Measuring Geopolitical Risk*

Dario Caldara† Matteo Iacoviello‡

September 27, 2017

Abstract

We present a monthly index of geopolitical risk (GPR index) based on a tally of newspaper stories that contain terms related to geopolitical tensions, and examine its evolution and effects since 1985. The GPR index spikes around the Gulf War, after 9/11, during the 2003 Iraq invasion, during the 2014 Russia-Ukraine crisis, and after the Paris terrorist attacks. Overall, the GPR index serves as a good proxy for movements in geopolitical risk over time. At the macro level, we find that: (1) Higher geopolitical risk leads to a decline in real activity and is associated with increases in the VIX and higher corporate credit spreads; (2) higher geopolitical risk leads to lower stock returns; (3) higher geopolitical risk leads capital to flow away from emerging economies and towards advanced economies.

PRELIMINARY

Latest version at https://www2.bc.edu/matteo-iacoviello/gpr_files/GPR_PAPER.pdf

*We thank Alessandra Bonfiglioli, Nick Bloom, Ricardo Correa, Nathan Converse, Chris Erceg, Bo Sun and seminar participants at the Federal Reserve Board, the Stanford Institute for Theoretical Economics, the Chengdu International Macro-Finance Conference, the CEF Conference in New York, the Central Bank of Chile, Pontificia Universidad Catolica de Chile, and Bocconi University. Joshua Herman, Lucas Husted, Andrew Kane, Aaron Markiewitz provided outstanding research assistance. All errors and omissions are our own responsibility. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

†Federal Reserve Board of Governors. Email: dario.caldara@frb.gov

‡Federal Reserve Board of Governors. Email: matteo.iacoviello@frb.gov

1 Introduction

Major geopolitical events are often regarded as setting the stage for radical changes in the business cycle and in financial markets, and geopolitical risks are often cited by central banks officials, the financial press, and business investors as key determinants of investment and spending decisions.[1] As a tell-tale sign of the importance of geopolitical risks, several companies profitably publish ratings, indexes, and intelligence reports containing detailed and colorful information on where it is safe, or dangerous, to do business. A Google search of the term "geopolitical risk" returned 229,000 hits as of September 2017, with the first 20 results containing the names of at least ten different companies, insurance agencies, and organizations that publish and sell indexes measuring geopolitical risk across countries, regions and over time.[2]

However, virtually all available indexes of geopolitical risk suffer from a variety of shortcomings that make them hardly usable for applied research. First, these indexes are often qualitative measures of geopolitical risk and are constructed in a subjective way, using a few colors, number scales from 1 to 5, or letter grading scales that are meant to tell the risky places and times apart from the safe ones. Second, these indexes either vary little over time, or they are available for only a short period. A notable example, the Doomsday Clock,[3] measures the countdown to a possible global catastrophe, with fewer minutes to midnight measuring higher risk, but the value of this index has changed only six times in the last twenty years. Third, when these indexes are quantitative, they are constructed by combining other variables which are meant to react to, rather than measure, geopolitical risks, such as gold, the VIX, the dollar index, and other financial market indicators.

In this paper we develop a monthly, quantitative index of global geopolitical risk –the GPR index– that aims at providing a complement to many available, qualitative indexes of geopolitical risk. Drawing on the methodology pioneered by Saiz and Simonsohn (2013) and Baker, Bloom, and Davis (2016), the index is constructed by counting the occurrence of words related to geopolitical tensions in leading international newspapers. The index is plotted in Figure 1. It spikes during the Gulf War and during the

[1] See for instance Greenspan (2002) and Carney (2016). In the 2017 Bank of England systemic risk survey, 61 percent of market participants cited geopolitical risk as one of the risks to the UK financial system (see http://www.bankofengland.co.uk/publications/Documents/other/srs/srs2017h1.pdf). In a Wells Fargo/Gallup May 2017 survey of more than 1,000 investors, 75 percent were very or somewhat worried about the impact of the various military and diplomatic conflicts happening around the world, ranking geopolitical risk ahead of political and economic uncertainty (see http://www.businesswire.com/news/home/20170613005348/en/Wells-FargoGallup-Survey-Geopolitical-Risks-Greater-Threat).

[2] For example, one can find, among the results of the search, the following companies providing businesses with the intelligence to make the right decisions in world of uncertainties: Marsh-McLennan, Control Risks Online, Zurich Insurance, Cambridge Econometrics, U.S. Energy Stream, Aon plc, Verisk Maplecroft, CSO Online, Euler Hermes, Risk Advisory, Strategic Risk.

[3] The Bulletin of the Atomic Scientists webpage has more details about the Doomsday Clock.

2003 invasion of Iraq.[4] It peaks around 9/11. More recently, it spikes up both during the Ukraine/Russia crisis, and around the Paris terrorist attacks, and during the escalation of the Syrian conflict in April 2017. An historical version of the index —dating back to 1899 and spanning over 120 years, and plotted in Figure 2— reaches its highest values at the beginning at the two World Wars, as well as at the onset of the U.S. involvement in them.

We study the relationship between the GPR index and business cycles and financial market movements. We find that a typical spike in the GPR —sized to match the jump after major episodes of higher geopolitical risk— depresses U.S. and advanced economies' industrial production by about 1 percent, and leads to a tightening in various measures of financial conditions. Most importantly, we find that what matters about geopolitical events is the uncertainty surrounding them: a version of the index that separates "geopolitical threats" from "geopolitical acts" shows that threats produce much more powerful effects than acts.

An important upshot of the index is that it quantifies episodes of geopolitical tensions that could hardly be predicted ahead of time and that are, by and large, exogenous to world economic conditions. The index is available at https://www2.bc.edu/matteo-iacoviello/gpr.htm.

Section 2 describes the construction of the geopolitical risk index. Section 3 presents the audit of the index, and compares the index to often–used proxies for movements in geopolitical risk. Section 4 studies the macroeconomic effects –on activity, capital flows and stock returns– of changes in geopolitical risk. Section 5 concludes.

2 Construction of the Geopolitical Risk Index

2.1 Basics

The goal of this paper is to understand the macroeconomic and financial consequences of major global political events and risks such as wars, terrorism, and regional political tensions. In recent years, these types of events and risks have been described by the press as "geopolitical," which motivates our decision to call the index the geopolitical risk (GPR) index.

The journalistic use of the term "geopolitical risk" does not meet any scholarly definition. In fact, a formal definition of geopolitical risk is not available, but rather seems to be an extension of "political risk" to major global events. Political risk is itself a fuzzy concept, and its definition has substantially evolved over history, with one definition referring to "political risk" as the risk an investment's returns

[4] Throughout the paper, we refer to historical events by adopting the naming convention followed by the Wikipedia entry for that event.

could suffer as a result of political changes or instability in a country (Sottilotta, 2013).

We define geopolitical risk as the risk associated with wars, terrorist acts, and tensions between states that affect the normal course of domestic politics and international relations. Both wars and terrorist events aim at achieving political or ideological goals through the use of violence, although they differ in scope, methods, and type of actors. Political tensions, even when initially confined inside one country, can also create geopolitical instability, to the extent that they increase the probability of future conflicts. The index aims at measuring, throughout history, these risks and associated instability. In practice, we construct the GPR index by counting the occurrence of words related to geopolitical tensions in leading newspapers. In particular, the GPR index reflects automated text-search results of the electronic archives of 11 national and international newspapers: The Boston Globe, Chicago Tribune, The Daily Telegraph, Financial Times, The Globe and Mail, The Guardian, Los Angeles Times, The New York Times, The Times, The Wall Street Journal, and The Washington Post. We calculate the index by counting the number of articles related to geopolitical risk in each newspaper for each month (as a share of the total number of news articles). The index is then normalized to average a value of 100 in the 2000-2009 decade,[5] so that a reading of 200, for instance, indicates that newspaper mentions of geopolitical risk in that month were twice as frequent as they were during the 2000s.[6] The search identifies articles containing references to any of the six categories of words reported in Table 1. We arrive at this set of words after a pilot audit study of articles mentioning geopolitical tensions – see the description of the audit in the next section –, and after isolating the most common unigrams and bigrams in geopolitics textbooks.[7] In doing so, we verify that the defining elements of geopolitics concern territory, countries, nations, and leadership, and that defining elements of geopolitical risk center around the risk of wars and terrorism. Building on this characterization, we construct a set of search terms that give us an index that is robust, sensible, and easily interpretable. As shown in Table 1, the first four categories of words are related to geopolitical threats and tensions, while the last two categories are related to geopolitical events and acts (of negative nature: the start of a war is part of the

[5] A monthly reading of 100 for the index corresponds to about 350 articles per month containing terms related to geopolitical risk.

[6] The companion dataset reports the total number of articles across newspapers in each month. The number of articles is trendless around 70,000 since the early 1990s. For one representative newspaper, this corresponds to about 200 articles per day, of which about one in 300 mentions, on average, terms related to geopolitical risk. As a comparison benchmark, one in 500 articles mentions the Beatles, and one in 300 articles mentions the Federal Reserve.

[7] For instance, the book "Introduction to Geopolitics" (Flint, 2016) contains 48,759 bigrams, of which the most common ones are "geopolit code", "world leader", "unit[ed] states", "world leadership", "war [on] terror", "modelski model", "geopolit agent", "cold war", "soviet union", "world war", and "foreign polic[y]". The book "The Geopolitics Reader" (Dalby, Routledge, and Tuathail, 2003), which is a compendium of 39 geopolitics essays written by different authors, contains 91,210 bigrams, of which the most common ones are "unit[ed] states", "cold war", "foreign polic[y]", "nation secur[ity]", "world war", "world order", "nation[al] state", "gulf war", "war II", and "nuclear weapon".

index, while the end of a war is not).[8] Group 1 includes search terms associated with explicit mentions of geopolitical risk, as well as mentions of military-related tensions involving large regions of the world and a U.S. involvement. Group 2 includes search terms directly related to nuclear tensions. Groups 3 and 4 include mentions related to war threats and terrorist threats, respectively. Finally, Groups 5 and 6 aim at capturing press coverage of actual geopolitical events (as opposed to just risks) which can be reasonably expected to lead to increases in geopolitical uncertainty, such as terrorist acts or the beginning of a war. We do this for two reasons. First, actual "negative"geopolitical events are often associated with higher risk. Hence, searching directly for events rather than risk can help obtaining a more precise identification of the timing of some risk-inducing shocks. Second, in assessing the impact of geopolitical risk, we want to control for the direct impact that the event itself might have. For this reason, in Section 2.3 we create two sub-indexes that measure threats and acts separately.

Figure 1 presents the benchmark raw index. The index is characterized by several spikes corresponding to key geopolitical events which led to higher tensions. The first spike is recorded in April 1986 and it corresponds to the US bombing of Lybia and the terrorist escalation that led to it. The second spike happens during the Kuwait invasion and the subsequent Gulf War. The index spikes again at the beginning of 1998, during a period of escalating tensions between the United States and Iraq. It then stays low until 9/11, when it spikes above 300. The index reaches its maximum during the 2003 invasion of Iraq. Since 2003, the index rises in correspondence of major terrorist events in Europe, like the March 2004 Madrid bombing, the July 2005 London bombing, and the November 2015 Paris terror attacks. The index also rises in 2014, amidst rising geopolitical tensions in Ukraine and Iraq, as well as during the rise of ISIS.

Figure 3 elaborates further on the contribution of each component to the index. Nuclear threats are disproportionately more important prior to the end of the Cold war, and gradually subside after 1989. By contrast, terror threats trend higher over the sample period, spiking after 9/11, and remaining at elevated levels ever since. Interestingly, while "war threats" and "war acts" appear to move somewhat in sync throughout the sample, mentions of "terrorist threats" seem to rise in proportion to mentions of actual "terrorist acts" since 9/11.

[8] We refrain from including in the search words containing proper nouns, except for "United States" and the names of the four largest continents. We do so because we want to guard our search against the possibility that terms related to geopolitical risks may cease to be such for a limited period of time. For instance, although the name "Adolf Hitler" may have correctly identified geopolitical risks in the 1930s, it is far more often associated today to books, movies and historical accounts of the World War II, and is best left out of the searches. The same is often true of many political leaders or figures once they die, or are jailed, or retire from office.

2.2 What Are We Measuring?

A crucial aspect that concerns the measurement of geopolitical risk is that exposure to such risk varies both geographically and by sector of the economy. The index captures geopolitical events and risks as perceived and chronicled in the press in English-speaking countries, and in particular in the United States: to construct the index we use eight American, two British, and one Canadian newspaper. Moreover, one of the search categories measures bilateral regional tensions between the United States and other regions of the world.[9]

At the same time, geopolitical events that involve the United States and their interests abroad have global implications. Furthermore, the newspapers we search have wide geographical coverage. Hence, the index should be also a good proxy for global geopolitical risk.

Additionally, the use of press coverage has the potential to induce fluctuations in, or exceedingly high levels of, the GPR index even if the underlying geopolitical risk factors remain constant due to either changes in geopolitical-related risk aversion of the public or to state-dependent bias in news coverage. For example, the high levels in the years following 9/11 (constructed from this U.S.–centric selection of newspapers) may reflect more public fear towards geopolitical tensions than actual risk. We do not attempt at separating the actual underlying risk factors from the public or from the media perception of the risk itself, as there are reasons to believe that the two are very highly correlated. To formally test this hypothesis, we purge our geopolitical risk index from any fluctuations in geopolitical risk that are driven by event that are exogenous to geopolitical risk and are likely to receive large media attention. Following Jetter (2017), we use natural disasters as an instrument that exogenously diminishes media coverage of other events, including geopolitical events. If underlying geopolitical risk is constant but moves only in reaction to media slant towards geopolitical risk, one should expect that natural disasters can explain geopolitical risk, and that the component of geopolitical risk that is explained by natural disasters is a measure of the media attention towards geopolitical risk. In Figure A.3 in the Appendix, we compare the GPR index with a news-based index of natural disasters, constructed using media mentions of words related to natural disasters (such as tsunamis, hurricanes, and earthquakes). The correlation between geopolitical risk and natural disasters is -0.07, and not statistically different from zero at conventional significance levels, thus suggesting that mentions of geopolitical risks are not systematically varying in response to other exogenous events that might divert media attention from chatter about geopolitical risk.

[9] In Caldara, Iacoviello, and Markiewitz (2017), we construct a country-specific geopolitical risk index – again, as perceived by the same English–speaking press for 16 emerging countries, drawing on mentions of geopolitical tensions in each country.

2.3 Geopolitical Threats vs. Geopolitical Acts

Figure 4 decomposes the index in two subindexes: the GPT index (GPR Threats) and the GPA index (GPR Acts). The GPR index captures a convolution of shocks to first and higher order moments of the distribution of geopolitical events. Spikes in risk often coincide with the realization of big events. We attempt to disentangle the direct effect of geopolitical events from the effect of pure geopolitical risks by constructing two indexes: Geopolitical threats (GPT): Search categories 1 to 4; Geopolitical acts (GPA): Search categories 5 and 6.

2.4 Long-span Historical GPR Index

Using a methodology similar to that of the benchmark index, we construct a long-span monthly GPR index dating back to 1900. We restrict the newspapers' coverage to only three journals for which we can retrieve data all the way back to the beginning of the 20th century, namely the New York Times, Chicago Tribune, and the Washington Post. Aside from using a smaller set of newspapers and a slightly larger set of synonyms (mostly added to capture shifts in the usage of particular words), the long-span index closely mimics the benchmark index for the period in which the coverage overlaps–their monthly correlation is 0.95–.[10]

Figure 2 displays the monthly historical index. Even here, every major spike in the index can be associated with episodes of rising geopolitical tensions. The index stays elevated during the two world wars, and peaks at the onset of both of them. Figure 5 further illustrates how the components in the measure of historical geopolitical risk have evolved over the last 115 years. Early in the sample, the index rises and stays high during World War I and World War II, and phrases directly related to the conflict itself dominate the index. The index stays at high levels between the 1950s and the 1980s, at a time when the threat of a nuclear war and rising geopolitical tensions between countries become more prevalent than wars themselves. Since the 2000s, terrorism events have come to dominate the index, alongside rising bilateral tensions among countries. Indeed, the index reaches the highest values at the start of World War I, the start of World War II, and around 9/11.

[10] We add to the search category "War Threats": {(war OR military) N/3 (crisis OR uncertain*)}, and {("war effort*" OR "military effort*") AND (risk* OR threat* OR fear*)}. We add to the search category "War Acts": {"state of war"}, and {"declaration of war"}.

3 Auditing the Index

The construction of any news-based index raises obvious concerns about accuracy and bias. In this section we explain how we address these concerns. To preview the results, human reading of 11,000 newspapers articles as well as comparison to external proxies confirms that the index accurately captures movements in geopolitical risk.

3.1 Human Reading of the Articles

We read and code articles drawn the newspapers in our sample from 1985 through 2016.[11] Our goal is to ensure that the index is an accurate, reliable and stable measure of geopolitical concerns. For this to happen, the search terms must be more likely to be employed when geopolitical concerns materialize than when they do not, and must be sufficiently varied to minimize sampling error. Importantly, the search terms do not need to be detailed enough to capture all articles related to geopolitical risk, but should create a proxy for underlying movements in geopolitical risk that maximizes the correlation between the index and rising geopolitical tensions, that is the phenomenon we want to capture.

The universe of newspaper articles used to construct the index contains about 65,000 news articles each month (sample $\mathcal{U}$), of which only a small fraction contains references to geopolitical issues.[12] We narrow this list down sampling from this universe a subset of articles that are more likely to be classified as GPR^H=1 by our computer algorithm. This subset (sample $\mathcal{E}$) contains about 8,000 articles per month –that is, 15% of the articles in sample $\mathcal{U}$–, and is constructed by extracting articles that contain any of the following words: *geopolitics*, *war*, *military*, *terrorism/t*. While the first word enters the universe almost by definition, the other three have come to dominate the language and the practice of geopolitics over the years. Wars, military actions and terrorism all represent violent means for states and organizations to claim and control territories, to achieve political ends, and to extend the boundaries of their influence.[13] The goal of our audit proceeds is twofold. Our first goal is to read and analyze the content of newspaper articles with the intent of coming up with search terms that allow us to selecting articles that mention geopolitical risk, and to exclude articles that do not not. Our second goal is to

[11] As in Baker, Bloom, and Davis (2016), the construction of the indexes only requires count of articles that contain predetermined expressions. By contrast, we need the full-text articles to carry out the audit. We could not access the full-text articles for the *Financial Times*, which limits the audit study to 10 out of 11 newspapers.

[12] This number corresponds to about 200 articles a day, on average, for each of the 11 newspapers in the sample.

