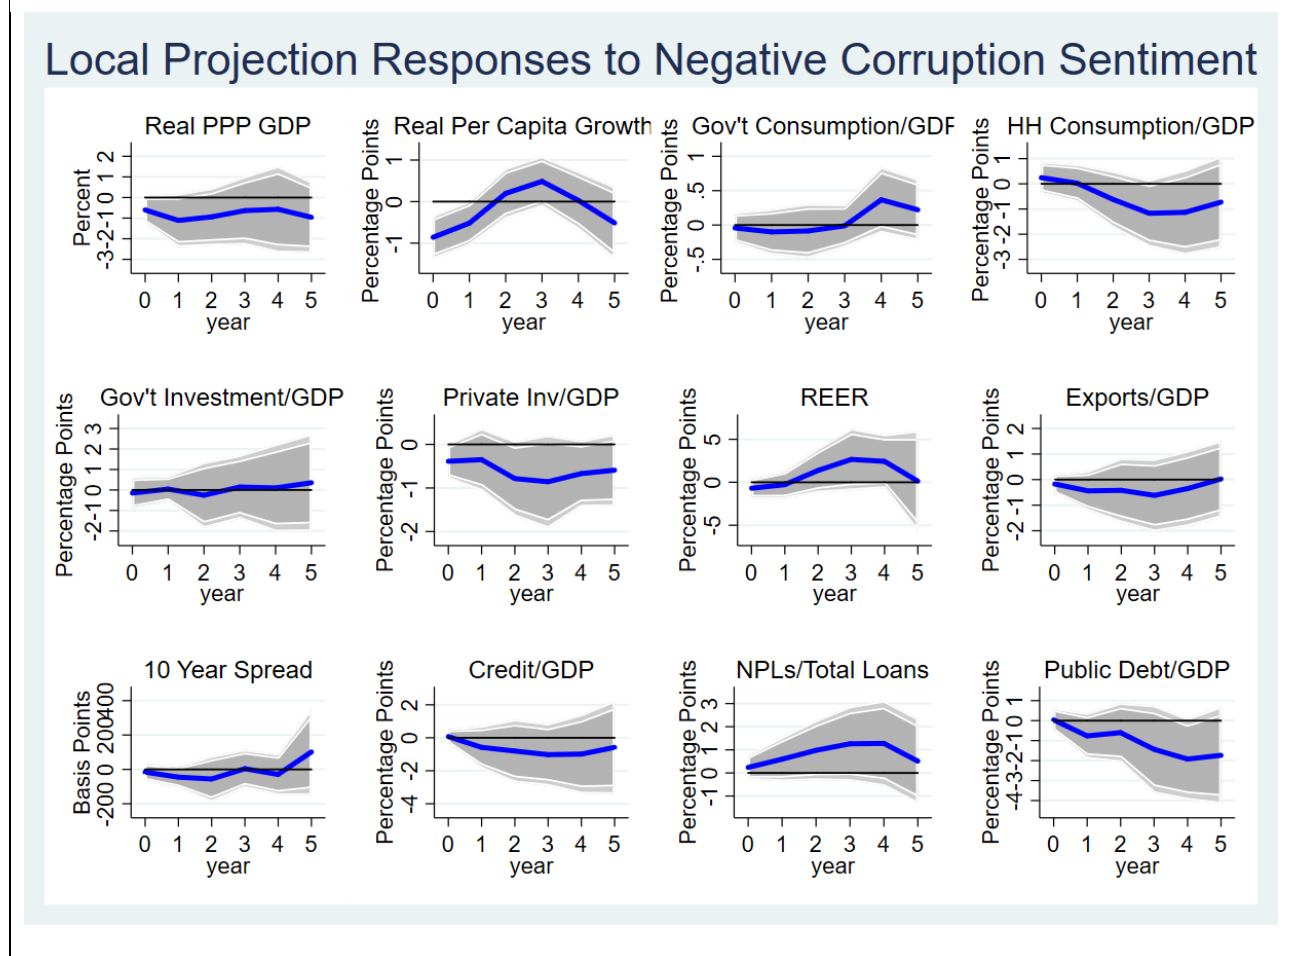
Local Projection Responses to Negative Corruption Sentiment



Our results show that an increase in negative corruption sentiment is associated with a decline in growth in emerging and low-income countries. Across the full panel of countries, periods with negative corruption sentiment that are at least one standard deviation above countries' mean are associated with an average drop in GDP per capita growth of 0.65 percentage point in $t+1$ (up from a hit of almost 0.50 in year 0) and the effect dissipates after $t+1$ (Figure 6). Emerging market and low-income countries (EMLICs) see larger effects compared to advanced economies, and in the latter the impact of negative sentiment is not significant. EMLICs also experience related falls in household consumption and private investment in percent of GDP. The magnitude of the full panel results and outsized growth effects on EMLICs are consistent with the findings of Hlatshwayo, et al., 2018, who find that shocks to their measure calibrated to be double the size as those considered by this paper (two vs. one standard deviation) are associated with lower real GDP per capital growth of two percentage points or roughly double the size of our findings. Comparison of the magnitudes to other papers in the literature is difficult as most papers employ continuous measures of corruption in their analyses rather than our event-based approach; however, Ugur's (2014) meta-analysis of 29 studies on corruption's effects on per capita growth similarly finds that corruption has a significant effect on low income countries while tending to have an insignificant effect in papers that consider a broader set of income groups.

