Climate Risk and Commodity Currencies*

Felix Kapfhammer† Vegard H. Larsen‡ Leif Anders Thorsrud§

This version: November 30, 2021

Abstract

Climate change increases the likelihood of extreme climate- and weather-related events, but also the pressure to adjust to a lower-carbon economy. We propose a novel measure of climate change transition risk and document that when it unexpectedly increases, major commodity currencies experience a persistent depreciation in line with traditional "Dutch disease" arguments. Furthermore, when expanding the analysis to a richer set of countries we find a significant negative correlation between a country's fossil fuel export dependency and exchange rate response following innovations in transition risk. None of these findings apply when existing climate risk proxies are used, suggesting that studies not distinguishing between different climate risk components might misinterpret the economic consequences of climate change.

JEL-codes: C22, C82, E71, F31, Q54

Keywords: Exchange Rates, Climate, Risk, Commodities

*This Working Paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. Comments from Jørgen J. Andersen, Hilde C. Bjørnland, Gunnar Bårdsen, Ragnar Torvik, and one anonymous referee improved the quality of this paper considerably. We also thank Farooq Akram, Georgia Bell, Leif Brubakk, Jon Ellingsen, Solveig K. Erlandsen, Leif Helland, and Espen R. Moen for valuable comments. Comments from conference participants at BI Norwegian Business School and Norges Bank also helped improve the paper. This work is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School. We are grateful to the *Dow Jones Newswires Archive* for sharing their data with us for this research project. Declarations of interest: none.

†Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School. Email: `felix.kapfhammer@bi.no`

‡Norges Bank and Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School. Email: `Vegard-Hoghaug.Larsen@norges-bank.no`

§Corresponding author. Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. Email: `leif.a.thorsrud@bi.no`

1 Introduction

The economic risks posed by climate change can be decomposed into at least three components (Carney, 2015): *physical risk* arising from climate- and weather-related events; *liability risk* arising if losses due to climate change are legally pursued with compensation demanded; *transition risk* resulting from adjustments towards a lower-carbon economy. Perhaps due to ease of measurement, the majority of empirical literature in finance and economics on the topic has mostly concentrated on the first of these components. Dell et al. (2014), Burke et al. (2015), Hong et al. (2019), and Choi et al. (2020) provide prominent and recent examples.

In this article we propose a novel measure of climate change transition risk and analyze how fluctuations in such risk affect fossil-intensive commodity currencies. This relationship has not been explored in the literature before but is interesting because it addresses how climate risk affects valuations at a national level, and because it can be argued that transition risk, as opposed to the other two climate risk components, is especially relevant for exchange rate developments in countries producing carbon-intensive commodities.

In particular, standard and well-known theories on positive changes in natural resource income predict both *spending* and *resource movement* effects and a persistent appreciation of the real exchange rate in response to resource gifts (see Bruno and Sachs (1982), Eastwood and Venables (1982), Corden (1984), and van Wijnbergen (1984) for early examples, and Torvik (2001), Charnavoki and Dolado (2014), Bjørnland and Thorsrud (2016), Arezki et al. (2017), and Bjørnland et al. (2019) for newer contributions to this literature). Traditionally, these gifts have been looked upon as either favorable shifts in the production function of the (booming) commodity-producing sector or discoveries of new resources. Climate change transition risk accommodates both, but as risks with the opposite sign.[1] For example, today it is well known that carbon budgets compatible with conventional temperature targets imply that new investments in high-carbon capital should be rapidly discontinued and that existing production technologies must be scaled down, retired before their "natural" end of life, or retrofitted at a cost (Campiglio and der Ploeg, 2021). However, the actual implementation and public support related to such changes are subject to uncertainty. As such, when the likelihood of change increases, i.e.,

[1] This does not rule out that high climate risk is associated with expectations of permanently lower commodity prices. As long as production technology has diminishing returns to scale, a persistent depreciation of the exchange rate is a common feature in theoretical models containing a reduction in natural resource income. Under the assumption of constant returns to scale in production, however, the equilibrium exchange rate will typically be determined only by the supply side of the economy, and commodity income does not matter (Rogoff and Obstfeld, 1996, Chap. 4). Still, even in this setting, transitional dynamics imply a real exchange rate depreciation, and the return to an equilibrium might take a very long time.

increased transition risk, the exchange rate should depreciate already today due to the forward-looking nature of foreign exchange markets.[2]

The challenge is how to measure transition risk. Unlike physical risk measures related to climate- or weather-related events, transition risk is not directly observed and must be proxied. To this end we contribute to the broader finance and economic climate literature by proposing a methodology for measuring this type of climate change risk.