[13] The choice of words is also supported by an analysis of the most common unigrams found in textbooks and books on the subject of geopolitics. For instance, the most common word roots in Colin Flint's textbook "Introduction to Geopolitics" (Flint, 2016) are "geopolit", "war", "nation", "terror", "polit", and "countri", and "global". Similarly, the most common word roots in Samuel Huntington's classic book on the 'Clash of Civilizations '(Huntington, 1997) are "civil", "war", "cultur", "polit", "power", "econom", "societi", "conflict".

perform a time series comparison of our (computer-generated) GPR index with an alternative index which makes no use of our search terms and is constructed reading and coding, using detailed guidelines and human judgment, month after month, a large number of articles.

Our initial audit proceeds as follows. We select 50 months at random from the universe of all months that contribute to the index, and for each month we select 50 articles at random from the approximately 8,000 articles in $\mathcal{E}$ corresponding to that month. Together with a team of research assistants, and aided by a detailed audit guide,[14] we read and code each of these 2,500 articles as $\text{GPR}^H = 1$ if they mention high or rising geopolitical tensions, and $\text{GPR}^H = 0$ otherwise. Based on this analysis, about half of the articles in $\mathcal{E}$ discuss high or rising geopolitical risks, with the remaining half percent being false positives. Additionally, the fraction of false positives is very volatile, with a monthly standard deviation of 17%, thus indicating that a very broad search is likely to be contaminated by high noise-to-signal ratio. For instance, the fraction of false positives averages around 80 percent in the months after the end of the Gulf War, when newspaper coverage of the Gulf War is very extensive, but a substantial majority of articles cites declining tensions, peace initiatives after Saddam Hussein's withdrawal from Kuwait, and the importance of the U.N. mandate to maintain peace and stability in the region. Figure A.5 in the Appendix compares the benchmark index with the naive GP index constructed counting the relative frequency of all articles mentioning any of the four keywords that enter the set $\mathcal{E}$, as well as with a WARMINUSPEACE index subtracting the word *peace* from the occurrence of the words *war* and *terrorism*.

Based on the results of this initial audit, we scan the full text of the abstracts of the 2,500 articles sampled from set $\mathcal{E}$ in order to identify differences between those that are coded $\text{GPR}^H = 1$ –set $\mathcal{E}^1$– and those that are coded as $\text{GPR}^H = 0$ –set $\mathcal{E}^0$–. Specifically, we use text analytics techniques to identify the bigrams that are more likely to show up in $\mathcal{E}^1$ than in $\mathcal{E}^0$. This process helps us to refine the set of words that are part of the benchmark GPR index, as discussed in the previous section. Specifically, we compare articles across sets $\mathcal{E}^1$ and $\mathcal{E}^0$ and note phrases (eventually listed in Table 1) that are far likely to identify articles belonging to set $\mathcal{E}^1$ than to set $\mathcal{E}^0$, as well as phrases that are often associated with false positives. For instance, most false positives occur at times of anniversaries (such as the centenary of the First World War in 2014), upon death of historical figures, at the time of books' publication, movies' releases and other art events that are connected to wars, terrorism, and other episodes of important historical relevance.[15] We repeat this process and refine the index until additional filters and search

[14] The audit guide is available at https://www2.bc.edu/matteo-iacoviello/gpr_files/audit_guidelines_GPR.pptx.

[15] The set of words that are highly likely to signal false positives are "civil war", "human rights", the word "war" in close proximity of the word "end" (end N/2 war), the phrase "air force", the words "movie", "film", "museum", "anniversary", "memorial" and "art". We exclude from the banchmark index articles containing these phrases.

adjustments bring only modest improvements in the error rate. For instance, the first audit shows that the error rate is high towards the end of 1989, when many newspaper articles written during with the fall of communism in Eastern Europe and the rise to power of Gorbachev in the Soviet Union herald a new era of peace and reduced geopolitical threats. Although in some instances these articles feature words such as "risks" and "threats" in proximity of words such as "reduced" or "diminished," or phrases such as "the threat of war is over," we were unable to develop simple text filters that would significantly reduce the error rate.[16]

In the second audit, we read articles from the set $\mathcal{G}$ of articles that belong to the benchmark (computer–generated) GPR index. The set $\mathcal{G}$ contains about 280 articles per month, that is, about 0.4% of the articles in the sample $\mathcal{U}$. As before, we read 50 articles from 50 different months, for a total of $2,500$ articles, and classify them as either belonging to set $\mathcal{G}^1$, if they discuss high or rising geopolitical tensions, or as belonging to set $\mathcal{G}^0$, if they do not. About 87 percent of the articles in the sample $\mathcal{G}$ are articles that an intelligent human would classify as mentioning high or rising geopolitical risks. For the 50 months that we sample, the correlation between the human-audited GPR index –that is, the product of GPR index and the fraction of articles belonging to $\mathcal{G}^1$ in that month– and the benchmark GPR index is 0.98. The error rate –the fraction of articles belonging to set $\mathcal{G}^0$– is essentially uncorrelated with GPR itself as well as with other macroeconomic variables.

In the final audit, we check that our computer–automated GPR classifications track the actual time-series variation in the intensity of concerns about geopolitical risk, as judged by intelligent humans. To this end, we randomly sample and read 6,125 articles from the GP set (that is, about 50 articles per quarter on average), compute in each quarter the fraction of articles that we code as $\text{GPR}^H = 1$, multiply this fraction by the GP index, and compare the resulting series with our computer-generated GPR index. Figure 6 shows that our computer index lines up well with an index that could be constructed by intelligent humans. The correlation between the two series is 0.837, a value that is remarkably high when one takes into account the sampling uncertainty. Figure A.4 shows that a high correlation—0.86 when we aggregate the data at annual frequency, and 0.78 at quarterly frequency—when we repeat the same exercise using historical data, and a random sample of 7,416 articles, from 1899 through 2017.

3.2 Comparison with Broader and Narrower Measures

As additional checks, we consider a broad index, a narrow index and a simple index. These indexes –shown as shares of total articles by month, rather than indexed to 100– are presented in Figure 7. As the figure shows, the index is robust to the inclusion and exclusion of specific phrases and synonyms.

[16] See e.g. The Meaning of Berlin, by Flora Lewis, published in the New York Times on Nov 26, 1989.

The **broad index** combines the search terms in Table 1 with *bigrams* that a machine learning algorithm separately applied to sets $\mathcal{E}^1$ and $\mathcal{E}^0$ signals as very likely indicators of rising geopolitical tensions. Specifically, we run two separate, automated text analyses of all the articles coded in the audit as either $\text{GPRE}^H = 1$ or $\text{GPRE}^H = 0$, and identify the most recurrent bigrams in both sets. Using Bayes' rule, we compute the odds ratio of an article being coded as $\text{GPRE}^H = 1$ instead of $\text{GPRE}^H = 0$ given that it contains each bigram, and add the ten bigrams (out of the 50 most recurring ones) with the highest ratio to the list of search terms.[17] While the resulting search criteria double the number of articles mentioning geopolitical risks, the resulting index is broadly in line with the benchmark one (correlation 0.92), thus suggesting that the index is robust to the exclusion of expressions that are likely to be associated with rising geopolitical tensions. The broad index exhibits a slightly lower correlation than the benchmark index with the human index, 0.82 when the data are aggregated at quarterly frequency.

The **narrow index** excludes articles containing search terms which, in the human audit, were often correlated with false positives. As before, we identify bigrams that are more likely than not to appear in set of articles that contains false positives. We then exclude from the index all the articles containing any of these bigrams.[18] The resulting search reduces by approximately 15 percent the number of articles mentioning geopolitical risks, but the resulting index is virtually indistinguishable from the benchmark one (correlation 0.997), thus suggesting that the index is sufficiently robust to the inclusion of words that are likely to be associated with false positives. Like the broad index, the narrow index also exhibits a slightly lower correlation than the benchmark index with the human index, 0.76 when the data are aggregated at quarterly frequency.

Finally, we construct a simple index that economizes on search terms, and is close in spirit to the methodology of Baker, Bloom, and Davis (2016). The **simple index** is based on articles that, instead of

[17] Prior to calculating the bigrams, we filter out stopwords and proper names (including names of countries, cities, and political organizations). For each bigram, we can calculate the probability that the bigram signals a $\text{GPRE}^H = 1$ article using Bayes' rule. For instance, the bigram "military aid" appears 25 times in the $\text{GPRE}^H = 1$ set, and 2 times in the $\text{GPRE}^H = 0$ set. Given that the $\text{GPRE}^H = 1$ and the $\text{GPRE}^H = 0$ sets contain 101,174 and 74,554 bigrams respectively, the unconditional probability that an article containing "military aid" discusses rising geopolitical risks is 0.90. The search terms included in the broad index are "military presence", "military assistance", "military aid", "peace process", "international terror", "administration official", "nuclear weapon", "weapon mass", "mass destruction". The tenth keyword, "terrorist threat", is included but redundant because it is also part of the benchmark index. These bigrams were found to have odds of belonging to set $\mathcal{E}^1$ instead of set $\mathcal{E}^0$ of more than 70 percent.

[18] Our benchmark index already excludes expressions that in our first audit were often associated with false positives, like "civil war" and "human rights", as well as words like movie, museum, anniversary, memorial and art. The additional bigrams that exclude an article from the count in the narrow index are "air force", "death penalty", "national guard", "supreme court", "justice department", "enemy combatant", "military commission", "military tribunal", "military civilian", "military loss", "defense department", "chief of staff", "law enforcement", "war crime". These bigrams were found to have odds of belonging to set $\mathcal{E}^1$ instead of set $\mathcal{E}^0$ of less than 40 percent.

belonging to any of the six potential search categories, contain at least one word from each of two sets of terms: the set $\mathcal{S}_1$, including {war OR military OR terrorism OR geopolitical}, and the set $\mathcal{S}_2$, containing {risk* OR concern* OR tension* OR uncertaint* OR threat*}. The correlation between the benchmark index and the simple GPR index is sizable, at 0.89. Although the principle of parsimony would make this index appealing (in particular, this index has the same quarterly correlation with the human index, at 0.84), we prefer the benchmark index because, among other things, it showed a lower error rate in pilot audits, and because it affords a natural decomposition into several search subcategories that is not afforded by the simple index.

3.3 Comparison with Alternative Measures of Geopolitical Risks

Several studies have used various proxies to construct quantitative measures of risk that speak to concerns specific to wars or terrorist-related events. A widely used database is the ICB annual database, which provides detailed information on 447 major international political crises that occurred over the period 1918–2006. This database has been used in the political science literature, as well as in studies on war and economics. One example is the work by Berkman, Jacobsen, and Lee (2011), who use the ICB database to create a crisis index that adds up, for each year, the total number of crises. Not surprisingly, such index is a very crude measure of geopolitical risk. The top panel of Figure 8 compares the historical GPR index with the count of international **crises per month** according to the ICB database (to reduce high–frequency noise, both indexes are plotted at quarterly averages). It seems somewhat heroic to argue that the mere count of international crises correlates well with geopolitical instability. For instance, the value of the index is below its historical average before and during the Second World War, thus failing to pass even the most basic of smell tests.

The second panel of Figure 8 compares our index to two alternative indicators that offer a different perspective on the threats coming from geopolitical violence. It plots **deaths** due to terrorist attacks (1) in the world and (2) in the United States and Europe combined. The latter are likely to receive more press coverage in the English-speaking press.[19] Both series appear to be uninformative about overall movements in GPR, although, as expected, they spike together around 9/11. However, the somewhat elevated levels of the GPR index in 2015 and 2016 appear to reflect rise in the worldwide number of deaths due to terrorism, alongside the heightened media attention on the conflicts in the Middle East.

The third panel of Figure 8 compares the GPR index with the **National Security** component of

[19] The data on terrorism-related deaths largely exclude wars, but the distinction appears mostly semantic as the dividing line between wars and terrorism has been blurred at least since 9/11. The data are from the Global Terrorism Database (GTD), which is an open-source database including information on terrorist events around the world.

the Economic Policy Uncertainty Index by Baker, Bloom, and Davis (2016). Like our measure, the national security EPU spikes during the Gulf War, after 9/11, and during the Iraq war. However, the GPR index seems to better capture other spikes in geopolitical uncertainty which are missed by the national security EPU. The correlation between the two measures is 0.69, a high but not extreme value. We believe that there is no reason to expect that the correlation between the two measures is higher than that. After all, the national security component of the EPU still captures uncertainty about policy responses to geopolitical events, which is a different concept from the uncertainty generated by the geopolitical events themselves.[20].

Finally, the last panel of Figure 8 compares the GPR index with an outside measure of political risk related to wars, the **U.S. External Conflict Rating (ECR)** constructed by the International Country Risk Guide (ICRG). The ratings constructed by the ICRG are largely subjective, as they are based on the insights of various analysts following developments in a particular country or region. The ECR measure moves only occasionally over the sample, changing on average once a year, with more pronounced movements and more frequent changes around 9/11, when both the GPR index and ICRG index spike. The correlation between the two series is 0.41.

3.4 Relation to Popular Measures of Economic Uncertainty

Figure 9 compares the benchmark index with other popular measures of uncertainty: VIX, gold, the EPU index of Baker, Bloom, and Davis (2016). As the Figure shows, there appears to be little comovement between our measure and these uncertainty measures.

In the top panel, the GPR index and the VIX share two common spikes in 1991, at the time of the Gulf War, and in 2001, after the 9/11 terrorist attacks. However, in both cases the correlation seems to run from our index to the VIX, rather than the other way round. A similar set of considerations applies in the middle panel, when comparing our index to the EPU index. Finally, there appears to be little correlation between the GPR index and the price of gold.

3.5 Additional Robustness Check

For the sake of robustness, we check the conditions advocated by Saiz and Simonsohn (2013) that must hold to obtain useful internet-based proxies for variables, such as geopolitical risk over time, that are

[20] Another database is the Militarized Interstate Disputes Dataset, which records all instances of when one state threatened, displayed, or used force against another. This dataset is only available with several years delay and is only available annually. See http://cow.dss.ucdavis.edu/data-sets/MIDs. Other papers have performed news–searches looking for early warning signals for wars. See for instance Chadefaux (2014).

otherwise difficult to measure. These checks are described and discussed in the Appendix.

4 Macroeconomic Effects of Geopolitical Risk

4.1 Overview

In this section, we investigate the macroeconomic and financial effects of geopolitical risk.[21] This analysis can be seen as a complement and extension to the work of several authors that have tried to measure the macroeconomic consequences of wars, terrorist attacks, and other forms of collective violence. Using panel annual data on a large set of countries, Blomberg, Hess, and Orphanides (2004) find that both wars and terrorism have an economically significant negative effect on growth, but the effect of terrorism is smaller and less persistent than that associated with wars. However, their measure of terrorism is a simple dummy variable that presents a significant limitation to their analysis. Similar results are also found in the cross-country study of Tavares (2004), who finds – using annual data from 1987 through 2001 – that terrorist attacks have a modest negative effect on growth. Glick and Taylor (2010) find large and persistent impacts of wars on trade, national income, and global economic welfare.

Our index has several advantage over these measures: it is available at a monthly frequency; it can separately measure risks and acts; it is a continuous measure that gives higher weight to larger events; it directly measures the public perception of these events; it is available in real time.

4.2 Global Effects

We first provide an encompassing overview of the global macroeconomic effects of geopolitical risk by estimating a monthly VAR from 1985M1 through 2015M12 which includes a broad set of real and financial variables. The model consists of eight variables: (1) the GPR index; (2) the option-implied volatility on the S&P 100 stock futures index constructed by the Chicago Board of Option Exchange (VIX); (3) the log of US industrial production; (4) the log of advanced economies industrial production; (5) the log of emerging economies industrial production; and 6) the log of the Brent price of oil expressed in real terms dividing by the U.S. CPI index; (7) Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity; (8) the Standard and Poors 500 index, expressed in

[21] Several papers have looked at the theoretical channels by which terrorism and war may affect economic activity. Lenain, Bonturi, and Koen (2002) describe three channels in which terrorism can influence activity: shrinking insurance coverage stemming from the perception of greater risk, higher trade costs, and stepped-up security spending. Eckstein and Tsiddon (2004) present a model where terror endangers life such that the value of the future relative to the present is reduced. In their model, a rise in terror activity reduces investment, income, and consumption.

real terms.[22] The IP indexes, the real price of oil, and the stock market index are linearly detrended prior to estimation. All VAR models presented in the paper are estimated using Bayesian techniques. We impose a Minnesota prior on the reduced-form VAR parameters by using dummy observations as in Del Negro and Schorfheide (2011). The resulting specification, which includes a constant, is estimated using six lags of the endogenous variables.[23]

We identify the structural shocks by using a Cholesky decomposition of the covariance matrix of the VAR reduced-form residuals, ordering the GPR index first. The ordering implies that the GPR index reacts contemporaneously only to its own shock. Hence, any contemporaneous correlation between the macro variables and the GPR index reflects the effect of the GPR on the macro variables. While the characteristics of the GPR index discussed in Section 2 lend support to this assumption, later we explore robustness to alternative Cholesky orderings.

The solid lines in Figure 10 show the median impulse responses to an exogenous increase in the GPR index of 153 points, while the shaded areas represent the corresponding 68-percent pointwise credible bands. The shock equals the average change in the index following the nine episodes of largest increases in geopolitical risk.[24] On the real side, IP in various geographical regions responds quickly to the increase in geopolitical risk. The decline in the U.S. bottoms out at 0.7 percent after about 1 year, and converges slowly back to trend. The dynamic response of IP in advanced economies is broadly similar. IP in emerging economies contracts too, bottoming out at 0.8 percent below trend 8 months after the shock, somewhat faster than in the U.S. and in the Advanced Economies.

The rise in the GPR index induces a small and short-lived increase in overall global risk, as proxied by the VIX. By contrast, the response of the stock market is economically and statistically significant, dropping by 2.5 percent on impact and remaining below trend for over a year. Similarly, a rise in GPR induces a persistent increase in corporate credit spreads, which peak at 10 basis points 6 months after the shock. Finally, the increase in the GPR leads to a decrease in oil prices, which bottom out at 8 percent below trend after 6 months. This result stands in contrast with much of the conventionally held view that higher geopolitical risk drives up oil prices persistently, a view that might reflect a selective memory that confounds all geopolitical tensions with "oil supply shocks driven by geopolitical tensions in the Middle East." The evidence from the VAR supports the result that higher geopolitical risk creates

[22] Caldara, Cavallo, and Iacoviello (2016) provide details on the construction of the IP indexes for advanced and emerging economies.

[23] The vector of hyperparameters of the Minnesota prior is $\lambda = [1, 3, 1, 1, 1]$. We use the first year of the sample as a training sample for the Minnesota prior. All the results reported in the paper are based on 10,000 draws from the posterior distribution of the structural parameters, where the first 2,000 draws were used as a burn-in period.

[24] These episodes are U.S. bombing of Lybia in 1986, the Kuwait Invasion, the Gulf War, the 1998 Iraq Disarmament Crisis, 9/11, risk of Iraq invasion in 09/2002, the 2005 London bombings, the Russian annexation of Crimea, and the 2015 Paris terrorist attacks.

recessionary effects, which impacts negatively on oil demand and prices.