Our results also show that the economic growth recovers in the medium run after the negative corruption sentiment shock hits. It could reflect the impact from countries adopting forceful anti-corruption measures to tackle the negative corruption sentiment shocks. For example, when the corruption case is identified by the news, government may conduct a legal investigation and put in place structures to avoid future scandals to address it. If successful, such actions could improve business climate and rebuild confidence, and thus boost economic growth in the medium run. However, if the anticorruption measures are not successful, the economic recovery could be prolonged. The economic recovery is shown to be stronger in advanced economies than in EMLICs (Figure 7 and Figure 8).

Figure 8. Cumulative Macroeconomic Impacts of Negative Corruption Sentiment, Emerging Market & Low-Income Countries



V. CONCLUSIONS

The paper constructs a sentiment-enhanced corruption perception index based on FT news articles over 2005–18 using sentiment analysis. The SECPI complements and correlates with existing popular corruption perception measures. It is highly sensitive to the current events (e.g., corruption investigation and election), and can be made available at a higher frequency with relatively low costs. The paper also shows that the SECPI correlates with variables related to governance, business climate, and poverty. The paper reinforces the argument that corruption perception, as captured by our index, has sizable negative impact on growth,

especially in emerging market and low-income countries. The high frequency SECPI provides a useful tool for investigating the short-term impacts of the corruption on economies. Looking ahead, the same techniques used in the paper can be applied to extract a wider range of opinions on corruption from diverse data sources, including local newspapers and social media posts.

References

- Achury, Carolina, Christos Koulovatianos, and John Tsoukalas. Political economics of external sovereign defaults. No. 508. Center for Financial Studies (CFS), 2015.
- Agarwal, Apoorv, Fadi Biadisy, and Kathleen Mckeown. "Contextual phrase-level polarity analysis using lexical affect scoring and syntactic n-grams." In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pp. 24-32. 2009.
- Aghion, Philippe, Ufuk Akcigit, Julia Cagé, and William R. Kerr. "Taxation, corruption, and growth." *European Economic Review* 86 (2016): 24-51.
- Akitoby, Bernardin, and Thomas Stratmann. "The value of institutions for financial markets: evidence from emerging markets." *Review of World Economics* 146, no. 4 (2010): 781-797.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. "Inference on treatment effects after selection among high-dimensional controls." *The Review of Economic Studies* 81, no. 2 (2014): 608-650.
- Burscher, Bjorn, Rens Vliegenthart, and Claes H. de Vreese. "Frames beyond words: Applying cluster and sentiment analysis to news coverage of the nuclear power issue." *Social Science Computer Review* 34, no. 5 (2016): 530-545.
- Cambria, Erik, Björn Schuller, Yunqing Xia, and Catherine Havasi. "New avenues in opinion mining and sentiment analysis." *IEEE Intelligent systems* 28, no. 2 (2013): 15-21.
- Campos, Nauro F., Ralitza D. Dimova, and Ahmad Saleh. "Whither corruption? A quantitative survey of the literature on corruption and growth." Discussion Paper No. 5334 (Bonn: Institute for the Study of Labor), 2010.
- Cavoli, Tony, and John K. Wilson. "Corruption, central bank (in) dependence and optimal monetary policy in a simple model." *Journal of Policy Modeling* 37, no. 3 (2015): 501-509.
- Deng, Shuyuan, Atish P. Sinha, and Huimin Zhao. "Adapting sentiment lexicons to domain-specific social media texts." *Decision Support Systems* 94 (2017): 65-76.
- Eshbaugh-Soha, Matthew. "The tone of local presidential news coverage." *Political Communication* 27, no. 2 (2010): 121-140.
- Fayad, Ghada, Chengyu Huang, Yoko Shibuya, and Peng Zhao. "How do member countries receive IMF policy advice: results from a state-of-the-art sentiment index." (2020).
- Fazekas, Mihály, and Gábor Kocsis. "Uncovering high-level corruption: cross-national objective corruption risk indicators using public procurement data." *British Journal of Political Science* 50.1 (2020): 155-164.