As a starting point we observe that the scientific discussion about climate change and the statistical evidence documenting it dates back several decades (Arrhenius (1896), Keeling (1970), Nordhaus (1977)), but that public awareness of climate change and its potential economic consequences seem to be of a much newer date. For this reason we share the view taken in, e.g, Nimark and Pitschner (2019), Larsen et al. (2020), and ter Ellen et al. (2020), where the media operate as "information intermediaries" between agents and the state of the world, and use news media coverage as a proxy for changing perceptions of transition risk in the public discourse. Indeed, a number of studies have documented that mass media coverage increases public awareness about environmental issues (Schoenfeld et al. (1979), P. (1986), Boykoff and Boykoff (2007), Sampei and Aoyagi-Usui (2009), Hale (2010)). This coverage naturally includes changes in the discussion and implementation of actual policies and changes in investor and consumer behavior, but also more silent features related to systematic directional modification of ideas and narratives as they are spread in the public discourse (Shiller (2017), Hirshleifer (2020)). As stated by Shiller (2001): "significant market events generally occur only if there is similar thinking among large groups of people, and the news media are essential vehicles for the spread of ideas".

Our underlying hypothesis is simple: When the association in media coverage between a given country and talk about the process of adjusting towards a lower-carbon economy is high, it signals transition risk that leads to a persistent depreciation of commodity currencies in line with the "Dutch disease" mechanism discussed above.

We operationalize the hypothesis using a unique and large corpus, i.e., text from over 20 million articles, of international business news provided by the *Dow Jones Newswires Archive* (DJ). This data is partitioned into monthly blocks and a neural network is used to construct word embeddings for each month in the dataset. Word embeddings represent words in vector space, and have, following the seminal contributions of Mikolov et al. (2013) and Mikolov et al. (2013), become a much-used tool in the Natural Language Pro-

[2] In Norway, a major fossil fuel commodity exporter, Norges Bank's reflections on recent exchange rate developments in their third Monetary Policy Report in 2019 are consistent with this type of reasoning: " The krone has been weaker for some time than projected in the Monetary Policy Report. [...] Prospects for lower activity in the petroleum sector and uncertainty about the need for restructuring in the Norwegian economy may also have weighed on the krone."

cessing (NLP) and Machine Learning (ML) literature. The reason is that they densely encode many linguistic regularities and patterns and allow for arithmetic operations capturing associative meaning. Accordingly, for each month in the sample, we derive the weighted sum of word vectors representing economic risk and climate change dimensions and regress these on word vectors for each country. The parameter estimates of these regressions measure how strong the association between a given country and climate risk is and how it varies across time.

To address the relationship between transition risk and real exchange rates we use Vector Autoregressive models and focus on the commonly used commodity currencies of Australia, Brazil, Canada, Malaysia, Mexico, Norway, Russia, and South Africa.[3] Our main results can be summarized as follows: First, taking into account the dynamic interactions between, e.g., commodity prices, business cycle developments, currencies, and climate risk, shows that transition risk is generally not significantly affected by the other variables in the system, whereas exogenous transition risk innovations generally lead to a significant and persistent exchange rate depreciation. On average, roughly five percent of the long-run variation in the real exchange rates can be attributed to transition risk innovations. Second, including a rich set of other countries in the analysis shows that there is a significant negative correlation between a country's fossil fuel export dependency and the real exchange rate response following innovations in transition risk. In line with the conjecture that transition risk should be particularly relevant for commodity currencies, the exchange rate responses for non-commodity currencies as a group are not significantly different from zero on average.

At an intuitive level, our measure of risk is intended to capture how transition risk is spoken about in the public discourse, but not necessarily how much it is talked about. For example, important climate events, such as the Paris Agreement, increase the frequency of climate-related words used in the press (Engle et al., 2020), but do not necessarily change how these words are used in context with other words and how different countries are written about in the press. In contrast, time variation in our suggested measures can be due to the fact that climate change-related words are used more or less frequently in relation to a country, or because transition risk-related words are used differently across time. A word such as "risk" might, for example, be used much more in relation to the word "climate" today than a decade ago. Likewise, a word like "green" might be more closely associated with "economics" now than before. The word embedding methodology we apply is designed to capture exactly these types of changes.

[3] As seen in Figure B.1, in Appendix B, these countries produce a substantial amount of carbon-intensive commodities and are thus particularly relevant from a transition risk perspective. Major (fossil fuel) commodity exporters that do not have floating exchange rates have been left out of the analysis.

Additional results indicate that the proposed transition risk measures add value. Comparing our measure with a number of commonly used alternatives, such as temperature anomalies, a news-based climate risk index recently proposed by Engle et al. (2020), or so-called Climate Change Performance Indexes, yield either insignificant or theory-inconsistent results in the commodity currency setting.

Finally, one implication of the underlying mechanism we build on is that increases in transition risk today should be associated with a future reduction in the activity level in the commodity sector, i.e., reduced supply. Because natural resource income is an important part of aggregate income creation in major commodity exporters, another implication is that aggregate stock market developments might become negatively affected as well. Corroborative results support both of these predictions. We find a negative relationship between country-specific commodity supply and unexpected increases in transition risk. We also document that an unexpected increase in transition risk tends to cause persistently lower aggregate stock market valuations.