All told, the VAR evidence presented above is suggestive of sizeable recessionary effects of an increase in GPR.

4.3 Threats vs Acts

Next, we evaluate the difference between innovations in the two broad components of the GPR index, GPT – Threats Index – vs GPA – Acts Index. We do so by estimating the same VAR of before, replacing the GPR index with these two components. We are interested in the effects of shocks to these two subcomponents. A natural identification scheme is one where we assume that GPT innovations do not affect GPA contemporaneously, whereas GPA innovations affect GPT contemporaneously.

The impulse responses to the GPT and GPA shocks are shown in Figure 11. The upshot of these impulse responses is that they lend support to the old idea, attributed to London financier Nathan Rothschild, that one should buy on cannons and sell on trumpets. An innovation to GPT (the pure risk component of the index) leads to a small but short–lived rise in GPA, and produces large and protracted recessionary effects, as well as a decline in stock prices. By contrast, a positive shock to the GPA leads to a small but short–lived decline in economic activity, whereas the stock market rises sharply one month after the shock.

The different responses to "threat" and "act" shocks are broadly consistent with the notion that threat shocks depress asset prices and economic activity because they reveal information about new adverse negative events that are not fully diversifiable, thus raising risk premia. By contrast, the "act" component of the shock may lead to a resolution of the uncertainty around a particular set of events, as well as to a coordinated policy response that ends up giving protection on the worst possible outcomes.[25]

4.4 Robustness

The 9/11 terrorist attacks are the episode that induced the largest increase in the index (the index rose to its highest level with the Iraq war two years later, but from already elevated levels). Moreover, 9/11 is potentially associated with a structural break in the index, which remained persistently elevated after the terrorist attacks converging towards a higher level compared to the first part of the sample. The 9/11 terrorist attacks were not a typical geopolitical event, as they targeted the world financial center and provoked direct financial effects by forcing the closure of the U.S. stock market and other major financial institutions. Additionally, the 9/11 attacks had also a great impact on news coverage

[25] See Pastor and Veronesi (2013) for a model of political uncertainty and risk premia.

of terrorism and more broadly of geopolitical events outside the U.S..

Figure 12 reports the impulse responses to a GPR shock of a VAR model that includes two separate dummies for September and October 2001,[26] and another VAR model uses a censored version of the GPR index. We construct the censored GPR index by setting to zero all observations on the GPR where the time series of residuals from an AR(1) regression of the GPR increases by less than 1.68 standard deviations. The resulting index, which consists of 9 episodes, captures large increases in geopolitical risk, and gives no weight to small variations in the index.

The impulse responses to a GPR shock in the model with the 9/11 dummies and the censored GPR are very similar to the baseline model. By censoring the GPR index, a GPR shock induces a very temporary increase in geopolitical risk, which is zero 3 months after the shock. The lack of persistence translates into a more muted response of the stock market index and of corporate spreads.

Finally, in Figure 13 we explore the robustness to an alternative Cholesky ordering where we order the GPR index after the VIX and oil prices, that is, last in the block of fast-moving variables. By construction, the impact response of the VIX is zero. This exclusion restriction lowers the entire dynamic path of the VIX, which declines marginally after the shock. However, despite the decline in the VIX, a shock to GPR has the same recessionary effects as in the benchmark model. Hence, this result highlights that while one transmission channel of geopolitical risk could be higher global volatility, it has a significant independent effect.

In Figure A.6 in the Appendix, we show that the VAR results are robust to including the EPU in the specification.

4.5 Geopolitical Risk and Capital Flows

We now present additional evidence on the economic consequences of changes in global geopolitical risk, by showing how geopolitical risks affect capital inflows in a sample of advanced and emerging economies. The procyclical and volatile nature of capital flows makes them a leading policy concern, especially in economies which rely heavily on external funds. We use country-level, quarterly data on capital inflows from the IMF's Balance of Payments Statistics database. The sample consists of 22 advanced economies, 23 emerging economies and the United States. Our sample covers the period from 1990Q1 through 2015Q4.

Our baseline specification tests whether movements in geopolitical risk have any explanatory power for gross capital inflows. We choose gross capital inflows, that is purchases of domestic asset by foreign

[26] We included a dummy for both September and October because news coverage of the 9/11 terrorist attacks and especially their geopolitical implications was higher in October than in September.

residents, in line with a large and growing body of empirical evidence that shows that gross capital flows respond systematically to changes in global conditions, and in line with the notion that a measure of global geopolitical risk is more likely to matter for the economic decisions of global investors on where to allocate capital across countries. Our regression takes the form:[27]

$$y_{i,t} = \alpha_i + \rho_1 y_{i,t-1} + \beta GPR_t + \Gamma X_t + u_{i,t}, \tag{1}$$

where $y_{i,t} = inflows_{i,t}/GDP_{i,t}$ are gross capital inflows divided by annualized GDP, α_i are country fixed effects and X_t is a vector of control variables. We estimate equation (1) separately for emerging economies, advanced economies excluding the United States, and the United States. Throughout, we assume that the effect of GPR on the capital inflows is equal within each group of countries. Table 2 reports the results for a specification that includes VIX to control for global economic risk, as well as lagged GDP growth, following the work of Ahmed and Zlate (2014). As expected, the coefficient on the VIX is negative across all specifications. Additionally, the coefficient on lagged GDP growth is positive.

Comparing columns (1) and (2) of the table, an increase in GPR has different effects on capital inflows in emerging versus advanced economies. In emerging economies, a 1 standard deviation increase in the GPR reduces capital inflows by 0.25 percentage points, while an equally sized increase in the VIX reduces capital inflows by 0.91 percentage points. By contrast, for advanced economies increases in VIX reduces capital inflows, whereas increases in GPR lead to an increase in capital inflows. In column (3), the effect of changes in geopolitical risk in the United States is modest, and not statistically different from zero.

4.6 Geopolitical Risk and Stock Returns

In this section, we examine the reaction of country-specific stock market returns to changes in geopolitical risk. We do so both for the full sample covered by the historical index, and for the most recent period covered by the benchmark index. The baseline specification echoes the work of Berkman, Jacobsen, and Lee (2011), who do a similar analysis using data from the International Crisis Database looking at the effects of the number of crises on world stock returns. In particular, we run the following set of

[27] See Forbes and Warnock (2012) and Ahmed and Zlate (2014) for recent discussion of the focus on gross vs. net capital flows. As discussed in Ahmed and Zlate (2014), whether to look at net capital inflows or gross capital inflows is an open question. Conceptually, global geopolitical risk should more directly relevant for the investment decisions of foreign investors. For this reason, we focus on gross capital inflows. Note that gross capital inflows measure flows of both debt and equity capital from foreign investors and lenders, and that capital withdrawals are subtracted.

country-by-country regressions

$$r_{i,t} = \mu_i + \alpha_i GPRSHOCK_t + \varepsilon_{i,t}, \tag{2}$$

where $r_{i,t}$ is stock market return in country i in month t, and $GPRSHOCK_t$ is the residual of an autoregressive process of order 1 estimated for the GPR index. The results are seen in Table 3. In line with expectations, geopolitical risks have a strong negative impact on stock returns in virtually all advanced economies. Using the benchmark sample, an increase in geopolitical risk of 100 units reduces world stock returns, on average, by 1.23 percent. As shown by the Table, the effects are significant both in the pre– and post–World War II sample.

5 Conclusions

We construct a monthly index of geopolitical risk. This index captures an important dimension of uncertainty: the risk that events might happen that disrupt the normal and peaceful course of relations across states, populations, and territories.

It is often said that most of the enduring progress of many advanced economies in the last century owes to peace, democracy, and the creation of institutions for the collective action of countries in response to social and economic threats, and that collective prosperity and economic development preserve peace, cooperation and goodwill among countries. Our index captures one of the extreme risks posed by humans to such progress.

References

Ahmed, Shaghil and Andrei Zlate (2014). Capital flows to emerging market economies: a brave new world? *Journal of International Money and Finance 48*, 221–248. [18]

Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2016). Measuring economic policy uncertainty*. *The Quarterly Journal of Economics 131*(4), 1593. [2, 8, 11, 13]

Berkman, Henk, Ben Jacobsen, and John B Lee (2011). Time-varying rare disaster risk and stock returns. *Journal of Financial Economics 101*(2), 313–332. [12, 18, 24]

Blomberg, S.Brock, Gregory D Hess, and Athanasios Orphanides (2004). The macroeconomic consequences of terrorism. *Journal of Monetary Economics 51*(5), 1007 – 1032. [14]

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello (2016). Oil price elasticities and oil price fluctuations. Mimeo, Federal Reserve Board. [15]

Caldara, Dario, Matteo Iacoviello, and Aaron Markiewitz (2017). Country-specific geopolitical risk. Mimeo, Federal Reserve Board. [6]

Carney, Mark (2016). Uncertainty, the economy and policy. *Bank of England*. [2]

Chadefaux, Thomas (2014). Early warning signals for war in the news. *Journal of Peace Research 51*(1), 5–18. [13]

Dalby, Simon, Paul Routledge, and Gearóid Ó Tuathail (2003). *The geopolitics reader*. Routledge. [4]

Del Negro, Marco and Frank Schorfheide (2011). Bayesian macroeconometrics. In Herman van Dijk, Gary Koop, and John Geweke (Eds.), *Handbook of Bayesian Econometrics*, pp. 293–389. Oxford University Press. [15]

Eckstein, Zvi and Daniel Tsiddon (2004). Macroeconomic consequences of terror: theory and the case of israel. *Journal of Monetary Economics 51*(5), 971 – 1002. [14]

Flint, Colin (2016). *Introduction to geopolitics*. Routledge. [4, 8]

Forbes, Kristin J. and Francis E. Warnock (2012). Capital flow waves: Surges, stops, flight, and retrenchment. *Journal of International Economics 88*(2), 235–251. [18]

Glick, Reuven and Alan M Taylor (2010). Collateral damage: Trade disruption and the economic impact of war. *The Review of Economics and Statistics 92*(1), 102–127. [14]

Greenspan, Alan (2002). Remarks by chairman alan greenspan: banking. In *Washington, DC: Federal Reserve Board (At the annual convention of the American Bankers Association, Phoenix, Arizona.* [2]

Huntington, Samuel P (1997). *The clash of civilizations and the remaking of world order*. Penguin Books India. [8]

Jetter, Michael (2017). The effect of media attention on terrorism. *Journal of Public Economics 153*, 32 – 48. [6]

Lenain, Patrick, Marcos Bonturi, and Vincent Koen (2002). The economic consequences of terrorism. [14]

Pastor, Lubos and Pietro Veronesi (2013). Political uncertainty and risk premia. *Journal of Financial Economics 110*(3), 520 – 545. [16]

Saiz, Albert and Uri Simonsohn (2013). Proxying for unobservable variables with internet document-frequency. *Journal of the European Economic Association 11*(1), 137–165. [2, 13, A.1]

Sottilotta, Cecilia Emma (2013). Political risk: Concepts, definitions. Technical report, challenges. Working Paper Series, roma, LUISS. [4]

Tavares, Jos (2004). The open society assesses its enemies: shocks, disasters and terrorist attacks. *Journal of Monetary Economics 51*(5), 1039 – 1070. [14]

Table 1: Search Categories for the benchmark GPR index

Search Category	Words
1. Geopolitical Threats	Geopolitical AND (risk* OR concern* OR tension* OR uncertaint*) "United States" AND tensions AND (military OR war OR geopolitical OR coup OR guerrilla OR warfare) AND ("Latin America" OR "Central America" OR "South America" OR Europe OR Africa OR "Middle East" OR "Far East" OR Asia)
2. Nuclear Threats	("nuclear war" OR "atomic war" OR "nuclear conflict" OR "atomic conflict" OR "nuclear missile*") AND (fear* OR threat* OR risk* OR peril* OR menace*)
3. War Threats	"war risk*" OR "risk* of war" OR "fear of war" OR "war fear*" OR "military threat*" OR "war threat*" OR "threat of war" ("military action" OR "military operation" OR "military force") AND (risk* OR threat*)
4. Terrorist Threats	"terrorist threat*" OR "threat of terrorism" OR "terrorism menace" OR "menace of terrorism" OR "terrorist risk" OR "terror risk" OR "risk of terrorism" OR "terror threat*"
5. War Acts	"beginning of the war" OR "outbreak of the war "OR "onset of the war" OR "escalation of the war" OR "start of the war" (war OR military) AND "air strike" (war OR battle) AND "heavy casualties"
6. Terrorist Acts	"terrorist act" OR "terrorist acts"

NOTE: This table lists the combination words searched in the construction of the GPR index. The exact search query, inclusive of words that are excluded from the search, is in the Appendix.

Table 2: GPR and Capital Flows

Variables	(1) Inflows/GDP Emerging Economies	(2) Inflows/GDP Advanced Economies	(3) Inflows/GDP United States
Lagged Inflows	0.290***	0.142**	0.499***
	(0.0595)	(0.0532)	(0.0863)
GPR Index, standardized	-0.248*	0.914***	0.415
	(0.132)	(0.317)	(0.377)
VIX, standardized	-1.081***	-1.496**	-1.051**
	(0.198)	(0.698)	(0.445)
Lagged GDP Growth	0.281***	1.967***	-0.271
	(0.0953)	(0.476)	(0.714)
Constant	-20.93**	-183.0***	32.64
	(9.536)	(47.13)	(71.64)
Observations	1,796	2,065	103
R-squared	0.157	0.047	0.333
Number of country	23	22	1
Country FE	YES	YES	—
Clustered SEs	YES	YES	—

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

NOTE: Source: Balance of payments (BOP) quarterly data collected from national sources. Gross inflows are BOP liabilities, and consist of the nonresidents purchases of domestic assets net of sales. The sample of advanced economies includes Japan, Germany, France, UK, Italy, Canada, Spain, Australia, Netherlands, Switzerland, Sweden, Belgium, Norway, Austria, Denmark, Greece. Finland, Portugal, New Zealand, Slovakia, Slovenia, Estonia. The sample of emerging economies includes China, Brazil, India, Korea, Mexico, Indonesia, Turkey, Argentina, Venezuela, South Africa, Thailand, Colombia, Malaysia, Israel, Chile, Philippines, Peru, Ecuador, Jordan, El Salvador, Russia, Saudi Arabia, and Latvia. We drop from the sample countries with a standard deviation of inflows to GDP larger than 50 percent (Hong Kong, Ireland, Singapore, Luxembourg and Iceland). GDP growth is quarter-on-quarter GDP growth, expressed in percentage terms. The sample goes from 1990Q1 through 2015Q4, for a maximum of 104 observations per country.

Table 3: GPR and Stock Market Returns

Country	GPR Benchmark			GPR Historical			GPR Historical, Post-1950		
	Coefficient	Std. Errors	Sample	Coefficient	Std. Errors	Sample	Coefficient	Std. Errors	Sample
Australia	-.45	.51	1985-2017	-.48	.34	1900-2017	.12	.55	1950-2017
Austria	-.87	.81	1985-2017	-1.6	.67	1922-2017	-1.14	.71	1950-2017
Belgium	-2.29	.72	1985-2017	-1.06	.5	1900-2017	-1.59	.6	1950-2017
Canada	-1.06	.43	1985-2017	-1.12	.37	1914-2017	-.89	.48	1950-2017
Denmark	-.94	.86	1985-2017						
Finland	-.02	1.17	1985-2017	.1	1.04	1900-2017	.1	1.04	1950-2017
France	-2.05	.9	1985-2017	-1.05	.54	1900-2017	-1.16	.8	1950-2017
Germany	-2.22	1.32	1985-2017	-.99	1.18	1923-2017	-.99	1.18	1950-2017
Greece	-4.57	1.36	1988-2017						
Ireland	-1.82	.93	1988-2017						
Italy	-2.54	.94	1985-2017	-1.38	.6	1905-2017	-2.21	.85	1950-2017
Japan	-.54	.59	1985-2017	.13	.49	1914-2017	-.34	.62	1950-2017
Netherlands	-2.22	.86	1985-2017	-.56	.6	1919-2017	-1.13	.75	1950-2017
New Zealand	.22	1.13	1985-2017						
Norway	-2.32	.82	1985-2017						
Portugal	-1.05	.94	1985-2017	-.36	1.13	1934-2017	-.36	1.13	1950-2017
Spain	-1.47	.82	1985-2017	-.46	.54	1915-2017	-.99	.74	1950-2017
Sweden	-1.48	.89	1985-2017	-.69	.45	1906-2017	-1.11	.66	1950-2017
Switzerland	-.88	.67	1985-2017						
United Kingdom	-1.67	.63	1985-2017	-.9	.36	1900-2017	-.97	.59	1950-2017
United States	-1.35	.59	1985-2017	-.82	.47	1900-2017	-.99	.51	1950-2017
World	-1.23	.59	1985-2017	-1.02	.52	1919-2017	-.83	.53	1950-2017

NOTE: Estimation of individual country effects of geopolitical risk on stock market returns. The table reports the estimated α's for the following model: $r_{i,t} = \mu_i + \alpha_i GPRSHOCK_t + \varepsilon_{i,t}$, where $r_{i,t}$ is stock market return in country i in month t, and $GPRSHOCK_t$ is the residual of an AR(1) process estimated for GPR divided by 100, so that the coefficient measures the percentage impact on stock returns from an innovation in GPR of 100 units. Standard errors are corrected for autocorrelation using the Newey-West method. Source: Morgan Stanley Capital International (MSCI) without dividends. Prior to 1970, world returns are calculated using the weighted average of country-specific returns. We use the same weights as Berkman, Jacobsen, and Lee (2011) which are as follow: United States 41%, United Kingdom 12%, Germany 8%, France 8%, Japan 6%, Italy 4%, Canada 3%, Switzerland 2.5%,The Netherlands 2.5%, Belgium 2%, Spain 2%, India 2%, Australia 2%, South Africa 2%, Denmark 1%, Norway 1%, and Sweden 1%.