- González-Bailón, Sandra, and Georgios Paltoglou. "Signals of public opinion in online communication: A comparison of methods and data sources." *The ANNALS of the American Academy of Political and Social Science* 659, no. 1 (2015): 95-107.
- Grimmer, Justin, and Brandon M. Stewart. "Text as data: The promise and pitfalls of automatic content analysis methods for political texts." *Political analysis* 21.3 (2013): 267-297.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. "Cultural biases in Economic Exchange?" *The Quarterly Journal of Economics*, Vol. 124 (3) (2009): 1095-1131.
- Gupta, Sanjeev, Hamid Davoodi, and Rosa Alonso-Terme. "Does corruption affect income inequality and poverty?" *Economics of Governance* 3, no. 1 (2002): 23-45.
- Hlatshwayo, Sandile, Anne Oeking, Manuk Ghazanchyan, David Corvino, Ananya Shukla, and Lamin Y. Leigh. *The Measurement and Macro-Relevance of Corruption: A Big Data Approach*. No. 2018/195. International Monetary Fund, 2018.
- Hopkins, Daniel J., and Gary King. "A method of automated nonparametric content analysis for social science." *American Journal of Political Science* 54, no. 1 (2010): 229-247.
- Hu, Minqing, and Bing Liu. "Mining and summarizing customer reviews." In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 168-177. 2004.
- Huang, Haizhou, and Shang-Jin Wei. "Monetary policies for developing countries: The role of institutional quality." *Journal of International Economics* 70, no. 1 (2006): 239-252.
- Kaufmann, Daniel. "Can corruption adversely affect public finances in industrialized countries." *Brookings Institution*, April 19 (2010).
- Kaufmann, Daniel, Aart Kraay and Massimo Mastruzzi (2010). "The Worldwide Governance Indicators: Methodology and Analytical Issues". World Bank Policy Research Working Paper No. 5430.
- Krosnick, Jon A. "Survey research." *Annual review of psychology* 50.1 (1999): 537-567.
- Liu, Bing. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge university press, 2020.
- Loughran, Tim, and Bill McDonald. "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks." *The Journal of Finance* 66, no. 1 (2011): 35-65.

Manning, Christopher, and Hinrich Schutze. *Foundations of statistical natural language processing*. MIT press, 1999.

Mithe, Ravina, Supriya Indalkar, and Nilam Divekar. "Optical character recognition." *International journal of recent technology and engineering (IJRTE)* 2, no. 1 (2013): 72-75.

Monroe, Burt L., Michael P. Colaresi, and Kevin M. Quinn. "Fightin'words: Lexical feature selection and evaluation for identifying the content of political conflict." *Political Analysis* 16, no. 4 (2008): 372-403.

Nguyen, Thien Hai, and Kiyoaki Shirai. "Topic modeling based sentiment analysis on social media for stock market prediction." In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1354-1364. 2015.

Nguyen, Thien Hai, Kiyoaki Shirai, and Julien Velcin. "Sentiment analysis on social media for stock movement prediction." *Expert Systems with Applications* 42, no. 24 (2015): 9603–9611.

Oliveira, Nuno, Paulo Cortez, and Nelson Areal. "The impact of microblogging data for stock market prediction: Using Twitter to predict returns, volatility, trading volume and survey sentiment indices." *Expert Systems with Applications* 73 (2017): 125-144.

Pang, Bo, and Lillian Lee. "A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts." *arXiv preprint cs/0409058* (2004).

Rajkumar, Andrew Sunil, and Vinaya Swaroop. "Public spending and outcomes: Does governance matter?" *Journal of Development Economics* 86, no. 1 (2008): 96-111.

Sapienza, Paola, and Luigi Zingales. "A trust crisis." *International Review of Finance* 12, no. 2 (2012): 123-131.

Soroka, Stuart N. "The gatekeeping function: Distributions of information in media and the real world." *The Journal of Politics* 74, no. 2 (2012): 514-528.

Soroka, Stuart N., Dominik A. Stecula, and Christopher Wlezien. "It's (change in) the (future) economy, stupid: economic indicators, the media, and public opinion." *American Journal of Political Science* 59, no. 2 (2015): 457-474.

Soroka, Stuart, Lori Young, and Meital Balmas. "Bad news or mad news? Sentiment scoring of negativity, fear, and anger in news content." *The ANNALS of the American Academy of Political and Social Science* 659, no. 1 (2015): 108-121.