These results have practical importance for policy makers, as exemplified by the quote above, and contribute to three different strands of the economic literature. First, our study speaks to a large literature on the macroeconomic consequences of natural resource income, with influential work already cited, but extends this line of research by taking into account transition risk resulting from adjustments towards a lower-carbon economy. Such risk, however, is also relevant more generally. For example, as recently illustrated by Carattini et al. (2021) using a U.S. calibrated E-DSGE model, transition risk can lead to macroeconomic instability and even recessions.[4]

Second, our study speaks directly to a growing literature on the pricing implications of climate risk and its components. Thus far, however, most of this literature has been concerned with pricing of firms and firm value.[5] In relation to commodity producers, the growing literature on stranded assets is closely connected to the transition risk concept (Ramelli et al. (2018), Atanasova and Schwartz (2019), van der Ploeg and Rezai (2020), Sen and von Schickfus (2020)). Atanasova and Schwartz (2019), for example, find that the

[4]The authors study a closed economy. In the model, transition risk innovations lead to adverse macroeconomic outcomes foremost because financial frictions slow down the transition from "brown" to "green" assets. An initial high "brown" asset share will amplify the effects and vice versa. See Carattini et al. (2021), and the references therein, for related work using environmentally augmented DSGE models (E-DSGE).

[5]See, e.g., Bolton and Kacperczyk (2020), Hsu et al. (2020), Freeman et al. (2015), Daniel et al. (2019), Batten et al. (2016), Andersson et al. (2016), In et al. (2017), and Krueger et al. (2020). The recent study by Cha et al. (2020) shares our focus on the foreign exchange market. They conduct an exploratory analysis of the responses of monthly U.S. dollar real exchange rates of 76 countries to global temperature shocks, i.e., physical climate risk, and find significant responses for roughly half of the countries in the sample.

growth of firms' fossil fuel reserves now has a negative effect on firm value, suggesting that capital markets treat fossil fuel as "stranded assets" in the transition to a lower-carbon economy. Thus, just as stranded assets might affect firms' value negatively because of climate risk, our results imply that this risk also negatively affects the pricing of exchange rates (and aggregate stock markets) in countries where natural resource income is important.

Third, this article speaks to a growing literature using tools from NLP and ML to facilitate and improve measurement in economics and other social sciences. For example, Kozlowski et al. (2019) use word embeddings to produce richer insights into cultural associations and categories than possible with existing methods in the field of sociology, while Thorsrud (2018), Larsen and Thorsrud (2019), Baker et al. (2016), and Hansen et al. (2018) use text as data to measure business cycle developments, uncertainty, and monetary policy. In particular, by focusing on climate change, this article relates to Engle et al. (2020) who propose a news-based climate risk measure for dynamically hedging climate change risk. However, their index does not directly distinguish between the three different types of climate change risks, and essentially measures *how much* climate change is focused upon in the news. In contrast, our word embedding approach measures in *what context* it is focused upon and aims to capture country-specific transition risks. As discussed above, these differences are important for describing exchange rate fluctuations.

The rest of this paper is organized as follows: Section 2 presents the textual data, the word embedding methodology, and the proposed transition risk measures. Section 3 describes the exchange rate modeling framework and presents the main results regarding transition risk and commodity currencies, while Section 4 expands the analysis by including either a broad set of non-commodity currencies or using alternative climate risk proxies. Section 5 concludes.

2 Climate risk and measurement

Below we describe the DJ corpus and how we apply a word embedding model to construct quantitative and country-specific climate change transition risk measures.

2.1 News coverage and word embeddings

The DJ corpus consists of roughly 23 million news articles, written in English, covering the period 2001 to 2019. The database contains a large range of Dow Jones' news services, including content from *The Wall Street Journal*. Arguably, the DJ does not fully reflect the public discourse. Still, news stories relevant for investors and agents in the international foreign exchange market are undoubtedly well covered by this type of business news. The *Dow Jones* company's flagship publication *The Wall Street Journal* is also one of the largest newspapers in the U.S. in terms of circulation. This means that it has a large footprint in both the U.S. and global media landscape and that important ongoing stories and discussions are well covered by this type of news outlet.

The news corpus is cleaned prior to estimation. We remove all email and web addresses, numbers, and special characters, erase punctuation, set all letters to lowercase, and remove words containing fewer than two or more than 15 letters. These feature selection steps reduce the size of the vocabulary to approximately 90000 unique terms. The dimension reduction facilitates estimation and is common in the literature. Finally, the corpus is partitioned into monthly blocks of articles. Each month of data contains between 42000 (2005M2) and 115000 (2013M3) articles.

The monthly blocks of data are used as input in a word embedding model. Word embedding models represent words as relatively small and dense vectors. The famous and widely used word2vec algorithm (Mikolov et al. (2013), Mikolov et al. (2013)) is one of many algorithms used to compute such vectors and is often denoted as a skip-gram model with negative sampling. In essence, the method uses a binary classification problem, asking "is the context word *co* likely to show up near the target word *ta*?", as a vehicle to compute the classifier weights that will be the actual word embeddings.