Figure 1: The Geopolitical Risk Index

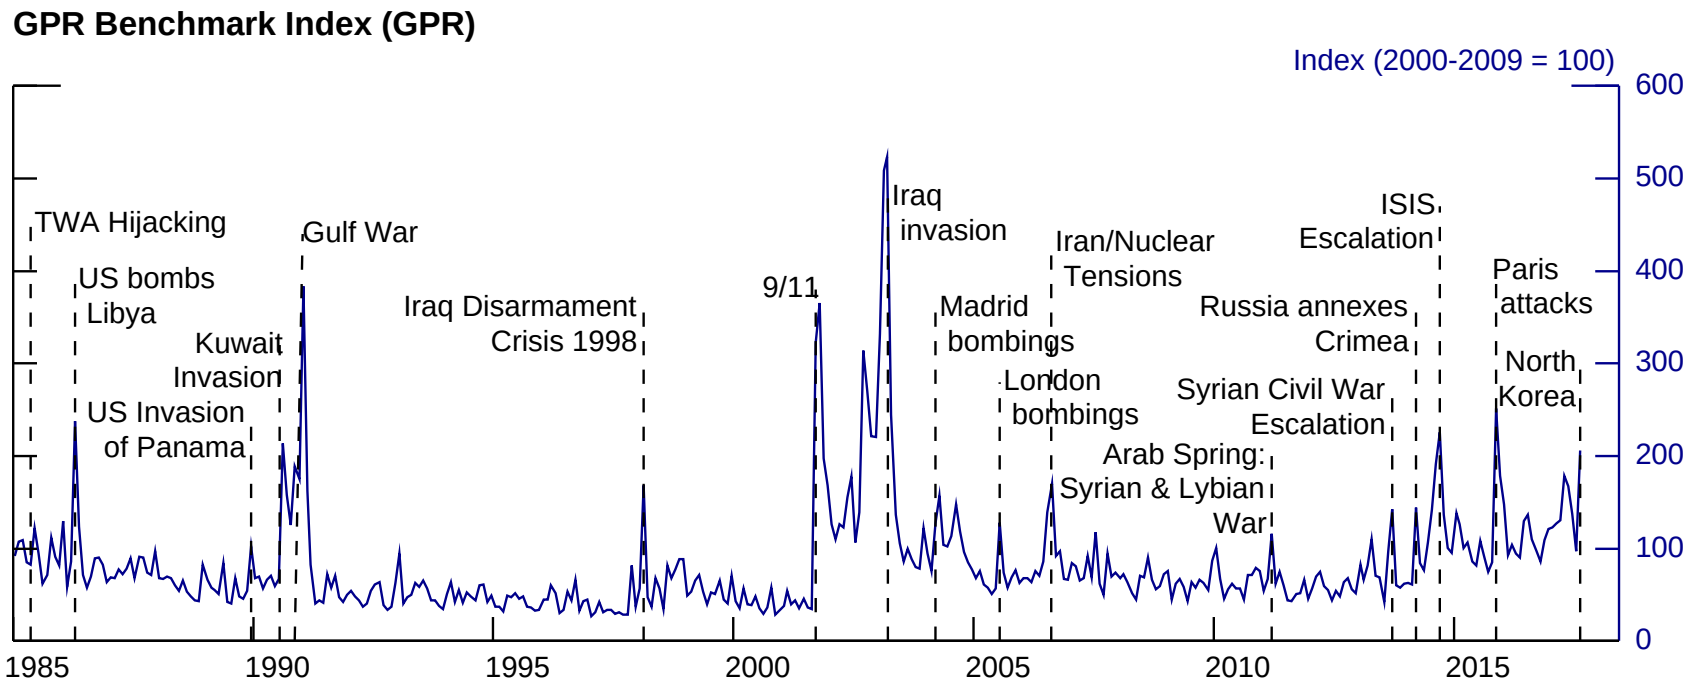

Note: The line plots the benchmark GPR index.

Figure 2: The Historical Geopolitical Risk Index

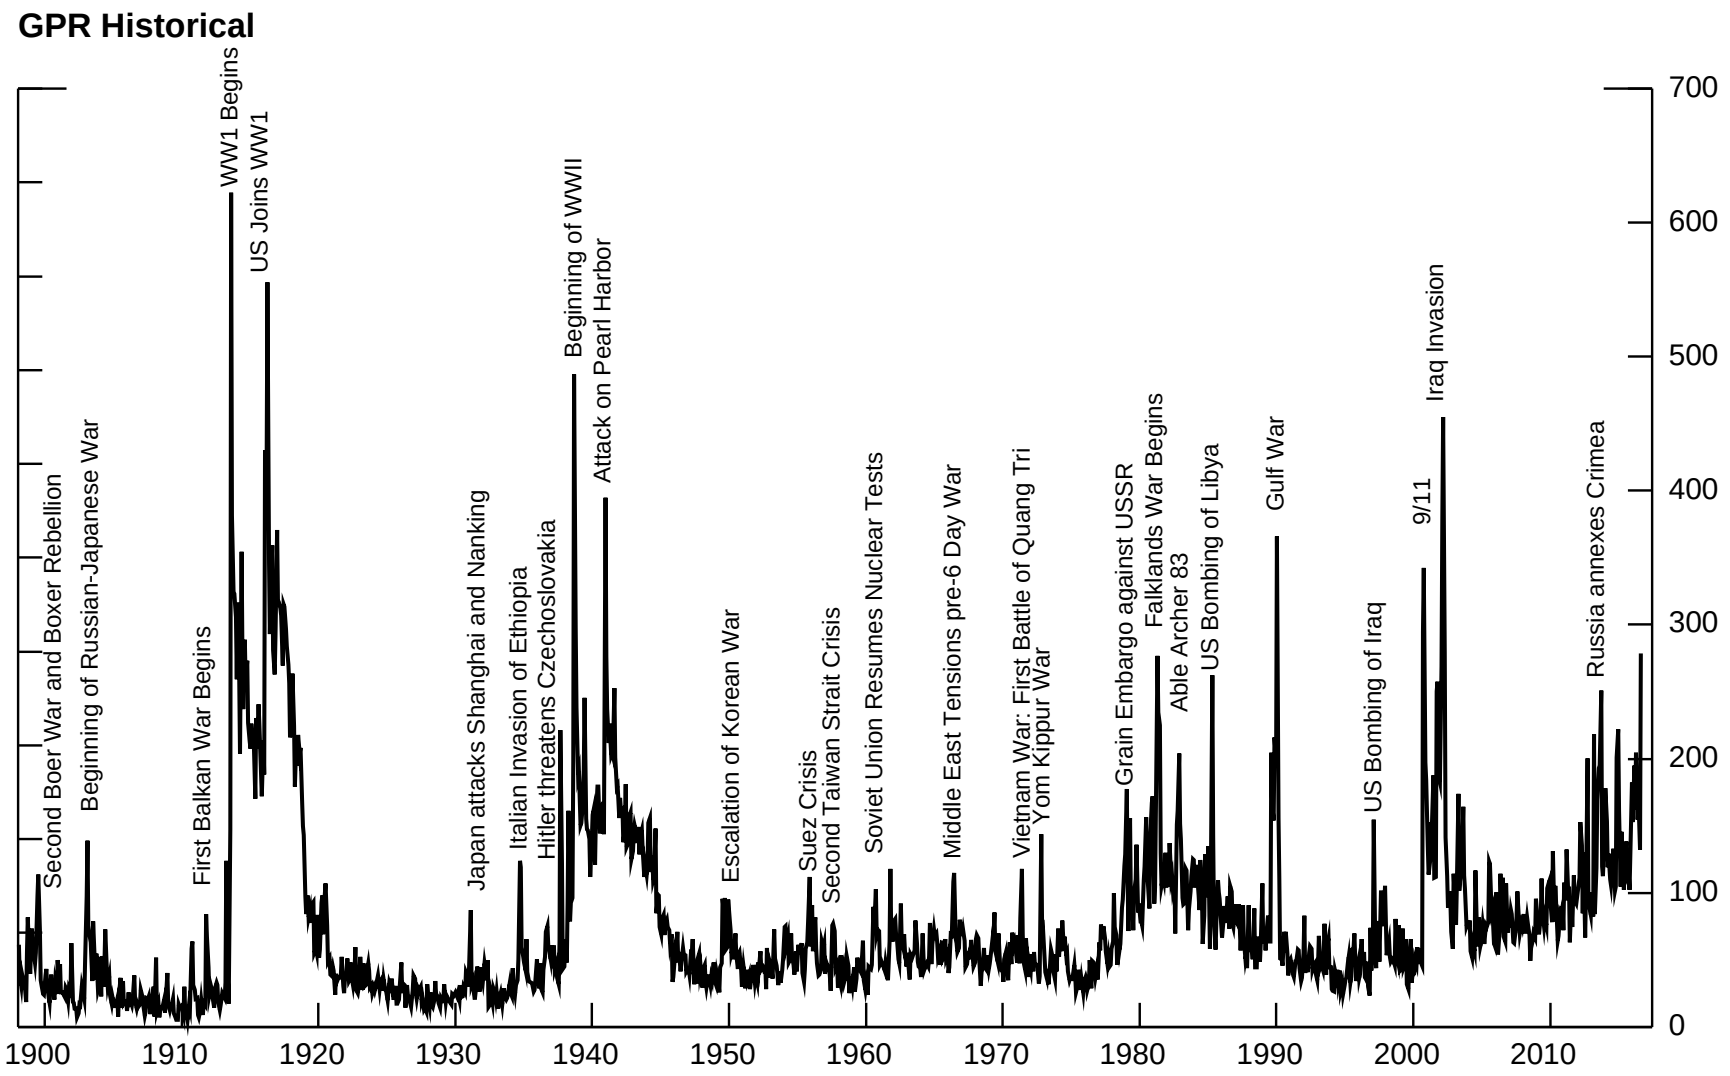

Note: The line plots the monthly GPR index since 1900.

Figure 3: The Geopolitical Risk Index: Contribution of Various Words

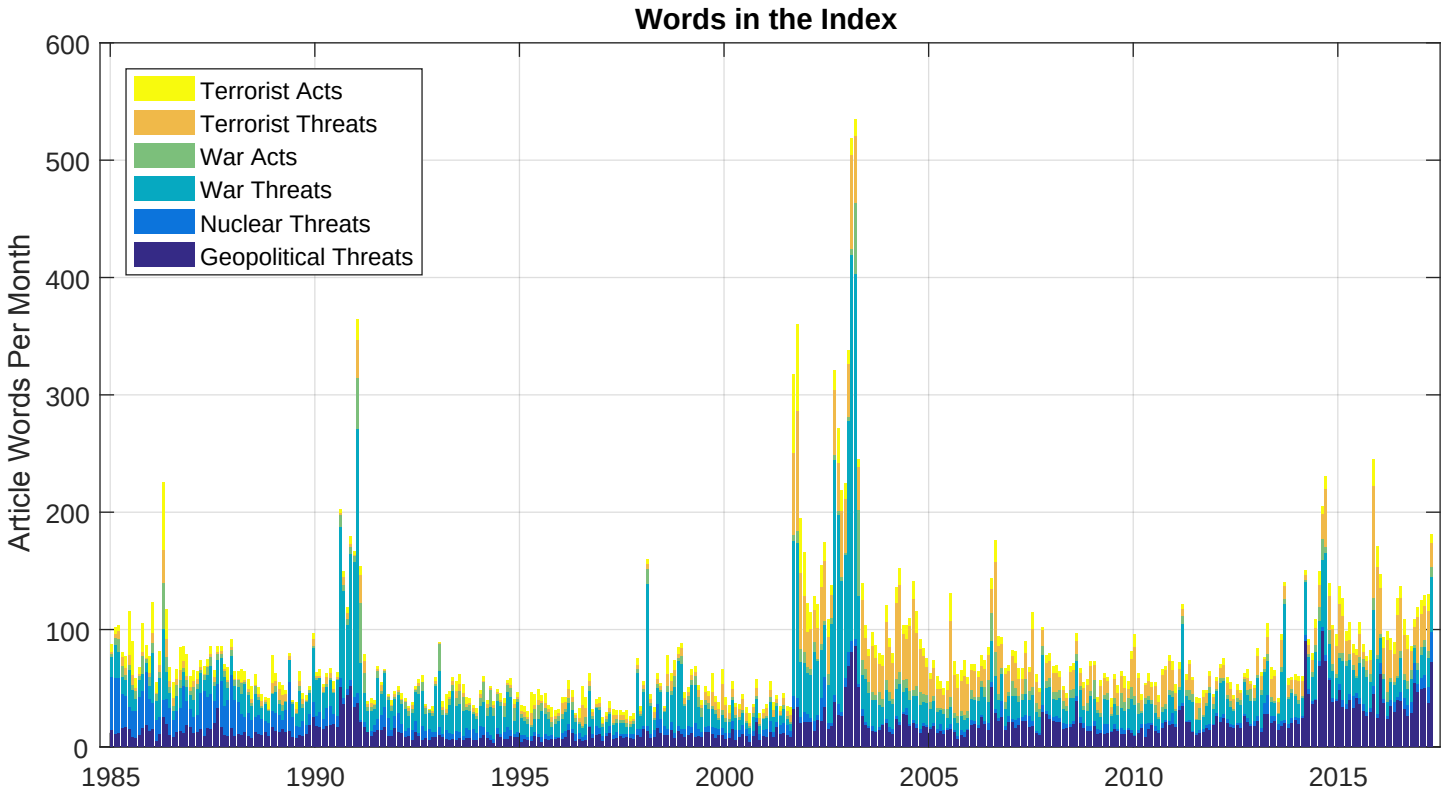

Note: The chart plots the cumulative contribution of the various phrases that enter the GPR index (top panel), and of some other phrases that do not enter the index (bottom panel). Higher geopolitical risk since the 2000s reflects increased mentions of both terrorist acts and terrorist threats, as well as an increased use of terms directly mentioning geopolitical uncertainties.

Figure 4: The Geopolitical Risk Index: The Two Sub-Indexes

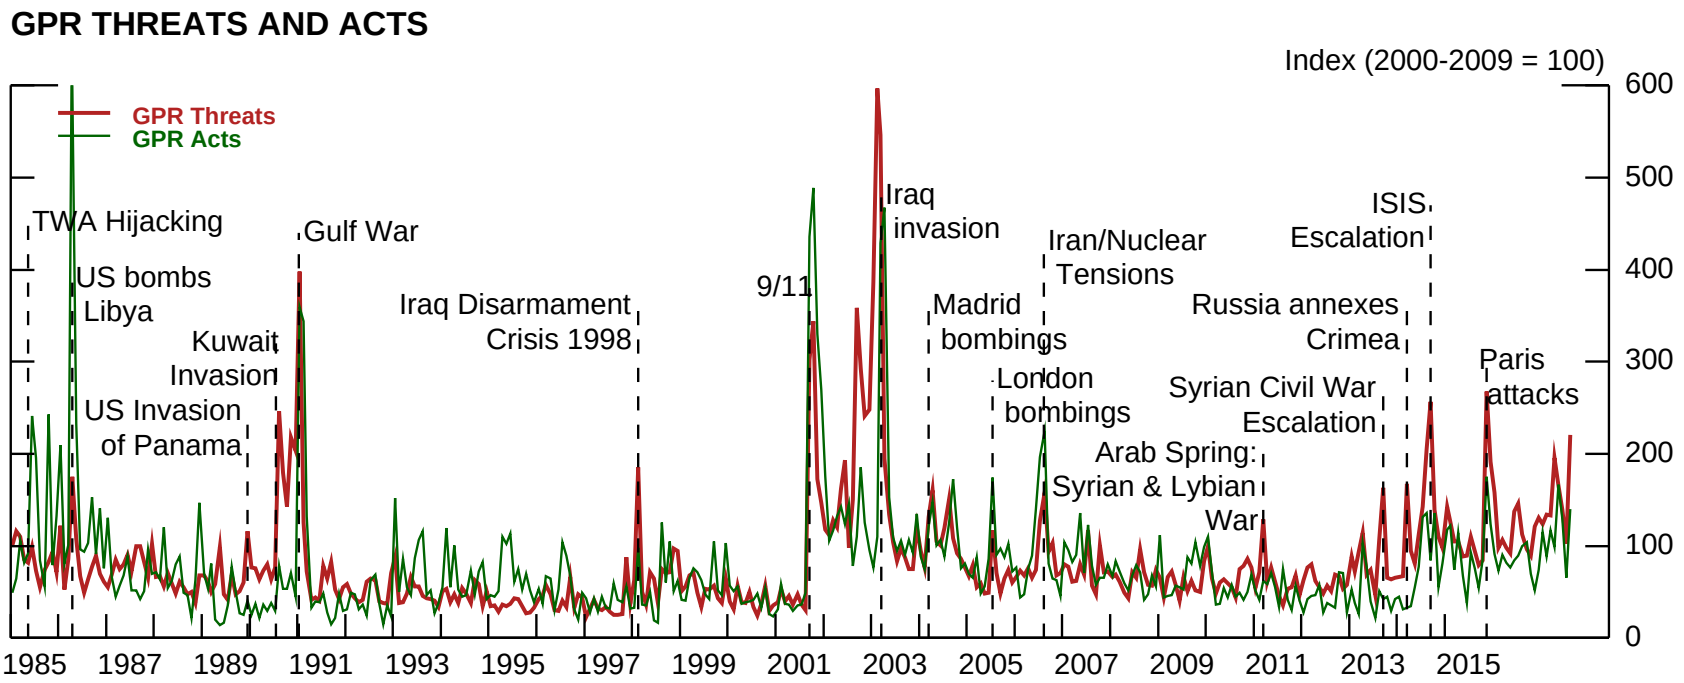

Note: This figure decomposes the benchmark index into two subcomponents. The GPR RISKS component measures only the risk aspect of geopolitical tensions, as measured by search terms 1, 2, 3, 4 and 5 in Table 1. The GPR ACTS component measures the actual occurrence of events, as measured by search terms 6, 7 and 8 in Table 1.