- Tanzi, Vito, and Hamid Davoodi. "Corruption, public investment, and growth." In *The welfare state, public investment, and growth*, pp. 41-60. Springer, Tokyo, 1998.
- Turney, Peter D. "Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews." *arXiv preprint cs/0212032* (2002).
- Ugur, Mehmet. "Corruption's direct effects on per - capita income growth: a meta - analysis." *Journal of Economic Surveys* 28, no. 3 (2014): 472-490.
- Van Atteveldt, Wouter, Jan Kleinnijenhuis, Nel Ruigrok, and Stefan Schlobach. "Good news or bad news? Conducting sentiment analysis on Dutch text to distinguish between positive and negative relations." *Journal of Information Technology & Politics* 5, no. 1 (2008): 73-94.
- Wilson, Theresa, Janyce Wiebe, and Paul Hoffmann. "Recognizing contextual polarity in phrase-level sentiment analysis." In *Proceedings of human language technology conference and conference on empirical methods in natural language processing*, pp. 347-354. 2005.
- Young, Lori, and Stuart Soroka. "Affective news: The automated coding of sentiment in political texts." *Political Communication* 29, no. 2 (2012): 205-231

Annex I. List of Variables for Local Projections

The sources of the data are in parentheses.

Real PPP GDP (IMF Vulnerability Exercise Fiscal Model Database)

Real Per Capita Growth (World Economic Outlook)

Current Account/GDP (IMF Vulnerability Exercise Fiscal Model Database)

General Government Expenditure/GDP (IMF Vulnerability Exercise Fiscal Model Database)

General Government Interest Expenditure/GDP (IMF Vulnerability Exercise Fiscal Model Database)

General Government Revenues/ GDP (IMF Vulnerability Exercise Fiscal Model Database)

Government Consumption/ GDP (IMF Vulnerability Exercise Fiscal Model Database)

Government Investment/ GDP (IMF Vulnerability Exercise Fiscal Model Database)

Net FDI/ GDP (IMF Vulnerability Exercise Fiscal Model Database)

Exports/ GDP (IMF Vulnerability Exercise Fiscal Model Database)

Imports/GDP (IMF Vulnerability Exercise Fiscal Model Database)

Trade Openness 10Y Average (IMF Vulnerability Exercise Fiscal Model Database)

Trading Partner Growth (IMF Vulnerability Exercise Fiscal Model Database)

Partner Import Demand (IMF Vulnerability Exercise Fiscal Model Database)

Real Effective Exchange Rate (IMF Vulnerability Exercise Fiscal Model Database)

Exchange Rate Depreciation (IMF Vulnerability Exercise Fiscal Model Database)

Exchange Rate, End of Period (IMF Vulnerability Exercise Fiscal Model Database)

PPP Exchange Rate (IMF Vulnerability Exercise Fiscal Model Database)

Terms of Trade Inflation (IMF Vulnerability Exercise Fiscal Model Database)

Inflation Rate (World Economic Outlook)

Reserves Growth, LCU (IMF Vulnerability Exercise Fiscal Model Database)

Remittances/GDP (IMF Vulnerability Exercise Fiscal Model Database)

Public External Debt/GDP (IMF Vulnerability Exercise Fiscal Model Database)

Public External Debt/Exports (IMF Vulnerability Exercise Fiscal Model Database)

Public Debt/ Revenues (IMF Vulnerability Exercise Fiscal Model Database)

Public Debt/GDP, 5y Change (IMF Vulnerability Exercise Fiscal Model Database)

Total Debt/GDP (IMF Vulnerability Exercise Fiscal Model Database)

Growth Deviation from 5y Avg (IMF Vulnerability Exercise Fiscal Model Database)

RGDP Growth Volatility (IMF Vulnerability Exercise Fiscal Model Database)
 Terms of Trade Volatility (IMF Vulnerability Exercise Fiscal Model Database)
 Exchange Rate Volatility (IMF Vulnerability Exercise Fiscal Model Database)
 Inflation Volatility (IMF Vulnerability Exercise Fiscal Model Database)
 Population 10Y Growth (IMF Vulnerability Exercise Fiscal Model Database)
 Population 5Y Growth (IMF Vulnerability Exercise Fiscal Model Database)
 Population (IMF Vulnerability Exercise Fiscal Model Database)
 Nat. Disaster Impact Growth (IMF Vulnerability Exercise Fiscal Model Database)
 Polity Score (Center for Systemic Peace)
 Checks & Balances Index (IMF Vulnerability Exercise Fiscal Model Database)
 Bureaucracy Quality (ICRG)¹⁹
 Democratic Accountability (ICRG)
 Ethnic Tensions (ICRG)
 Investment Prof (ICRG)
 Law & Order (ICRG)
 Military in Politics (ICRG)
 Religious Tensions (ICRG)
 Socioeconomic Conditions (ICRG)
 Monetary Union Membership, Dummy (IMF Vulnerability Exercise Fiscal Model Database)

¹⁹ Country Data Online, The PRS Group, Inc., www.prsgroup.com