In our setting, this approach has two particularly appealing features. First, estimated word embeddings encode many linguistic regularities and patterns and allow for arithmetic operations that can capture associative meaning. A famous example is "king" − "man" + "woman" ≈ "queen", where the word vector "king" and the difference between "woman" and "man" pulls the resulting vector in the royal and feminine directions, respectively, with the end product being close to the actual vector for the word "queen". For our purpose the royal and feminine dimensions are not relevant, but capturing the associative meaning of words that, taken together, point in the (latent) transition risk dimension is. As such, and as explained in greater detail in Section 2.2, given a set of words of interest, using word embeddings together with arithmetic operations is a potentially well suited tool.

Second, when estimating word embedding models, running text can be used as implicit supervised training. This avoids the need for any sort of hand-labeled supervision signal and makes the methodology flexible and user friendly in many different contexts. In contrast, popular NLP methods such as the Latent Dirichlet Allocation topic model, which has been applied in a number of recent economic studies (see, e.g., Larsen and Thorsrud (2019) and Hansen et al. (2018)), is a completely unsupervised methodology where the user needs to define the meaning of the estimated topics ex-post. Similarly, applying the commonly used word count approach (see Gentzkow et al. (2019) and the references therein) would not allow the researcher to capture the associate meaning between words and how this might change over time.

More formally, given a target word ta and a context word co, the probability that the word co is a real context word for ta is $P(+|ta, co)$ and that it is not a real context word is $P(-|ta, co) = 1 - P(+|ta, co)$. The intuition for the skip-gram model is then that a word is likely to occur near the target if its embedding is similar to the target embedding, where similarity is approximated by the dot product of the word vectors for co and ta. This yields the likelihood

$$L(\theta) = \sum_{(ta,co)\in +} P(+|ta, co) + \sum_{(ta,co)\in -} P(-|ta, co), \tag{1}$$

which for one target/context word pair (ta, co) can be written as:

$$L(\theta) = log\frac{1}{1 + e^{-co \cdot ta}} + \sum_{i=1}^{k} log\frac{1}{1 + e^{n_i \cdot ta}}, \tag{2}$$

where k denotes the context window for which the co words occur relative to the target word ta, and the logistic (or sigmoid) function is used to turn the similarity measure between the word vectors for co and ta into probabilities. The last term in (2) relates to the negative sampling part of the skip-gram model name. As running text is used as input to the model, only positive examples are present and negative examples need to be generated and added to the data. These terms are commonly called noise terms (n_i). For each target word, it is common to add k noise words.

Maximizing (2), and learning the latent word vectors, can be achieved using different methods. Here we use a simple two-layered neural network. This method is fast, efficient to train, and available in many software packages. We set the context window $k = 5$, we restrict the word embedding length $d = 100$, and the network is trained for five epochs on every monthly partition of the data.[6]

2.2 Word embeddings and climate risk

In the context of "Dutch disease economics", we think about the process of adjustment towards a lower-carbon economy as a "concern related to the fossil fuel producing sector of the economy" due to the adoption of "policies and behavior reducing environmental

[6]Estimation of word embeddings implicitly uses information embodied in the so-called Pointwise Mutual Information (PMI) matrix, which contains the likelihood of words co-occurring in the corpus. While word embeddings derived from neural networks are unknown to many economists, a potentially more familiar version would be to extract principal component estimates from the corpus' PMI matrix and let the factor loadings represent the word embeddings. However, this simpler method enforces a more restrictive functional form, and does not scale easily to a large corpus.

and climate impact". To provide a quantitative measure encapsulating this definition we use the linguistic regularities and patterns encoded in the estimated word embeddings together with arithmetic operations. The intuition for this approach is very much the same as in the royal example above.

More precisely, we first define five word-based categories representing economic risk and climate change dimensions. This is illustrated in Table 1. The sum of the *concern*, *fossil fuels*, and *economy* categories results in a vector intended to point in a direction encompassing "concern related to the fossil fuel producing sector of the economy". I.e, we want to capture economic concerns and not all other types of concerns. And we want those economic concerns to be related to fossil fuel production. The terms $climate^+$ and $climate^-$ add the climate change dimension, encompassing "policies and behavior reducing environmental and climate impact". "reducing" in this definition is captured by looking at $climate^+ - climate^-$. If the words defining $climate^+$ are used more in association with the economic risk dimension than the words defining $climate^-$, we think of this as pushing the aggregated word embedding into a more climate friendly part of the vector space. Finally, so as not to over-weight the economic risk dimension at the expense of the climate change dimension, the two dimensions are given equal weight before adding them together.

To capture the monthly association between countries and the word vector representing an economic climate change risk dimension, i.e., transition risk, we solve

$$TR_t \equiv \hat{\beta}_t = \arg\min S(\beta_t) \quad \text{and} \quad S(\beta_t) = \left\| country_t - transition\ risk_t \times \beta_t \right\|^2, \tag{3}$$

where the word vector for $country_t$ is given in Table 1, and β_t is the association between country c and climate risk. Although $\hat{\beta}_t$ is estimated using the OLS estimator on each monthly partition of the sample, the subscript t is used to highlight that this relationship potentially changes across time.[7]

An increase in β_t means that transition risk increases because the country is becoming more associated with concerns about the process of adjusting towards a lower-carbon economy. While such an increase might be good for the climate itself, economic theory predicts it will put downward pressure on commodity currencies because it implies reduced activity in the commodity-producing sector of the economy.