Figure 5: The Historical Geopolitical Risk Index: Contribution of Various Words to the Index

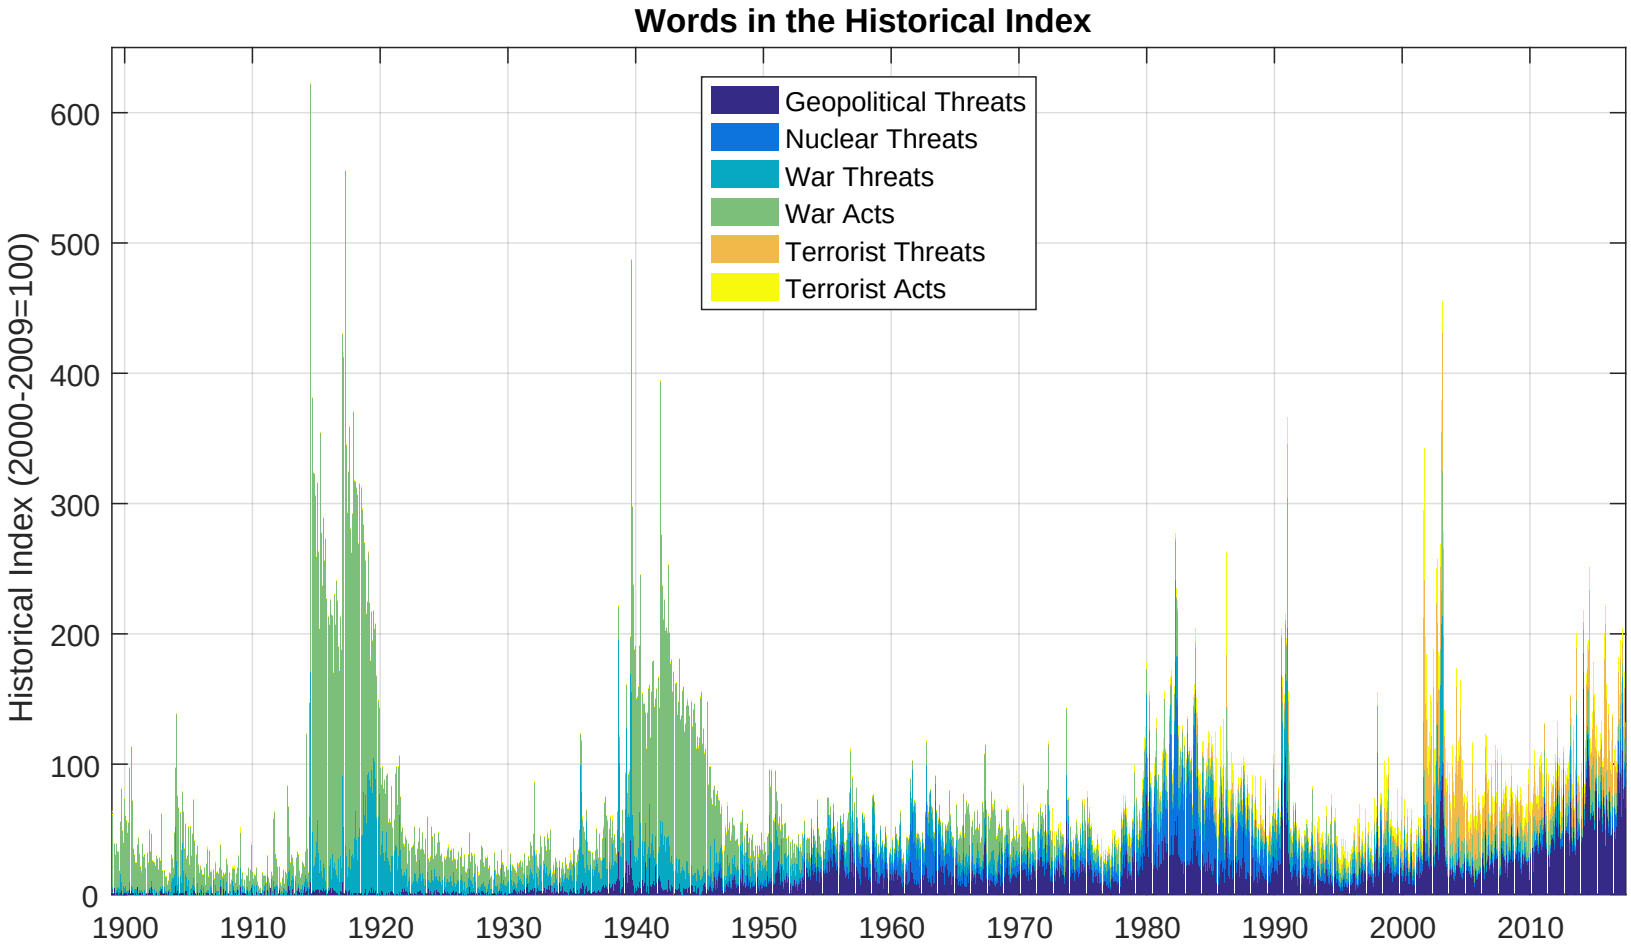

Note: The line plots the GPR index since 1899, broken down by the contribution share of each of its search categories to the total index.

Figure 6: Human and Computer-Generated GPR Indexes

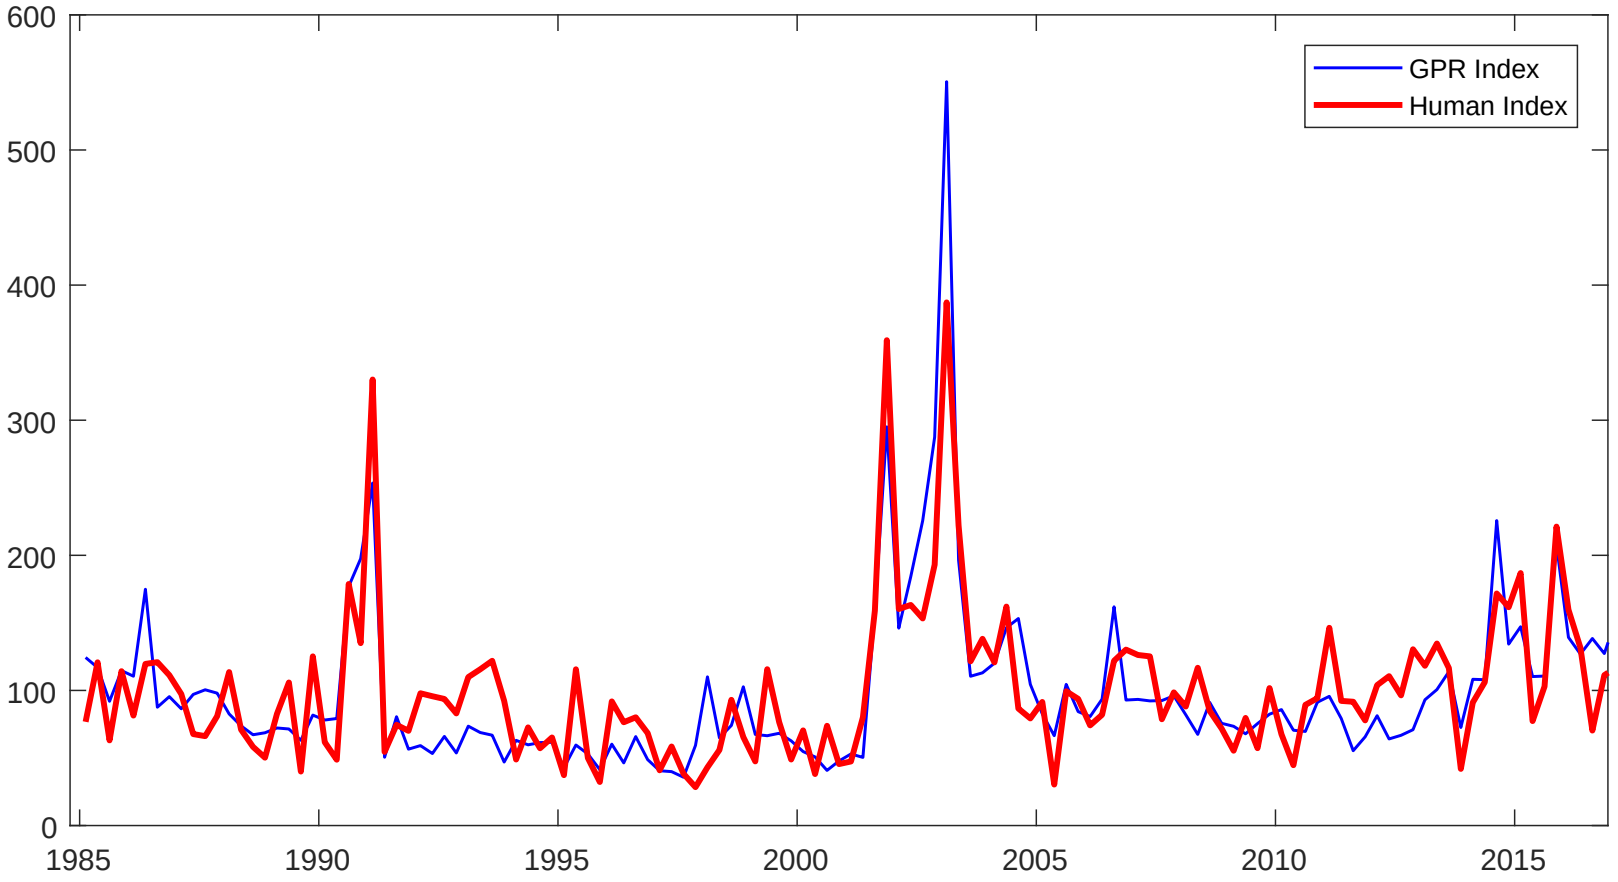

Note: Time-series comparison from 1985Q1 to 2016Q4 based on 6,125 articles. The series are plotted quarterly to reduce sampling variability. Both series are normalized to 100 from 1985 to 2016.

Figure 7: The GPR Index and three Alternatives

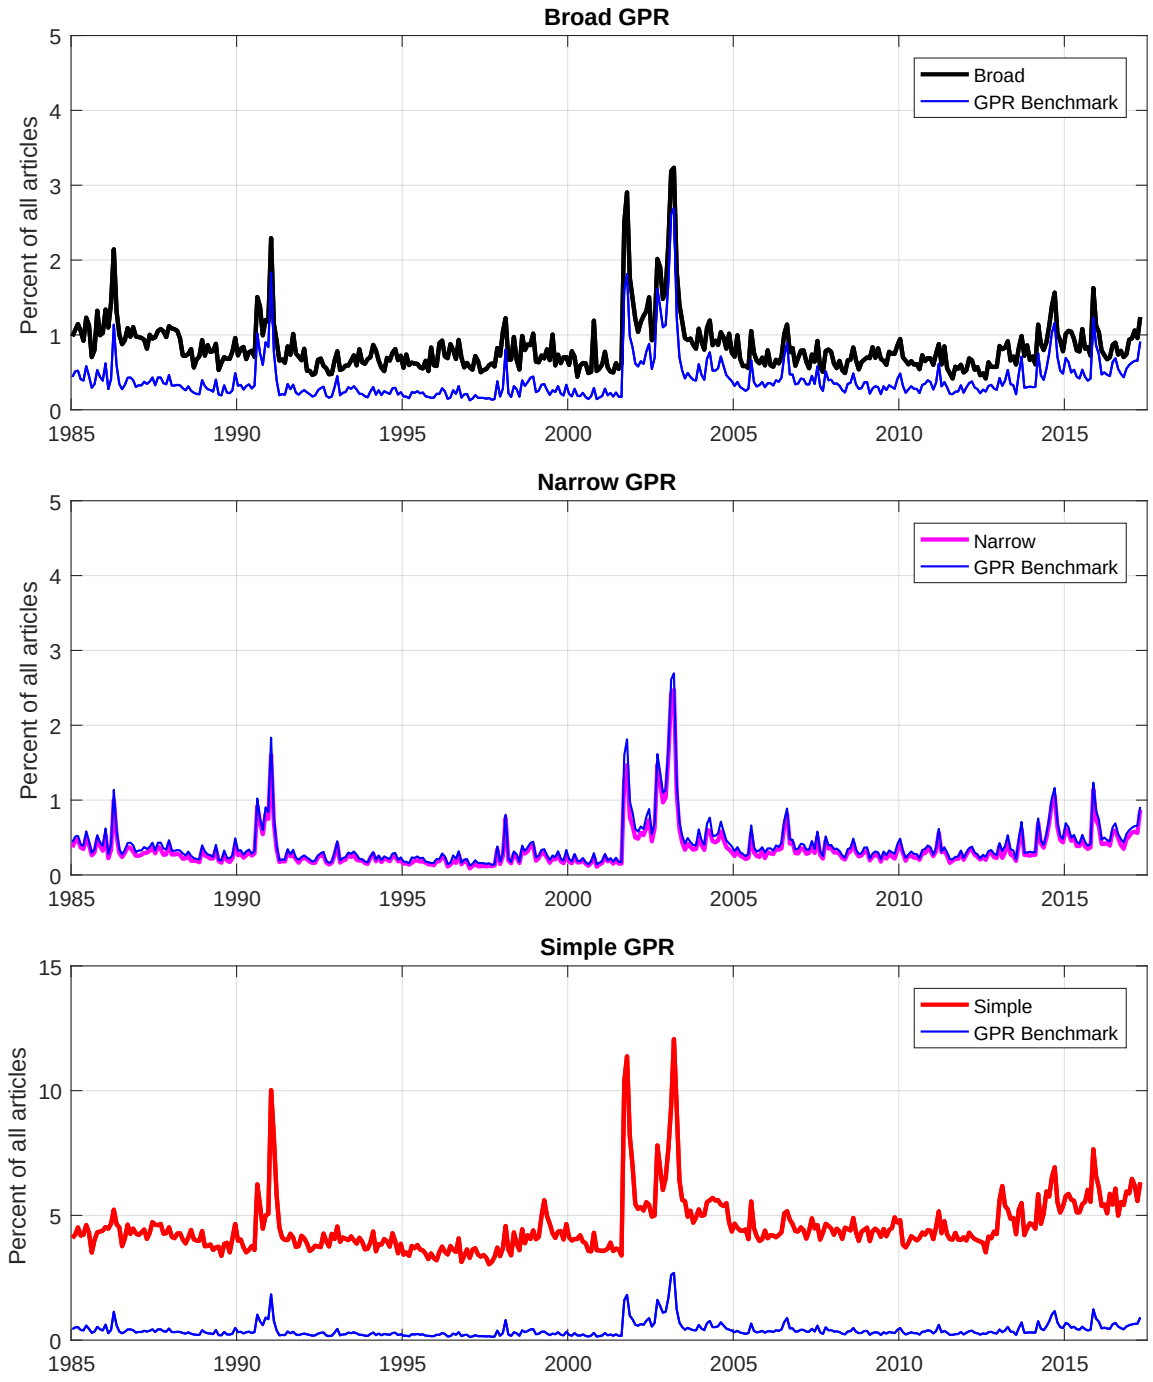

Note: The broad GPR combines the search terms in Table 1 with bigrams that a machine learning algorithm signals as very likely indicators of rising geopolitical tensions. The narrow GPR excludes articles containing the search terms which, in the human audit, were most likely associated with false positives. The simple GPR is based on articles that contain at least one word from each of two sets of terms: the set $\mathcal{S}_1$, including {war OR military OR terrorism OR geopolitical}, and the set $\mathcal{S}_2$, containing {risk* OR concern* OR tension* OR uncertaint* OR threat*}

Figure 8: The Geopolitical Risk Index and Other Proxies for Geopolitical Risk

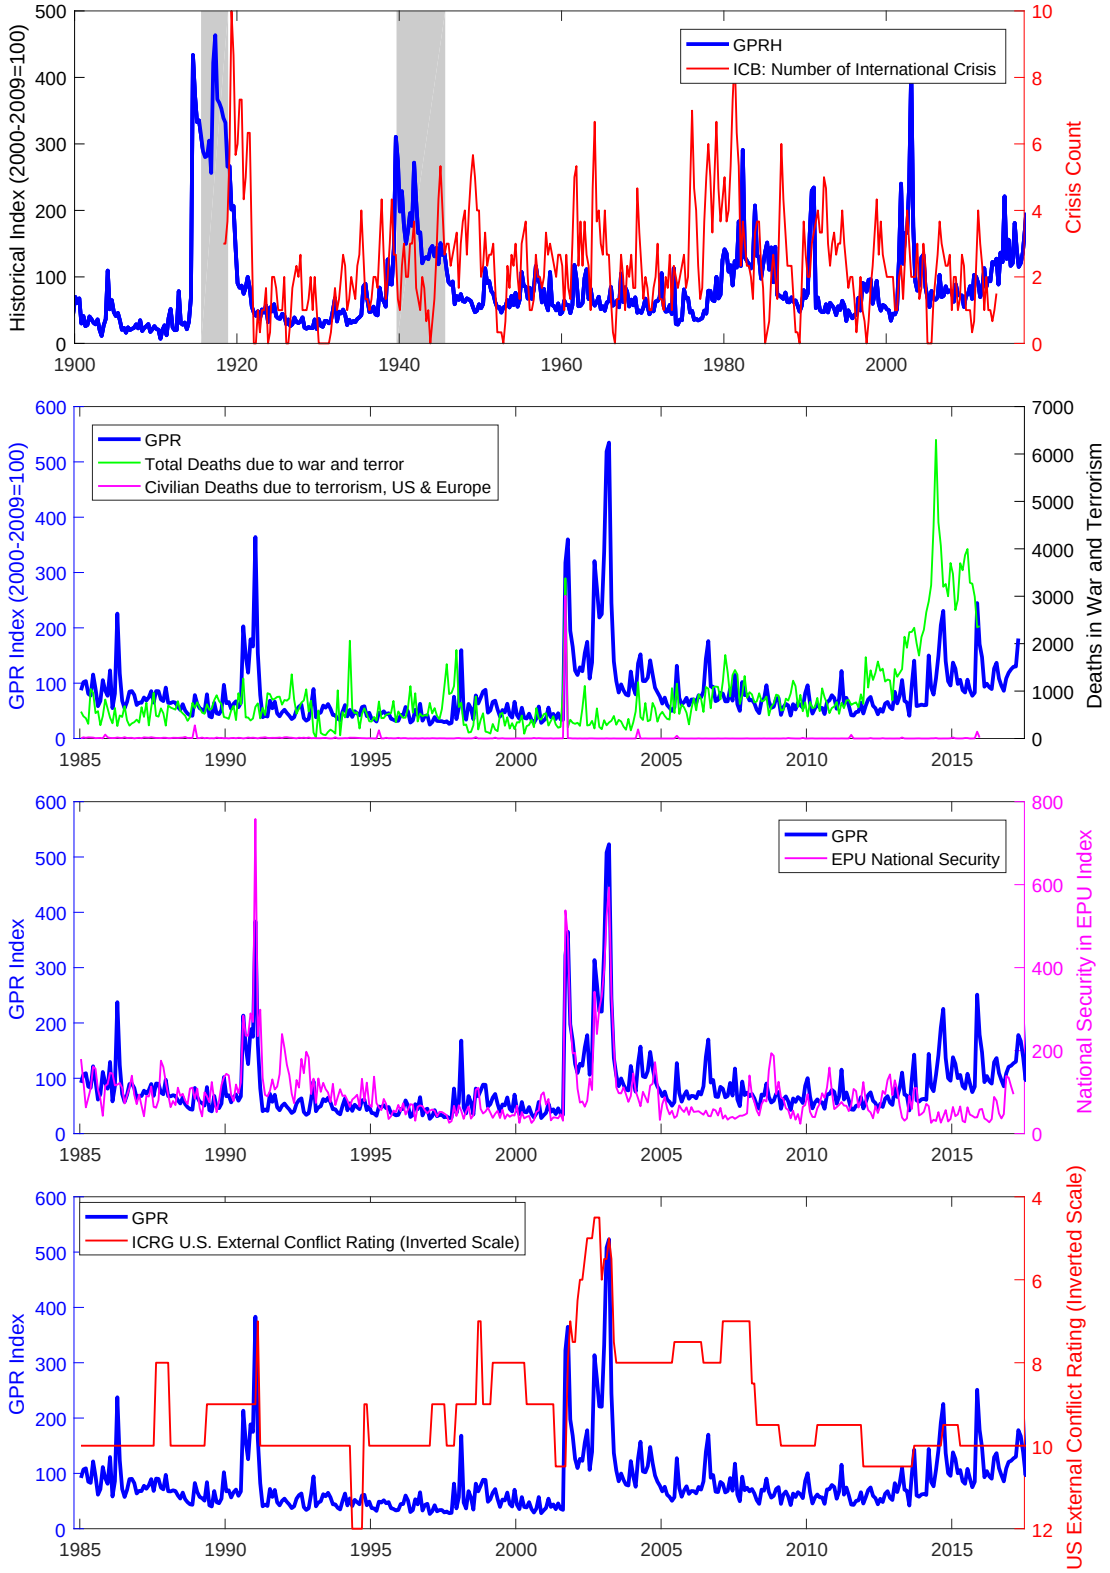

Note: See text. In the first panel, shaded areas represent the two World Wars.