We emphasize three points about this construction. First, because of differences in policies, public perception, and consumer and investor behavior across countries, the

[7] Kozlowski et al. (2019) apply a similar approach to uncover changes in cultural associations and categories. For example, to determine the gender association for the word "tennis", they project the word embedding for this word onto the gender dimension of, e.g., "man" – "woman", and document how the resulting projection changes through time, where a more negative projection coefficient implies a more feminine association.

Table 1. Constructing transition risk indexes from word embeddings. The upper part of the table reports the key dimensions of the transition risk definition used in this article. Category names are printed in bold and the associated words (i.e., word vectors) are listed in the right side of the table. The lower part of the table reports the words (word vectors) used to define each country. In the case of South Africa, the corpus has been cleaned prior to estimation by joining terms, e.g., instead of representing "South Africa" as a bi-gram it is collapsed to one token "SouthAfrica".

Key dimensions		Words
$\boldsymbol{concern_t}$	=	$\frac{1}{n_1}(concern_t + concerned_t + risk_t + risky_t + uncertain_t + worried_t + worrying_t)$
$\boldsymbol{fossil\ fuel_t}$	=	$\frac{1}{n_2}(extract_t + mine_t + fossil_t + fuels_t + fuel_t + oil_t + crude_t + petroleum_t + coal_t + lignite_t)$
$\boldsymbol{economy_t}$	=	$\frac{1}{n_3}(economy_t + economic_t + economics_t + business_t + sector_t + sectors_t)$
$\boldsymbol{climate_t^+}$	=	$\frac{1}{n_4}(climate_t + green_t + clean_t + renewable_t + oxygen_t + recycling_t + ecosystem_t + cooling_t + protect_t)$
$\boldsymbol{climate_t^-}$	=	$\frac{1}{n_5}(emissions_t + dirty_t + fossil_t + dioxide_t + methane_t + pollution_t + warming_t + exploit_t)$

$$transition\ risk_t \approx \frac{1}{2}\underbrace{(concern_t + fossil\ fuel_t + economy_t)}_{\text{economic risk dimension}} + \frac{1}{2}\underbrace{(climate_t^+ - climate_t^-)}_{\text{climate change dimension}}$$

Countries ($country_t$)		
Norway	=	$\frac{1}{n}(norway_t + norwegian_t)$
Mexico	=	$\frac{1}{n}(mexico_t + mexican_t)$
Malaysia	=	$\frac{1}{n}(malaysia_t + malaysian_t)$
Canada	=	$\frac{1}{n}(canada_t + canadian_t)$
Australia	=	$\frac{1}{n}(australia_t + australian_t)$
South Africa	=	$\frac{1}{n}(southafrica_t + southafrican_t)$
Brazil	=	$\frac{1}{n}(brazil_t + brazilian_t)$
Russia	=	$\frac{1}{n}(russia_t + russian_t)$

degree of transition risk is not only time-varying, but also country-specific. Second, the individual words in each category in Table 1 are averaged to construct one word vector for each category. This ensures that the methodology is robust to the exact words, and the number of words, allocated to each category.[8] Finally, although the transition risk measures constructed here are motivated by the commodity currency setting, the methodology and intuition is general and potentially useful in a wider set of applications.

To construct confidence intervals for the TR_t estimates, we follow Kozlowski et al. (2019) and conduct subsampling (Politis and Romano, 1994). For the 90% confidence interval, the corpus (for any given month) is randomly partitioned into 20 subcorpura, and the word2vec algorithm is run to produce the word embedding matrix for each partition of the data. Then, the error of the projection statistic TR_t for each subsample s is

[8] Performing over 30000 random leave-one-word-out (of each category) permutations of the words listed in Table 1, and computing a transition risk measure for each unique combination of words, does not materially affect the TR_t estimates. Irrespective of country, the median correlation is never below 0.94, see Table B.1 in Appendix B.

$e^s = \sqrt{\tau_s}(TR_t^s - TR_t)$, where τ_s and TR_t^s are the number of texts and the solution to (3), respectively, in subsample s. Then, the 90% confidence interval spans the 5th and 95th percentile variances, defined by $TR_t + \frac{e^{s(19)}}{\sqrt{\tau}}$ and $TR_t - \frac{e^{s(2)}}{\sqrt{\tau}}$, where $e^{s(2)}$ and $e^{s(19)}$ denote the 2nd and 19th order statistic associated with the lower and upper bound of the confidence interval.

Figure 1 reports the country-specific transition risk measures together with the estimated uncertainty. As clearly seen in the graphs, the transition risk measures are very precisely estimated. It is also clear that there is large cross-country variation in the degree of risk across time. For Canada and Norway, for example, the degree of transition risk is generally higher in the latter half of the sample than previously, while the developments in, e.g., Australia and Malaysia are more u-shaped. However, for most countries the risk estimates peak sometime after 2013.

To further build intuition for the word embedding approach, looking at two (parts of) sentences in the news when transition risk is measured as low and high might be illustrative. The first sentence is sampled from August 2002, and reads "...any worsening of the economic climate in Norway, particularly a further deterioration in the credit cycle...". The second sentence was printed in April 2007, and reads "...Norway will be at the forefront of international climate efforts and will take a leading role in the development of a new binding climate agreement...". Clearly, both sentences are about Norway. Still, although the word *climate* is used in the first sentence, the news is obviously not related to climate change transition risk. In contrast, the usage of the word *climate* in the second sentence is closely related to transition risk. Consistent with these examples, the estimated transition risk index for Norway is also rather low in 2002, but high in 2007.

In studies using text as data, it is common to annotate graphs like those in Figure 1 with historical events to informally validate how plausible the estimates are from a narrative perspective. Such an approach is less suited here. The reason is that TR_t measures the association between a country and transition risk, and not how much climate risk is talked about per se. In other words, whereas events likely affect *how much* different topics are talked about in the public discourse, the events might not change *how and in which context* these topics are talked about. Still, for completeness, Figure 1 is annotated with important climate events suggesting at least some correlation between such events and high levels of transition risk.[9]

[9]Some transition risk fluctuations definitely have a more ambiguous interpretation. The large increase in risk for Russia in 2014, for example, might be due to a large increase in the association between Russia and risks due to conflict, or alternatively, concerns about future Russian gas supply to continental Europe. Only the latter interpretation has a plausible relationship with our definition of transition risk. For these reasons we also control for alternative uncertainty measures in the exchange rate models used in later sections.