Figure 9: The Geopolitical Risk Index and other Measures of Risk

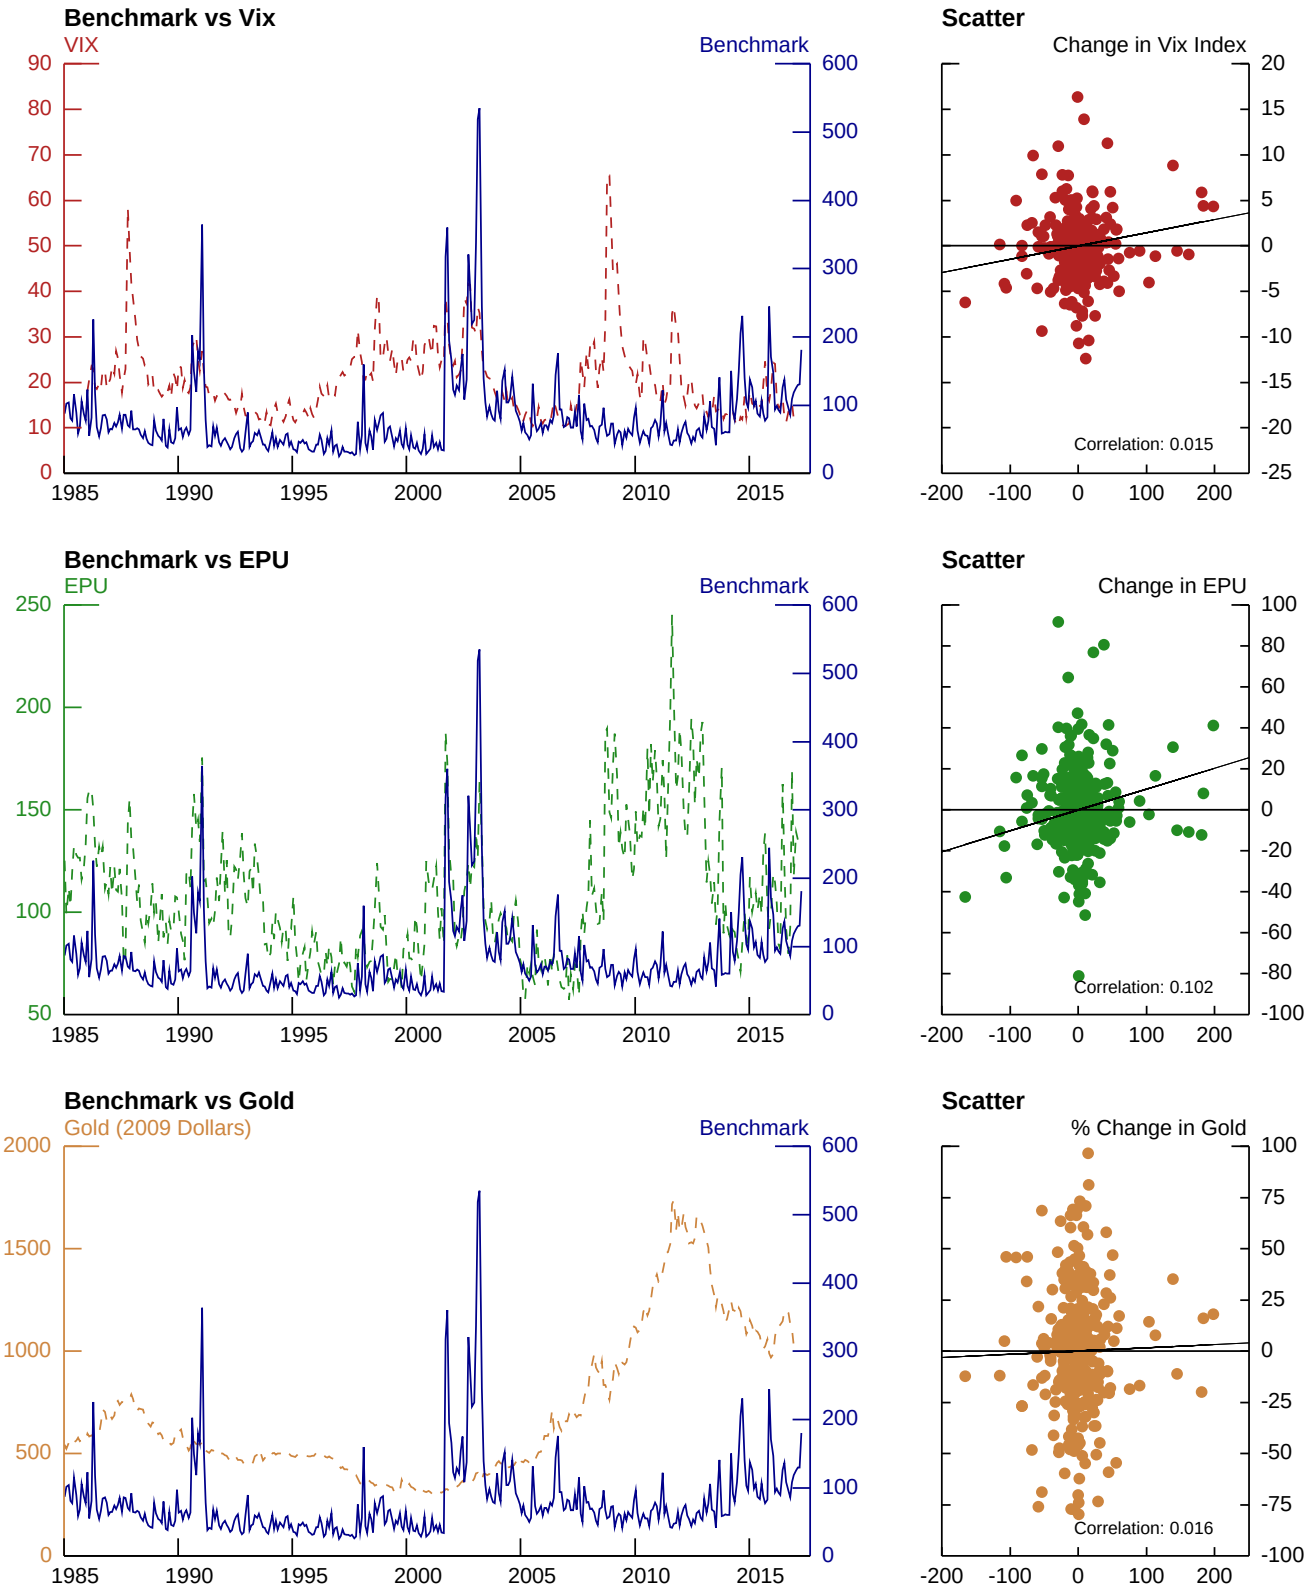

Note: This figure compare the GPR index with VIX (Haver Code: SPVXO), the EPU Index, and with the price of Gold.

Figure 10: The Impact of Increased Geopolitical Risk

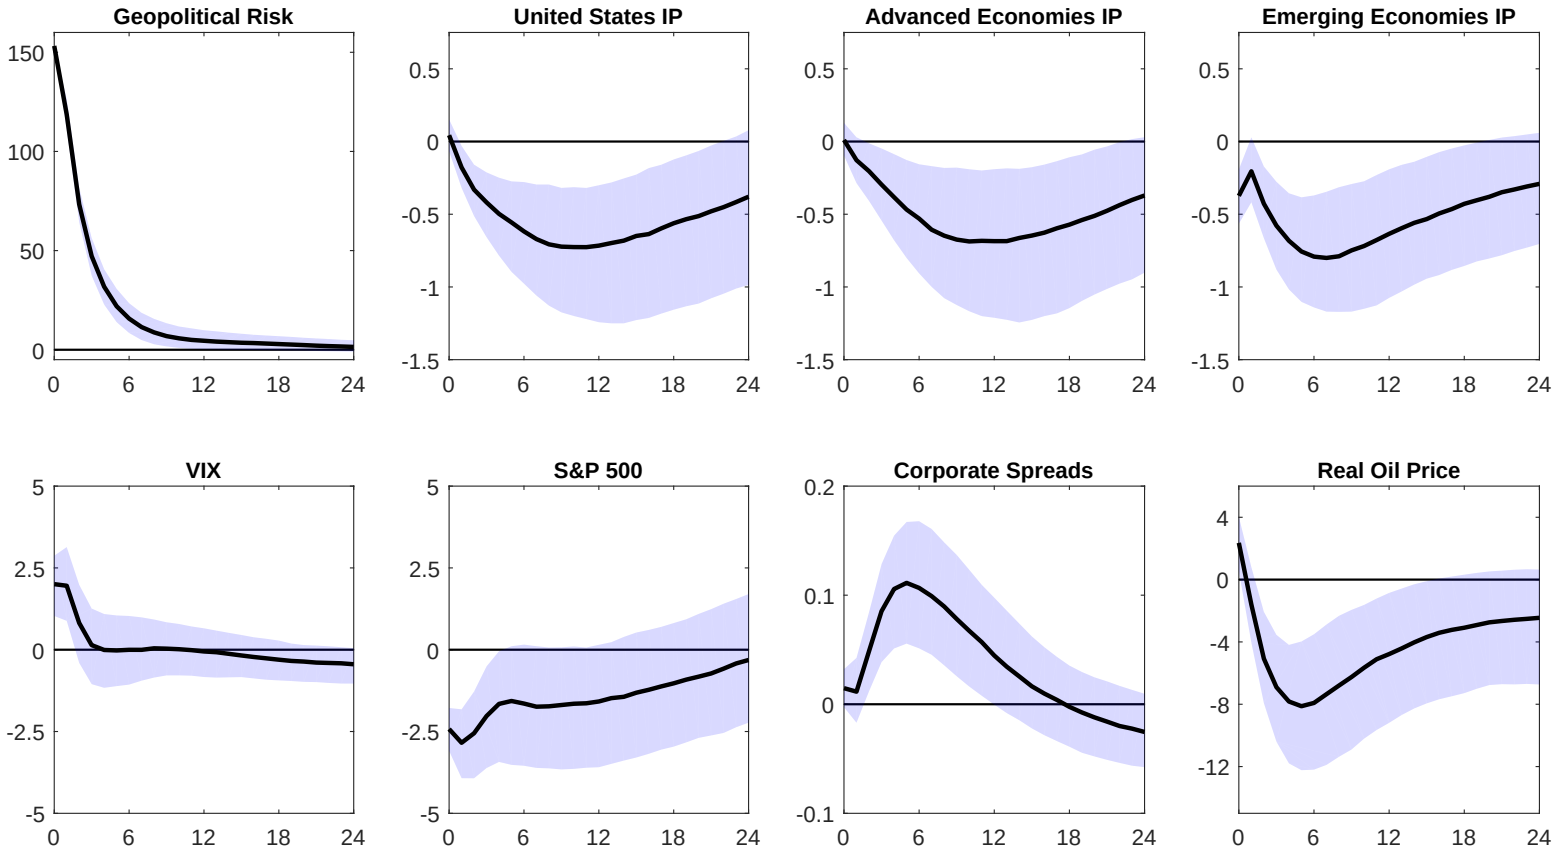

Note: The black solid line in each panel depicts the median impulse response of the specified variable to a rise of 120 points in the GPR index, while the shaded bands represent the 68 percent pointwise credible set. GPR index and the VIX are in levels, the corporate spread in basis points. All responses are measured in percent. The horizontal axis measures months since the shock.

Figure 11: The Impact of Increased Geopolitical Risk: Acts vs Threats

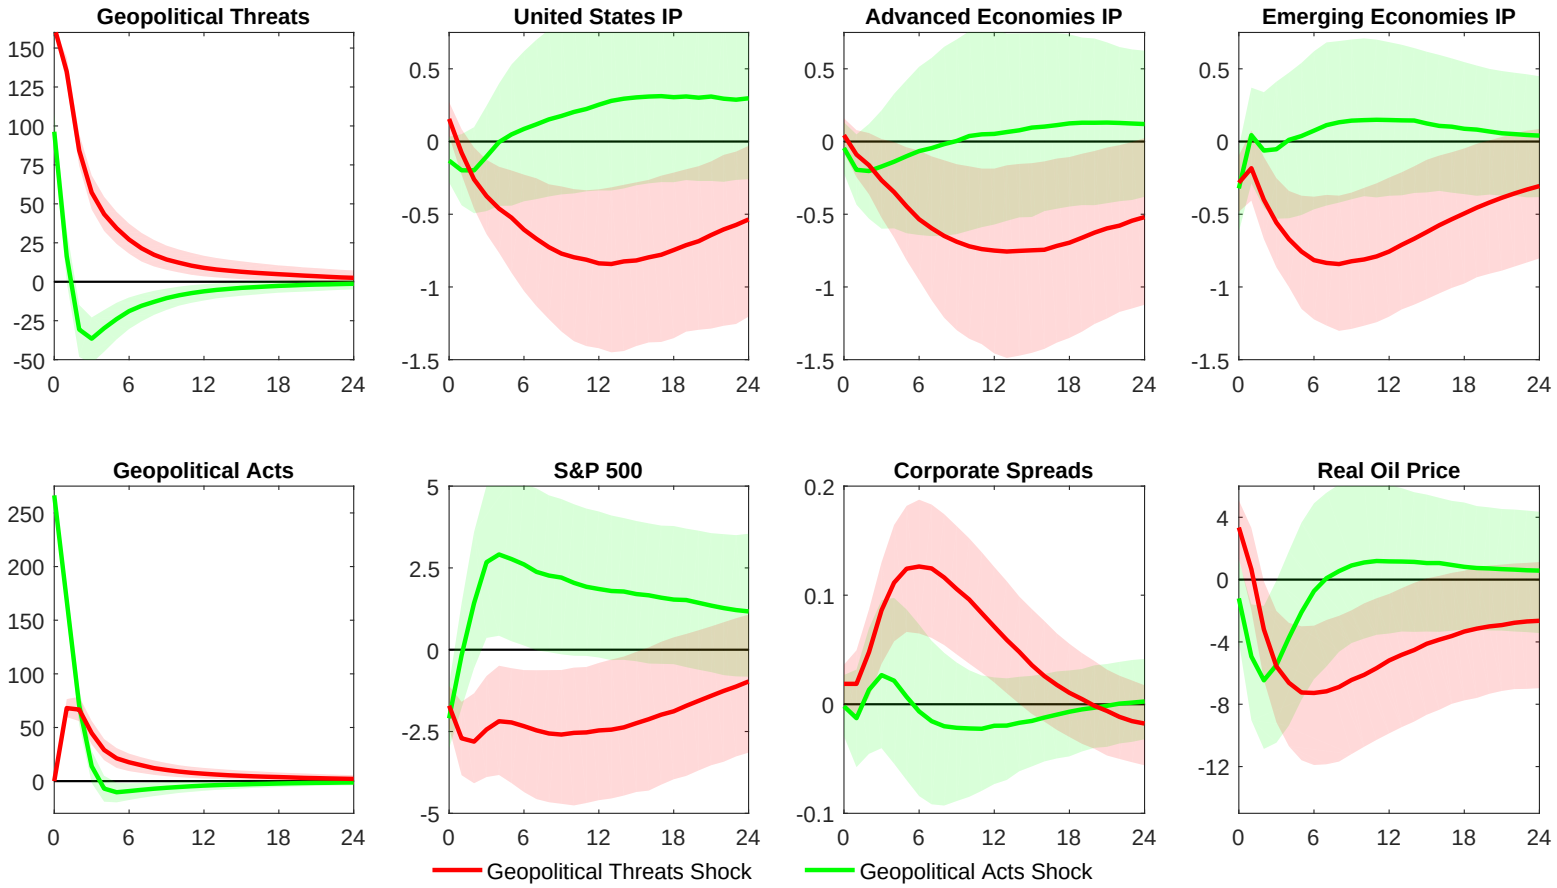

Note: VIX included in the VAR but not shown. GPR indexes in levels, corporate spread in basis points. All responses are measured in percent. The horizontal axis measures months since the shock.