(a) Australia

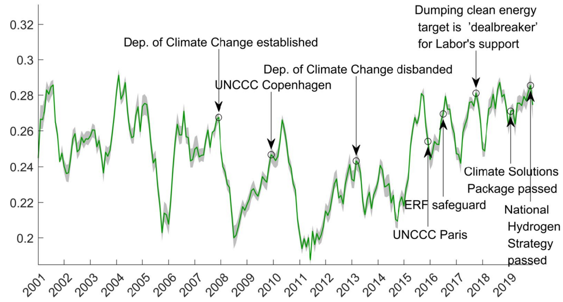

(b) Brazil

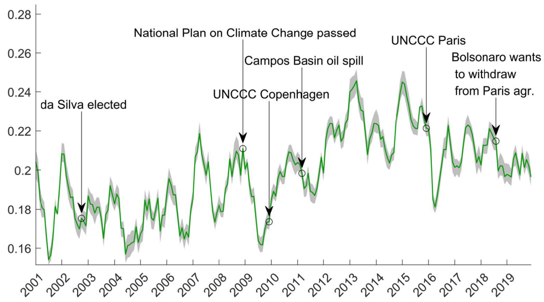

(c) Canada

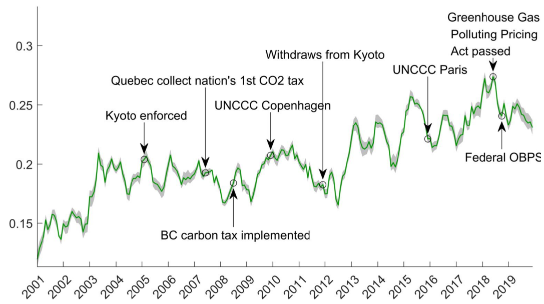

(d) Mexico

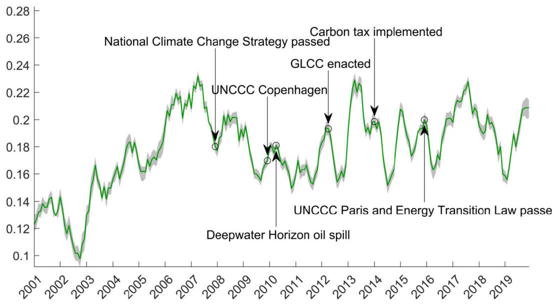

(e) Malaysia

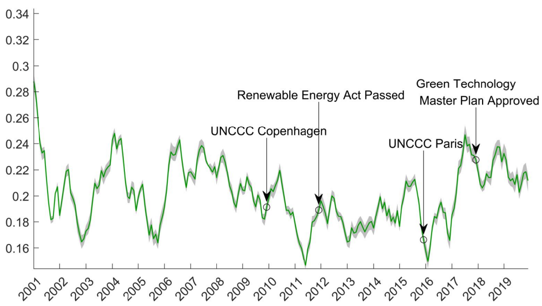

(f) Norway

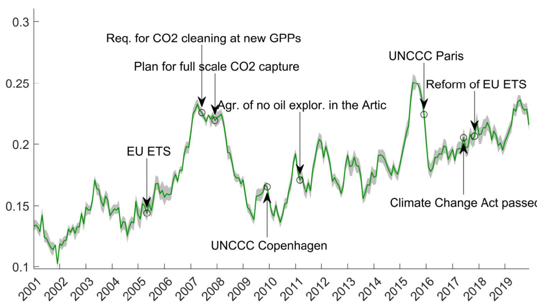

(g) Russia

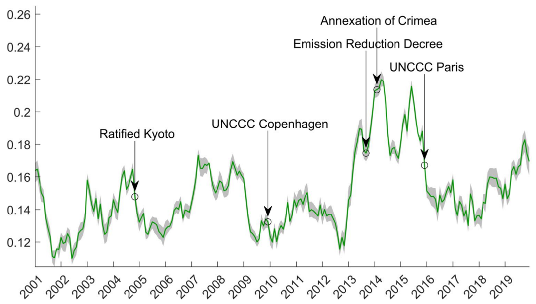

(h) South Africa

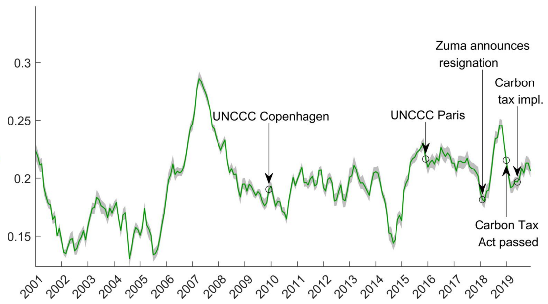

Figure 1. Climate change transition risk. The green lines show the mean estimates. The gray color shadings cover the 90% confidence intervals. The annotations report some important international and domestic political climate change events. For visual clarity, the raw TR_t series are smoothed using moving averages with a window size of seven months.

Another way to informally validate the constructed transition risk measures is to analyze how they correlate with one of the most direct and widely used measures of cli-

Table 2. Transition risk and temperature anomaly correlations. The first row reports the correlation between the raw series. The second column reports the correlation when a Hodrick–Prescott filter (Hodrick and Prescott, 1997), with a smoothing parameter set to 1600, is used to extract the low-frequency fluctuations from the series. *, **, and *** denote the 10%, 5%, and 1% significance level, respectively. Figure B.2, in Appendix B, visualizes these correlation patterns, and graphs the temperature anomaly series together with our measures of transition risk.

	Australia	*Brazil*	*Canada*	*Malaysia*	*Mexico*	*Norway*	*Russia*	*SouthAfrica*
Raw	0.04	0.17***	-0.02	0.07	-0.08	0.10	0.11*	0.08
HP-filtered	0.40***	0.52***	0.02	0.45***	-0.20***	0.63***	0.53***	0.21***

mate change, namely temperature anomalies (see, e.g., Deschenes and Greenstone (2007) and Kumar et al. (2019)). For this purpose we collect statistics from the GISS Surface Temperature Analysis (see Appendix A) and use the longitude and latitude resolution provided in that database to construct country-specific monthly time series of abnormal temperature fluctuations. Table 2 shows that the correlations are high and significant for at least six of the countries in our sample, and particularly so when looking at the low-frequency movements in the series. Still, for Canada and Mexico the correlation patterns are weak or even negative. While this lack of correlation can be explained (or excused) by a number of factors, such as media bias, measurement error, or that the two measures potentially capture two different climate risk components, a more formal way of validating our transition risk measures is to use economic theory and the predicted relationship with commodity currencies. We turn to this next.

3 Climate risk and commodity currencies

Can climate change transition risk explain commodity currency developments? We address this question using a Vector Autoregressive (VAR) modeling framework. Unlike more stringent theoretical models, as pointed out by, e.g., Rossi (2013), this approach provides a reasonable good fit to historical data, and takes into account the dynamic interaction between a set of potentially endogenous variables.

The VAR model can be written as:

$$\boldsymbol{y}_{c,t} = \boldsymbol{A}_{c,1}\boldsymbol{y}_{c,t-1} + \ldots + \boldsymbol{A}_{c,p}\boldsymbol{y}_{c,t-p} + \boldsymbol{D}_c\boldsymbol{x}_t + \boldsymbol{e}_{c,t} \quad \boldsymbol{e}_{c,t} \sim i.i.d.N(0, \boldsymbol{\Sigma}_c) \tag{4}$$

where c and t denotes the country and time indexes, p is the number of lags, and $\boldsymbol{D}_c$, $\boldsymbol{A}_{c,1}, \ldots, \boldsymbol{A}_{c,p}$, and $\boldsymbol{\Sigma}_c$ are matrices of suitable dimensions containing the model's unknown parameters. $\boldsymbol{y}_{c,t}$ is a vector containing endogenous variables and $\boldsymbol{x}_t$ is a vector of exogenous variables (including a constant).