Figure 12: The Impact of Increased Geopolitical Risk: 9/11 and a Censored Version

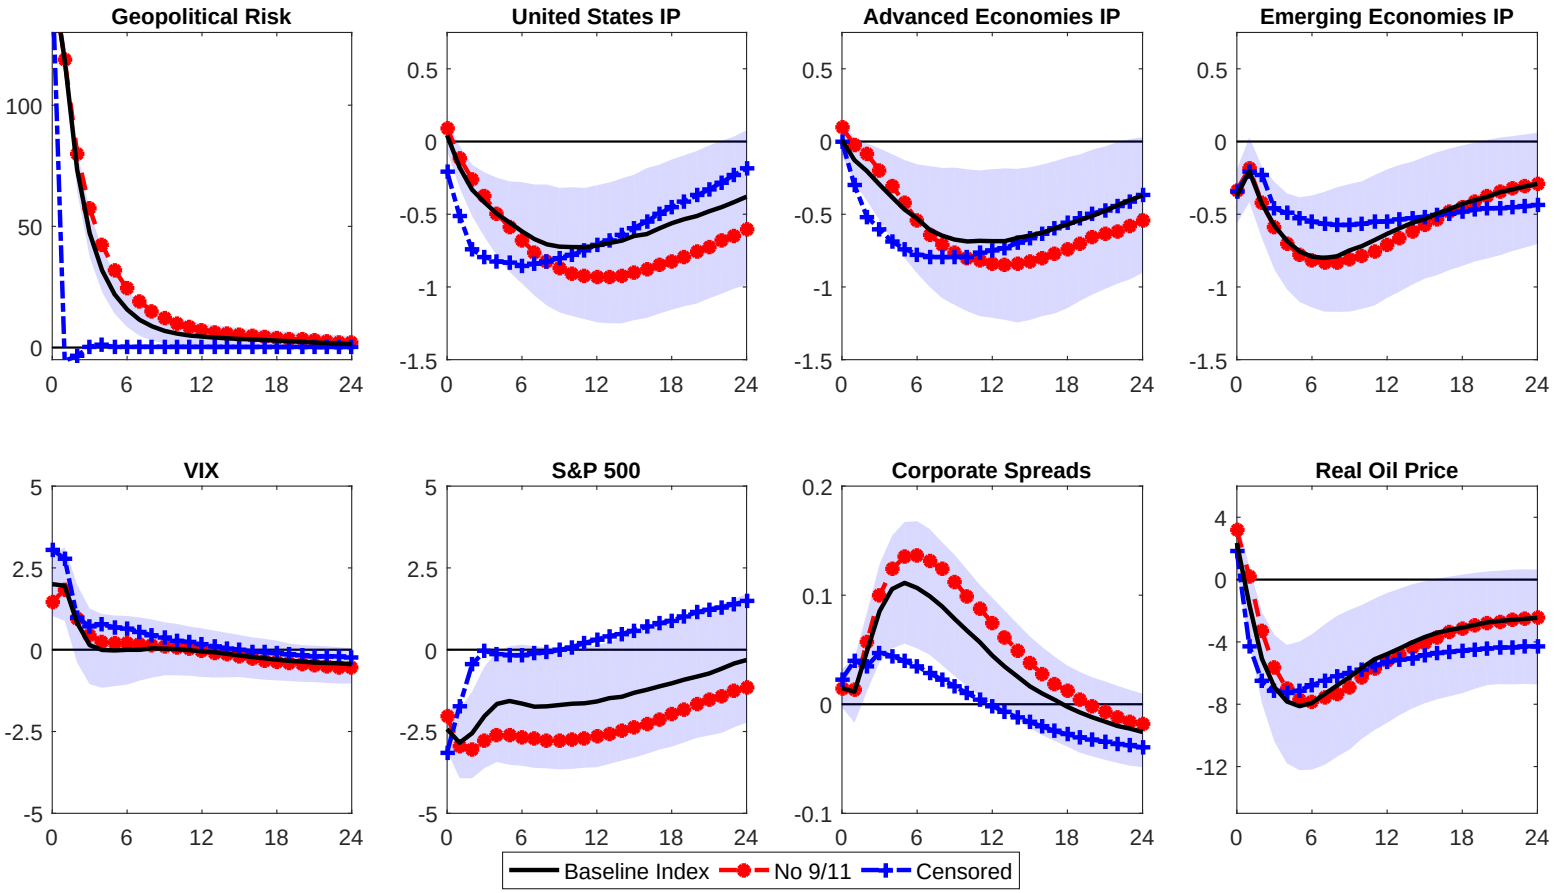

Note: The black solid line in each panel depicts the median impulse response of the specified variable to a rise of 120 points in the GPR index, while the shaded bands represent the 68 percent pointwise credible set. The horizontal axis measures months since the shock.

Figure 13: The Impact of Increased Geopolitical Risk: Alternative Ordering

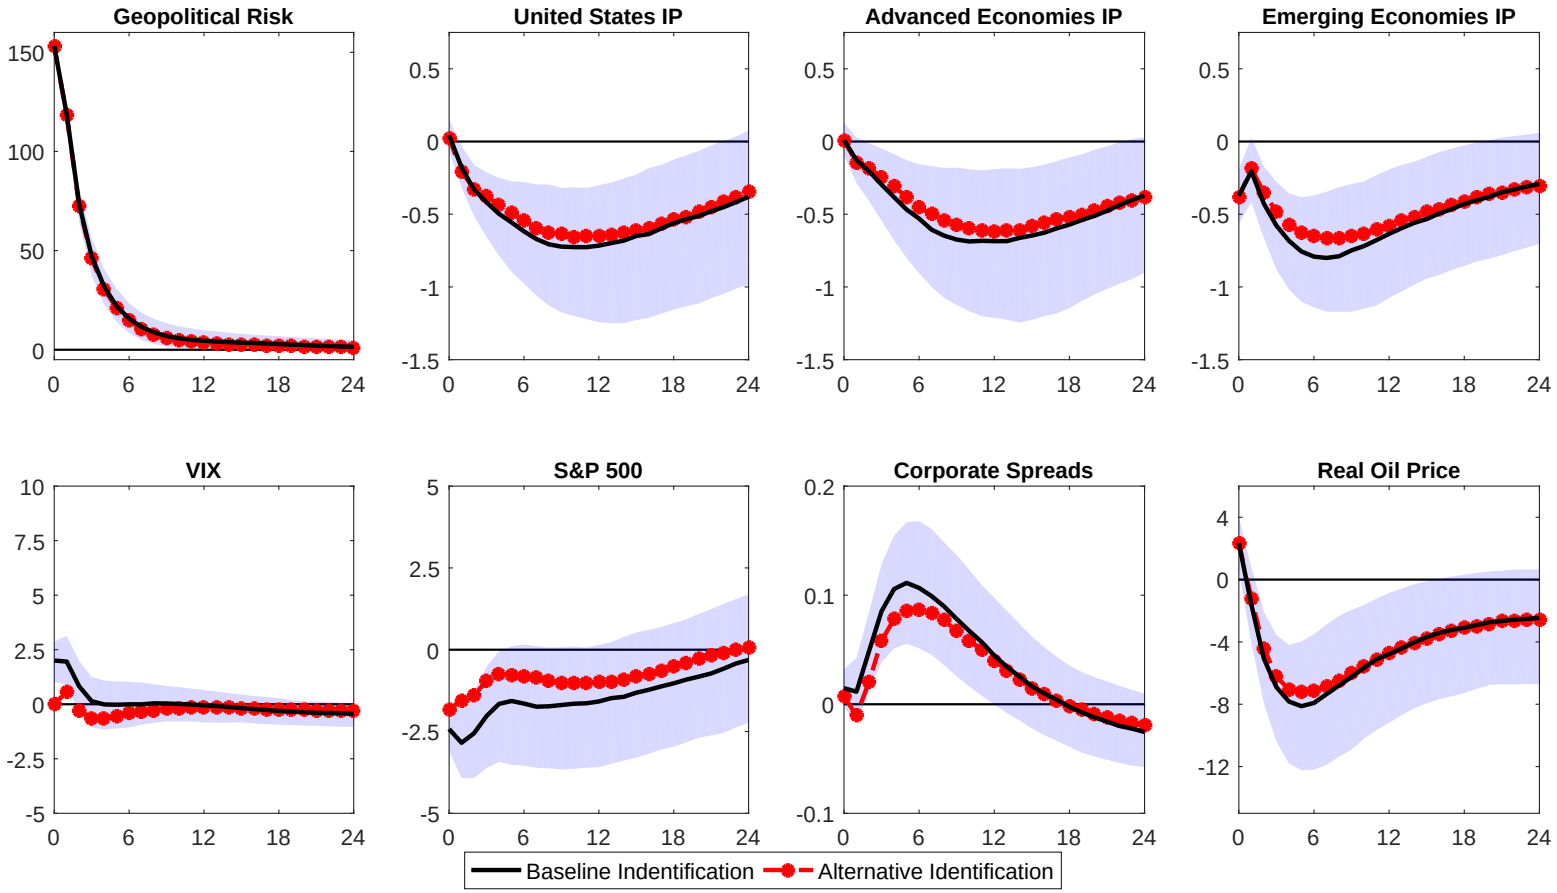

Note: The black solid line in each panel depicts the median impulse response of the specified variable to a rise of 120 points in the GPR index, while the shaded bands represent the 68 percent pointwise credible set. The horizontal axis measures months since the shock.

Appendix

A Additional Details on the Construction of the Index

The benchmark index is constructed by running a search query in the ProQuest Newsstand Database. We search the archives of the following newspapers. Boston Globe (start date availability: 1/1/1985); Chicago Tribune (1/1/1985); The Daily Telegraph (4/1/1991); Financial Times (5/31/1996); The Globe and Mail (1/1/1985); The Guardian (8/18/1992); Los Angeles Times (1/1/1985); The New York Times (1/1/1985); The Times (4/18/1992); Wall Street Journal (1/1/1985); and The Washington Post (1/1/1985). For our benchmark index, we run the search query shown in Figure A.1. The index is the aggregate number of articles returned by the search, divided by the total number of articles, indexed to 100 in the 2000–2009 decade. The newspaper-specific indexes are shown for robustness in Figure A.2. As the top left panel shows, coverage of geopolitical risks aligns with the benchmark GPR index for the three general interest newspapers that we use in the construction of the historical index. As the top right panel shows, coverage of geopolitical risks is slighly higher than the average for the two business newspapers in the sample, the Wall Street Journal and the Financial Times. Coverage also lines up with the average for the two U.S. newspapers not included in the historical index (bottom left panel). Coverage of geopolitical events by non-U.S. general interest newspapers lines up with the average, but is slightly more volatile (bottom right panel). We use the search tips and wild cards discussed at `https://proquest.libguides.com/proquestplatform/tips`.

B Additional Checks

Saiz and Simonsohn (2013) propose a number of data checks to confirm that the GPR index is a useful proxy for the phenomenon we want to measure, geopolitical risk. We describe how we perform this data check below.

1. We verify that our search terms are more likely to be used when geopolitical risk is high than when it is low (*Data check 1: Do the different queries maintain the phenomenon and keyword constant?*, and *Data check 3: Is the keyword employed predominately to discuss the occurrence rather than non-occurrence of phenomenon?*). Across all the documents in our human audit, we found that 86 percent of articles measure high geopolitical risk, whereas only 4.3 percent of these

articles measure declining tensions. We therefore conclude that increases in GPR are far more likely to lead to the use of our preferred search terms.

2. The GPR index is a frequency, thus satisfying data check 2 (*Data check 2. Is the variable being proxied a frequency?*).

3. We verify that the average number of documents found is large enough for variation to be driven by factors other than sampling error (*Data check 4: Is the average number of documents found large enough [...]?*). In particular, we verify that spikes in GPR are easily attributable to well-defined historical events not only at monthly, but also at daily frequency. For instance, the monthly data show that our index spikes in April 1986, mostly following the events that culminated with U.S. air strikes against Libya on April 15. However, the index also spikes, within the month, on April 8, when the U.S. accused Muammar el-Qaddafi of sponsoring terrorist acts aimed at Americans (such as the Berlin discotheque bombing which occurred on April 5). It also spikes on April 18, when British police found a bomb in a bag that was taken onto an El Al aircraft.

4. We verify that measurement error is low enough (*Data check 3*, and *Data check 5: Is the expected variance in the occurrence-frequency of interest high enough to overcome the noise associated with document-frequency proxying?*), by choosing combinations of search terms that – unlike with a single keyword or a bigram – are unlikely to be used outside of the realm of rising geopolitical risk . For instance, the naive geopolitical risk index GPRE is nearly as high in March 1991 as in January 1991, whereas the benchmark GPR index is four times as low. This occurs because the naive index fails to account for the fact that many articles comment the aftermath of the Gulf war, but do not explicitly mention threats or risks, something that our index takes into account.

5. We construct and examine narrow and broad indexes around the benchmark index, thus satisfying data check number 5.

6. Because our interest as economists is to look at the economic effects of higher geopolitical risk, we need to guard against the possibility that words related to geopolitical tensions are more likely to be mechanically used during, say, recessions, even if recessions are caused by geopolitical tensions. To address this concern, we construct a version of the GPR index that excludes the search terms "economy" OR "stock market*" OR "financial market*" OR "stock price*". Although 19 percent of articles in GPR are filtered out once this criterion is included, the resulting index (GPR-NOECON) is virtually indistinguishable from the benchmark index (their correlation is 0.989). Additionally, the share of articles in the GPR index mentioning economic words is uncorrelated

with real outcomes (the correlation with the change in log industrial production is 0.03), thus allaying concerns that an omitted variable biases the results both "geopolitical risk" and economic outcomes (*Data check 6: [...] Does the chosen keyword have as its primary or only meaning the occurrence of the phenomenon of interest?*, and *Data check 7: [...] Does the chosen keyword also result in documents related to the covariates of the occurrence of interest?*).

7. We use the naive GPRE index as a placebo document-frequency variable in our VAR analysis. In particular, there is the possibility that it is not geopolitical risks per se that are bad, but the overall tendency to discuss geopolitical events during recessions. We show that adding GPRE to the VAR does not change the predictive power of GPR in explaining GDP drops. (*Data check 8: Are there plausible omitted variables that may be correlated both with the document-frequency and its covariates?*)

Figure A.1: Search Query for the Benchmark GPR Index

```
pub.Exact("Boston Globe" OR "Chicago Tribune" OR "The Daily Telegraph" OR "Financial
Times" OR "The Globe and Mail" OR "The Guardian" OR "Los Angeles Times" OR "New York
Times" OR "The Times" OR "Wall Street Journal" OR "The Washington Post") AND
DTYPE(article OR commentary OR editorial OR feature OR front page article OR front
page/cover story OR news OR report OR review) AND (("United States" AND tensions AND
(military OR war OR geopolitical OR coup OR guerrilla OR warfare) AND ("Latin America"
OR "Central America" OR "South America" OR Europe OR (Africa NOT "South Africa") OR
"Middle East" OR "Far East" OR Asia)) OR (geopolitical AND (risk* OR concern* OR
tension* OR uncertaint*)) OR (("nuclear war" OR "atomic war" OR "nuclear conflict" OR
"atomic conflict" OR "nuclear missile*") AND (fear* OR threat* OR risk* OR peril* OR
menace*)) OR ("war risk*" OR "risk* of war" OR "fear of war" OR "war fear*" OR
"military threat*" OR "war threat*" OR "threat of war" OR (("military action" OR
"military operation" OR "military force") AND (risk* OR threat*))) OR ("terrorist
threat" OR "terrorist threats" OR "menace of terrorism" OR "terrorism menace" OR
"threat of terrorism" OR "terrorist risk" OR "terror risk" OR "risk of terrorism" OR
"terror threat" OR "terror threats") OR ("beginning of the war" OR "outbreak of the
war" OR "onset of the war" OR "escalation of the war" OR "start of the war" OR ((war
OR military) AND "air strike") OR (war AND "heavy casualties") OR (battle AND "heavy
casualties")) OR ("terrorist act" OR "terrorist acts") NOT ("civil war" OR "human
rights" OR (end N/2 war) OR "air force" OR movie OR film OR museum OR anniversary OR
memorial OR art))
```

Note: The index is the share of number of articles returned by the search above, divided by the total number of articles, indexed to 100 in the 2000–2009 decade.

Figure A.2: Share of GPR Articles by Individual Newspaper

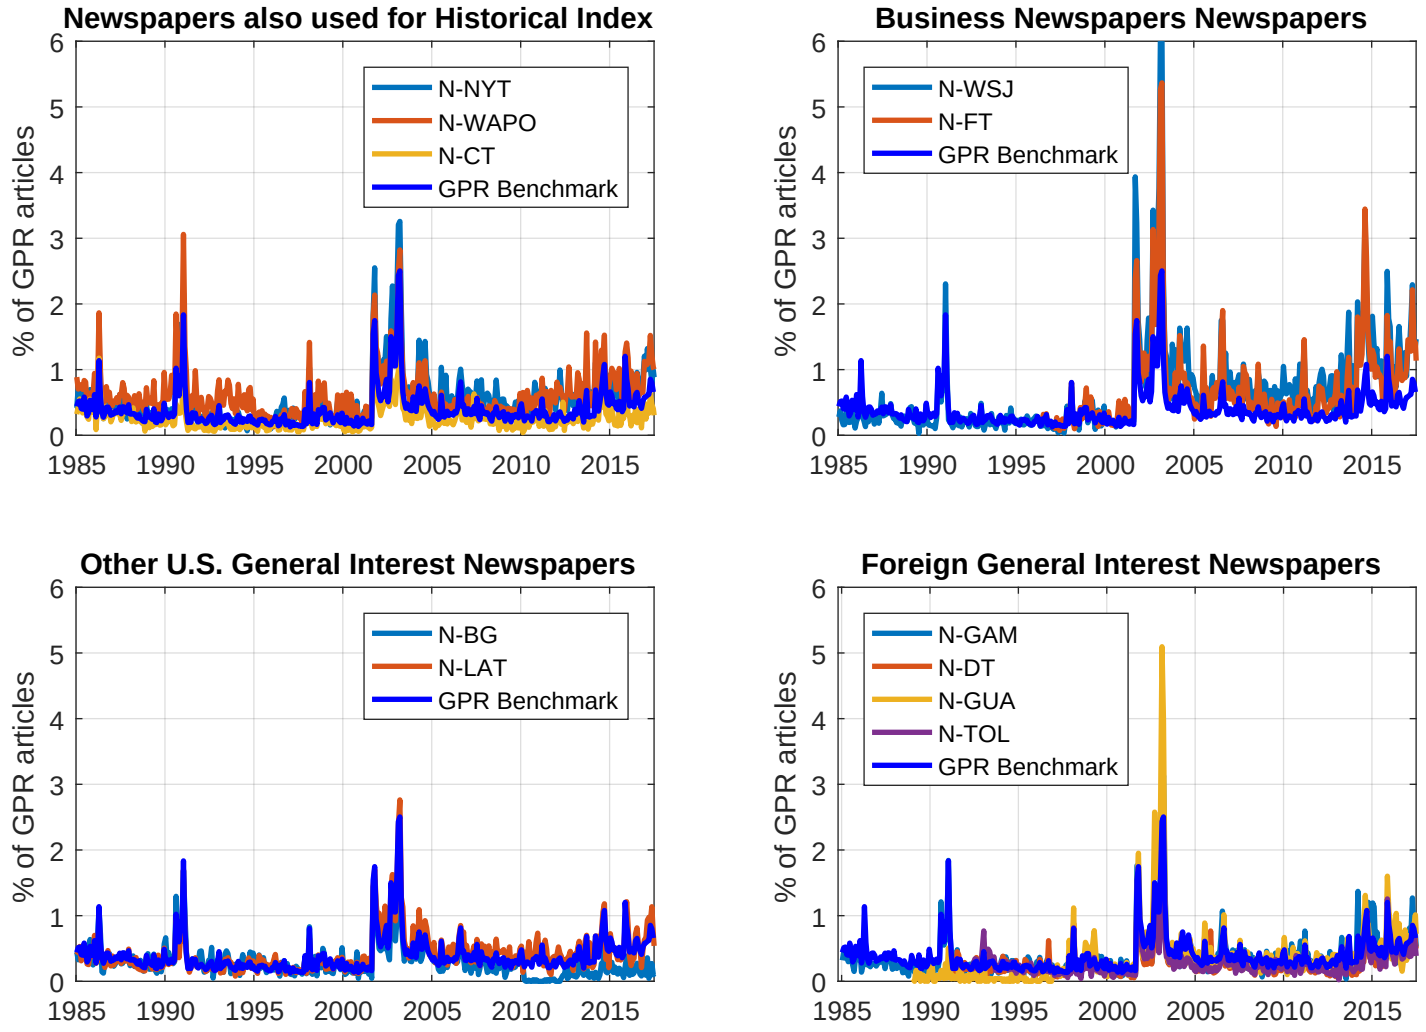

Note: Each panel plots the share of articles containing words related to geopolitical risk for each of the 11 newspapers.

Figure A.3: GPR and Natural Disasters Index

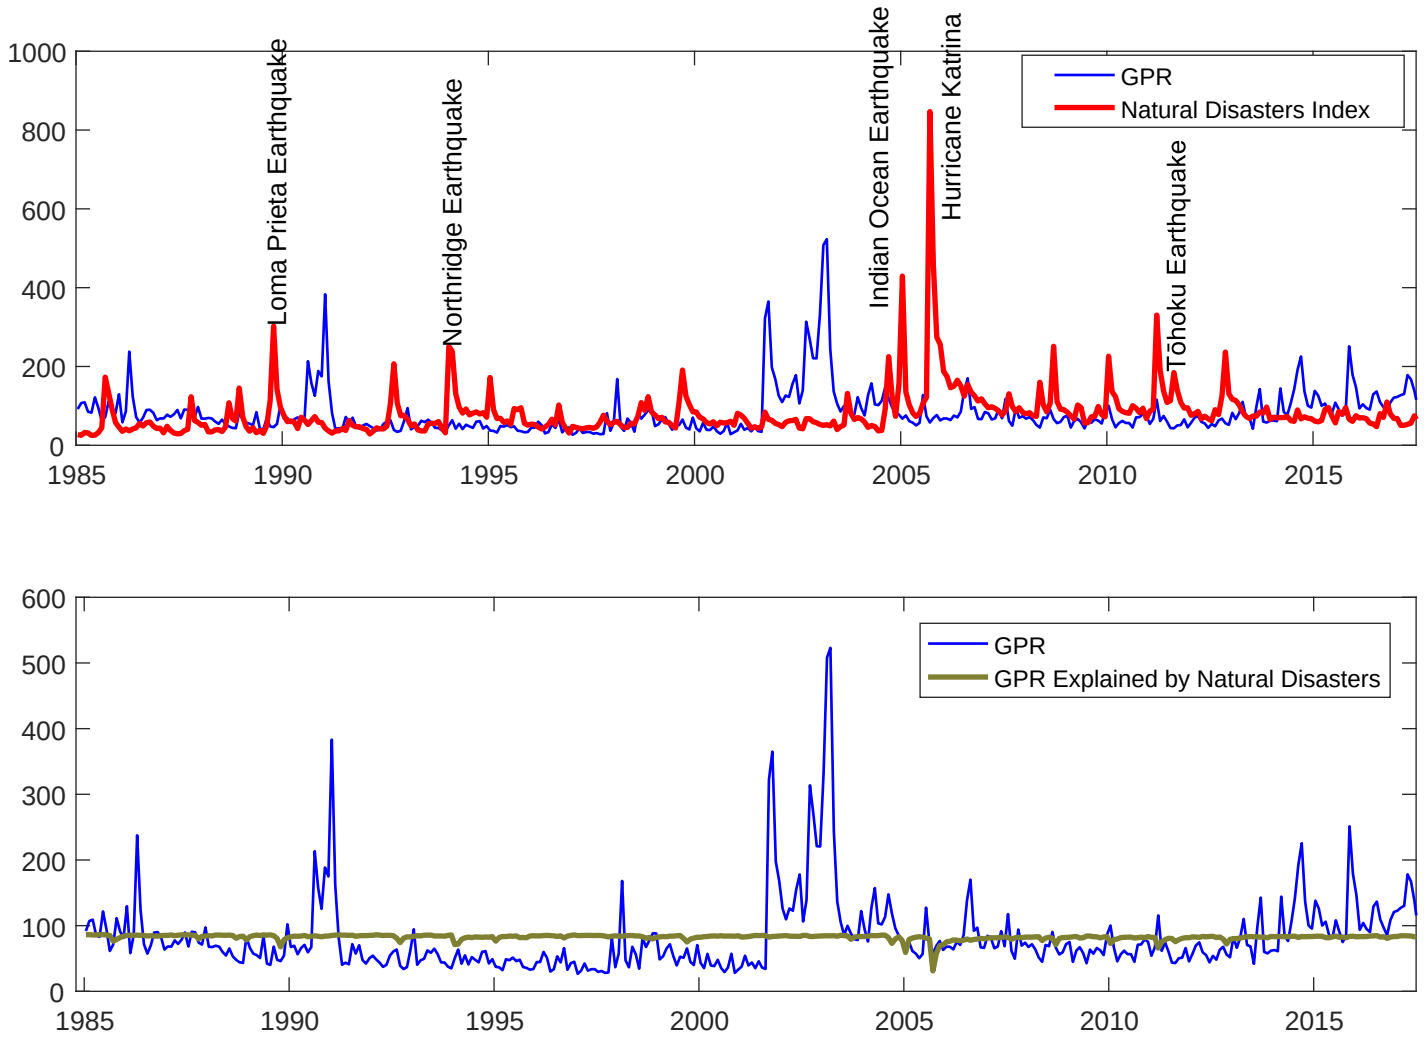

Note: The top panel of the figure compares the benchmark GPR index with a news-based index of natural disasters. The bottom panel of the figure shows the GPR together with the fitted value of a regression of the GPR against a constant and the Natural Disasters Index. If the GPR is driven not by actual events but by shifting media attention towards geopolitical events, one would expect that natural disasters move the GPR, even absent actual movements in underlying geopolitical tensions. However, as shown by the figure, the GPR hardly moves in response to natural disasters.

Figure A.4: Human and Computer-Generated Historical GPR Indexes

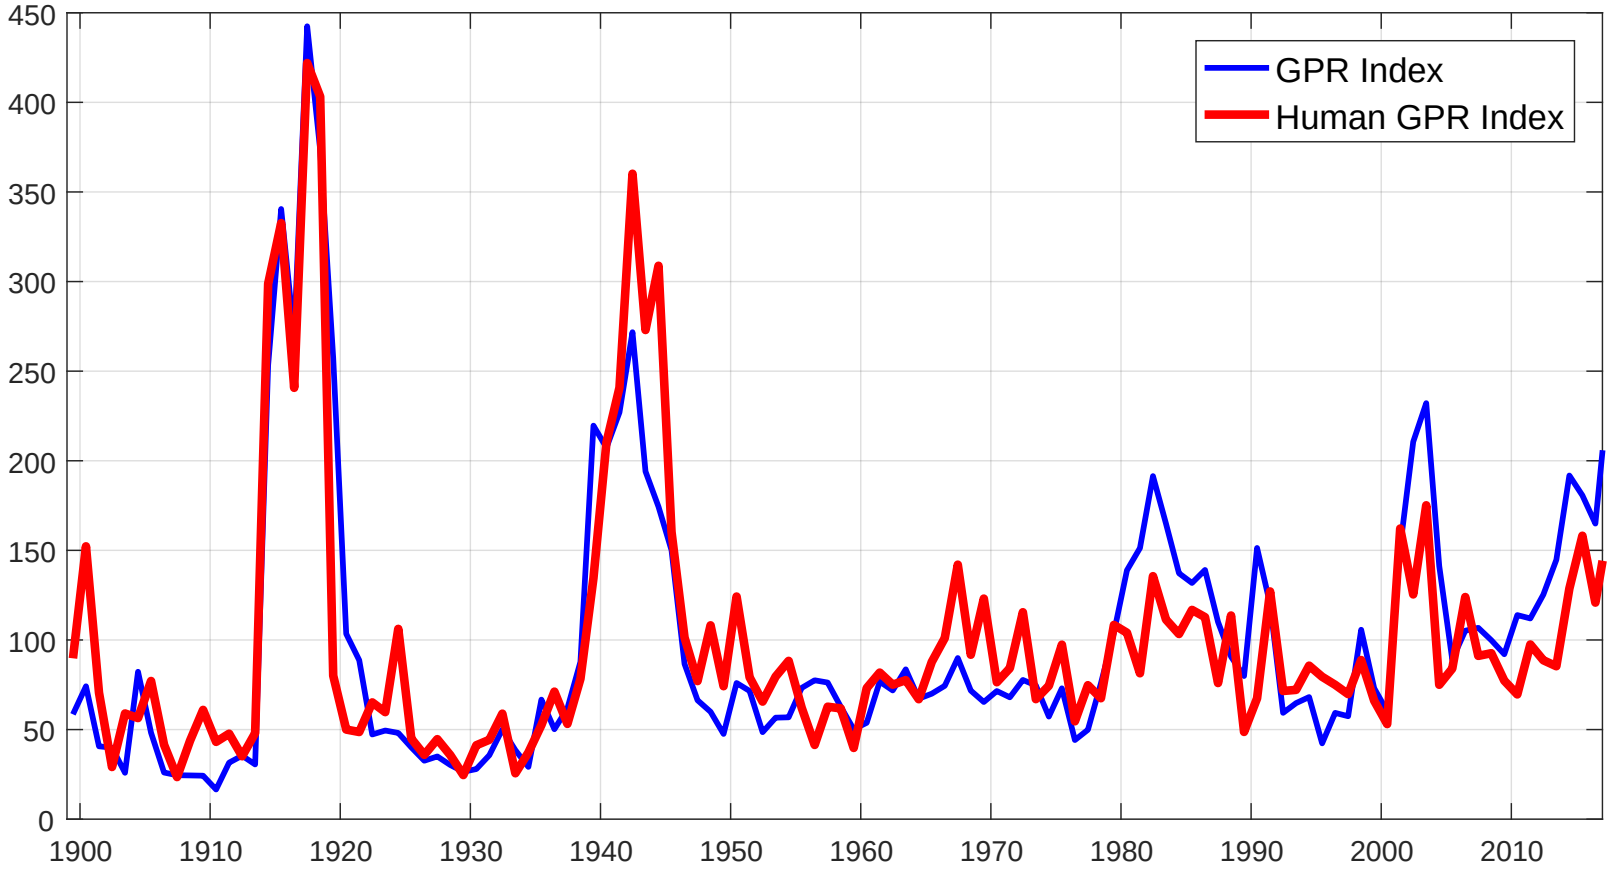

Note: Time-series comparison from 1899 to 2016 based on 7,416 articles. The series are plotted annually to reduce sampling variability. Both series are normalized to 100 throughout the sample.

Figure A.5: GPR, GPRWARMINUSPEACE and GPR Naive

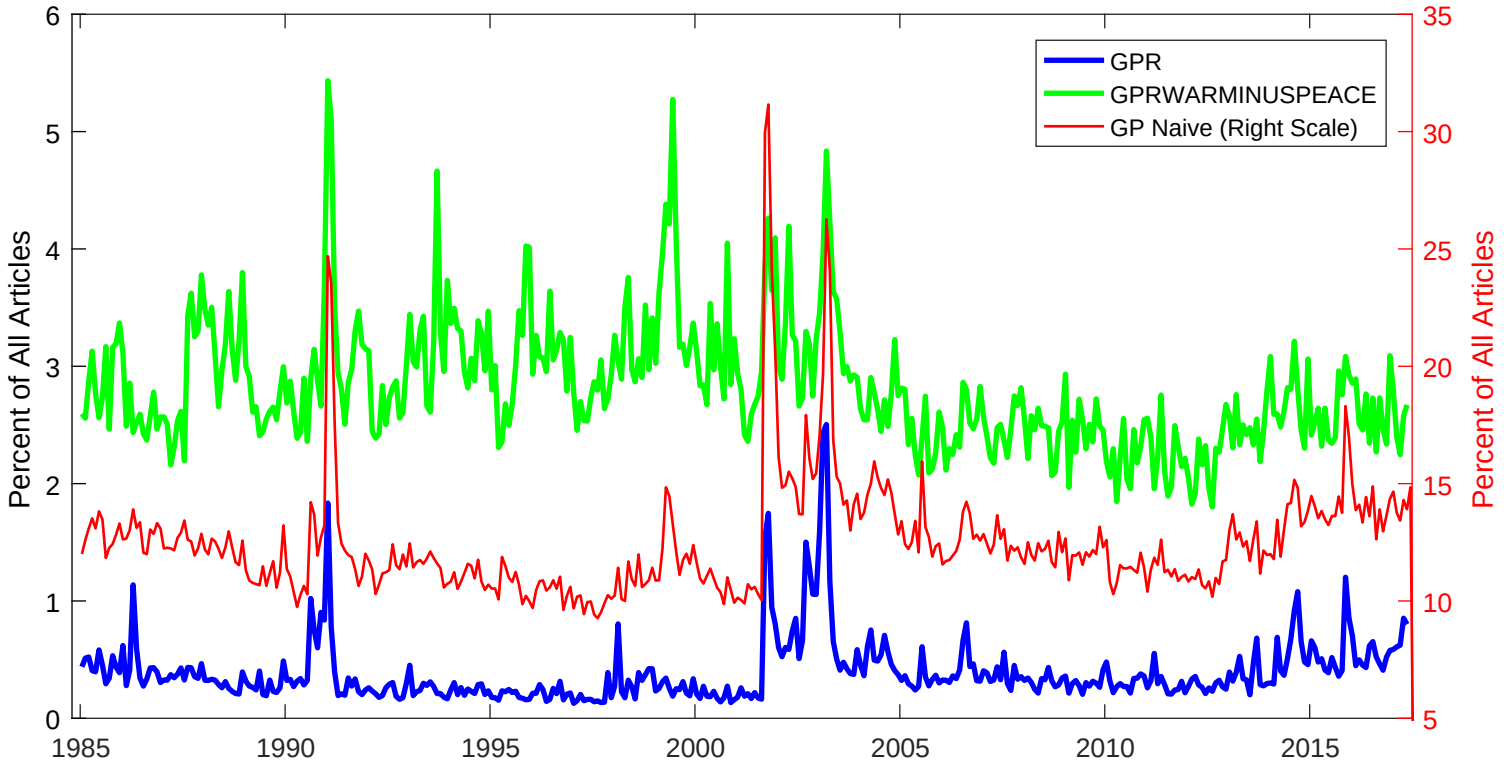

Note: This figure compares the benchmark GPR index with the naive GPRE index constructed counting the relative frequency of all articles mentioning any of the four keywords that enter the set $\mathcal{E}$, as well as with a GPRWARMINUSPEACE index subtracting the word peace from the occurrence of the words war and terrorism. Note that the indexes are expressed as a share of the total number of articles.

Figure A.6: The Impact of Increased Geopolitical Risk: Adding EPU to the Baseline VAR

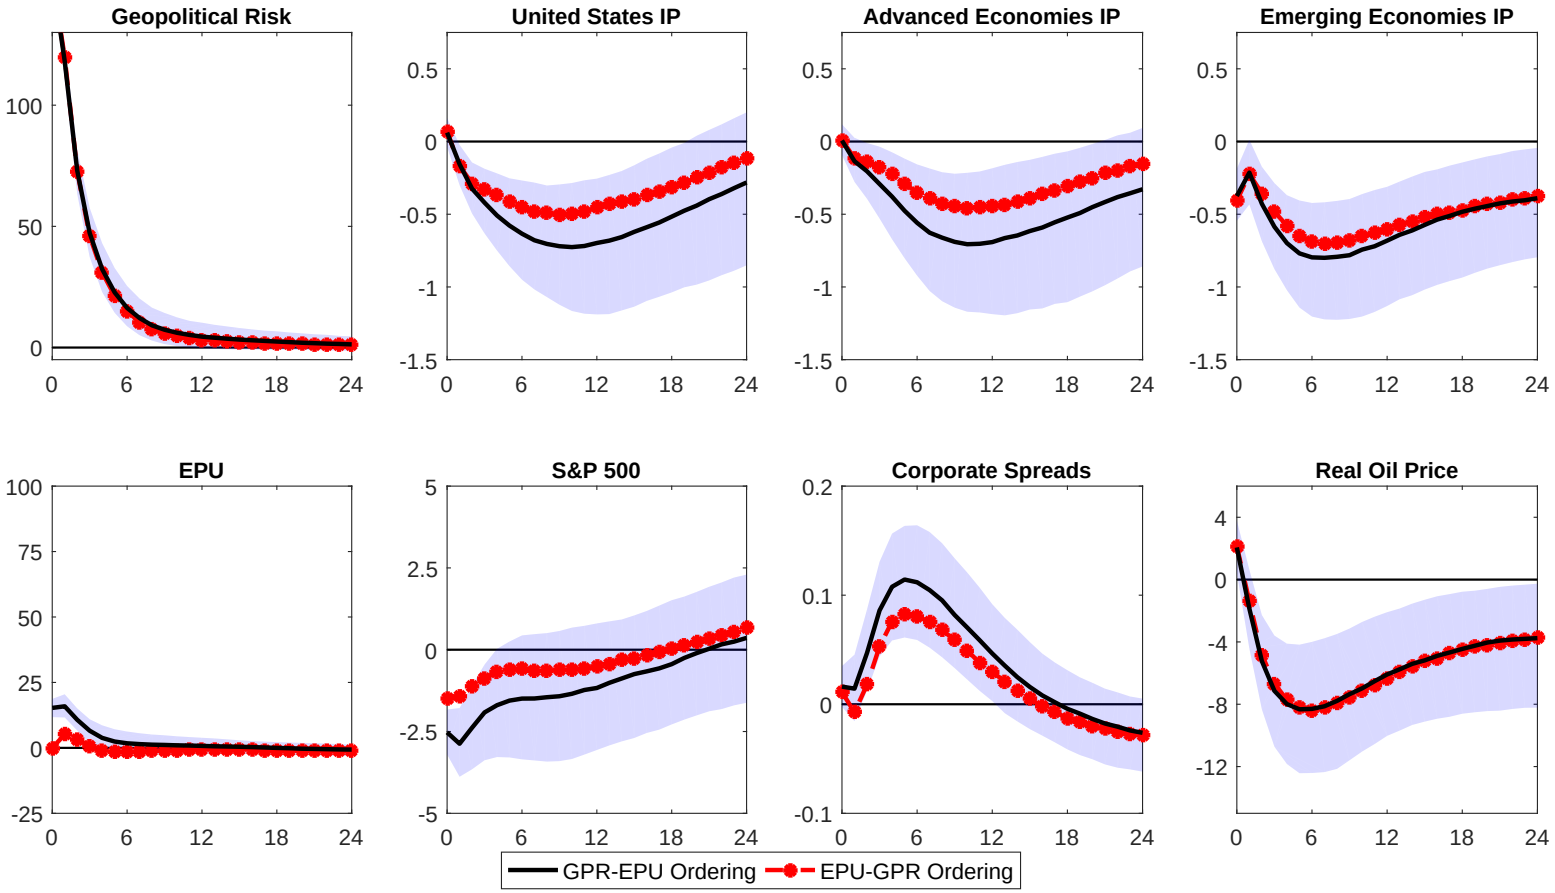

Note: The black solid line in each panel depicts the median impulse response of the specified variable to a rise of 150 points in the GPR index, while the shaded bands represent the 68 percent pointwise credible set. The horizontal axis measures months since the shock.