For commodity-exporting economies, and for data sampled at monthly frequency, commonly used explanatory variables include a commodity price index to capture exogenous terms-of-trade shocks (Chen and Rogoff, 2003), short-run interest rate differentials to capture deviations from uncovered interest rate parity, and business cycle indicators to capture growth prospects (Amano and van Norden (1995), Akram (2004), Bodart et al. (2012), Ferraro et al. (2015), Zhang et al. (2016), Kohlscheen et al. (2017)). Newer studies also often include some measures of uncertainty to capture "flight-to-quality" effects in times of trouble, such as financial crises, wars, and conflict (Forbes and Warnock (2012), Rey (2015), Goldberg and Krogstrup (2018), Caldara and Iacoviello (2018), Akram (2020)). Accordingly, in our benchmark specification we let $\boldsymbol{y}_{c,t} = \begin{bmatrix} TR_{c,t} & BC_{c,t} & r_{c,t} & ComX_{c,t} & REER_{c,t} \end{bmatrix}'$, where $TR_{c,t}$ is transition risk, $BC_{c,t}$ is a business cycle index, $r_{c,t}$ is the short-run real interest rate differential, $ComX_{c,t}$ is the real commodity price index, and $REER_{c,t}$ is the real exchange rate. Finally, $\boldsymbol{x}_t = \begin{bmatrix} 1 & UNC_t & GPR_t \end{bmatrix}'$, where UNC_t is a measure of financial uncertainty and GPR_t is a measure of geopolitical risk.

The real effective exchange rates are obtained from BIS, while the commodity price indexes are obtained from Gruss and Kebhaj (2019).[10] We construct the real interest rate differentials using trade weights, and the business cycle indicators are obtained from OECD's panel of leading indicators and their business tendency survey. For countries where this variable is not available, we use the year-on-year growth in industrial production. The benchmark uncertainty measure is the VIX_t, while the geopolitical risk variable is obtained from Caldara and Iacoviello (2018). In the interest of conserving space, a more detailed description of the traditional economic variables is relegated to Appendix A.

While the model structure in (4) allows for a rich description of the dynamic relationship between the variables, our focus is on how unexpected transition risk innovations affect the real exchange rate. For this purpose we identify exogenous innovations, $\boldsymbol{\varepsilon}_{c,t}$ through the relationship $\boldsymbol{\varepsilon}_{c,t} = \boldsymbol{P}_c \boldsymbol{e}_{c,t}$ where $\boldsymbol{P}_c$ is a lower triangular matrix derived from $\boldsymbol{P}_c \boldsymbol{P}_c' = \boldsymbol{\Sigma}_c$. We do not take a strong stand on whether transition risk is contemporaneously unaffected by shocks to the other variables in the system, and therefore identify transition risk innovations by ordering climate risk either first or last in the system. These two alternative identification assumptions accommodate a view where transition risk is treated either as contemporaneously exogenous to the remaining variables in the system or as completely endogenous. As we document below, however, our qualitative results are

[10] $ComX_t$ takes into account the basket of commodities produced by country c, and is constructed using time-varying net-export shares. As discussed in Gruss and Kebhaj (2019), different findings across studies regarding the relationship between commodity prices and currencies might simply reflect differences in how the commodity price indexes are defined. Our main results regarding transition risk and exchange rates presented in later sections are robust to using the alternative commodity price indexes derived by Gruss and Kebhaj (2019).

not affected by the particular ordering, suggesting that transition risk is fairly exogenous to the other economic indicators in the very short-run.

To allow for a reasonable degree of persistence, we set $p = 12$, standardize all data prior to estimation and use data covering 2002M1 to 2019M12. This ensures that the same amount of data is available for all the countries, and it is a period in which many of the countries in the sample either directly or indirectly have an inflation-targeting monetary policy regime.

3.1 Pooled and partially pooled estimates

We begin by considering two panel VAR versions of (4), pooling information from the different units to leverage the cross-sectional information in the data. In the first specification we assume full homogeneity across units, implying that parameters are identical across countries. In the second specification we relax the homogeneity assumption and allow for random effects and cross-sectional heterogeneity by adopting a hierarchical prior approach developed by Jarociński (2010). To favor a parsimonious model structure, parameter estimates are obtained for both specifications by sampling from the posterior distribution using a Minnesota type prior variance-covariance matrix (Litterman, 1986). Since both of these specifications are fairly standard in the literature, we relegate a more detailed description of the models to Appendix C.[11]

Figure 2 summarizes our main results and reports the response of the REER to a one standard deviation climate change transition risk innovation. Figure 2a shows the pooled responses, while Figures 2b-2i show the results when allowing for cross-sectional heterogeneity. Two main conclusions stand out.

First, in line with earlier theoretical arguments, an exogenous transition risk innovation leads to a persistent and significant depreciation of the real exchange rate. This holds both for the pooled and random effect specifications. The sizes of the responses are also economically significant. For the pooled estimates, for example, a one standard deviation innovation in transition risk leads to a roughly eight percent depreciation of the REER at the one year horizon.

Second, treating transition risk as either completely contemporaneously exogenous or endogenous to the other variables in the VAR system does not matter qualitatively

[11] A battery of tests give inconsistent results across countries, regarding both the existence of variable unit roots and the degree of cointegration. In an earlier working paper version of this paper we document that all our main conclusions apply when we instead estimate the long-run relationship between real exchange rates and transition risk using a single equation framework and the Dynamic Ordinary Least Squares (DOLS) estimator (Stock and Watson, 1993) or Autoregressive Distributed Lag (ARDL) models (Pesaran and Shin, 1998). See also Section 4, where we estimate (4) individually for each country using standard OLS.

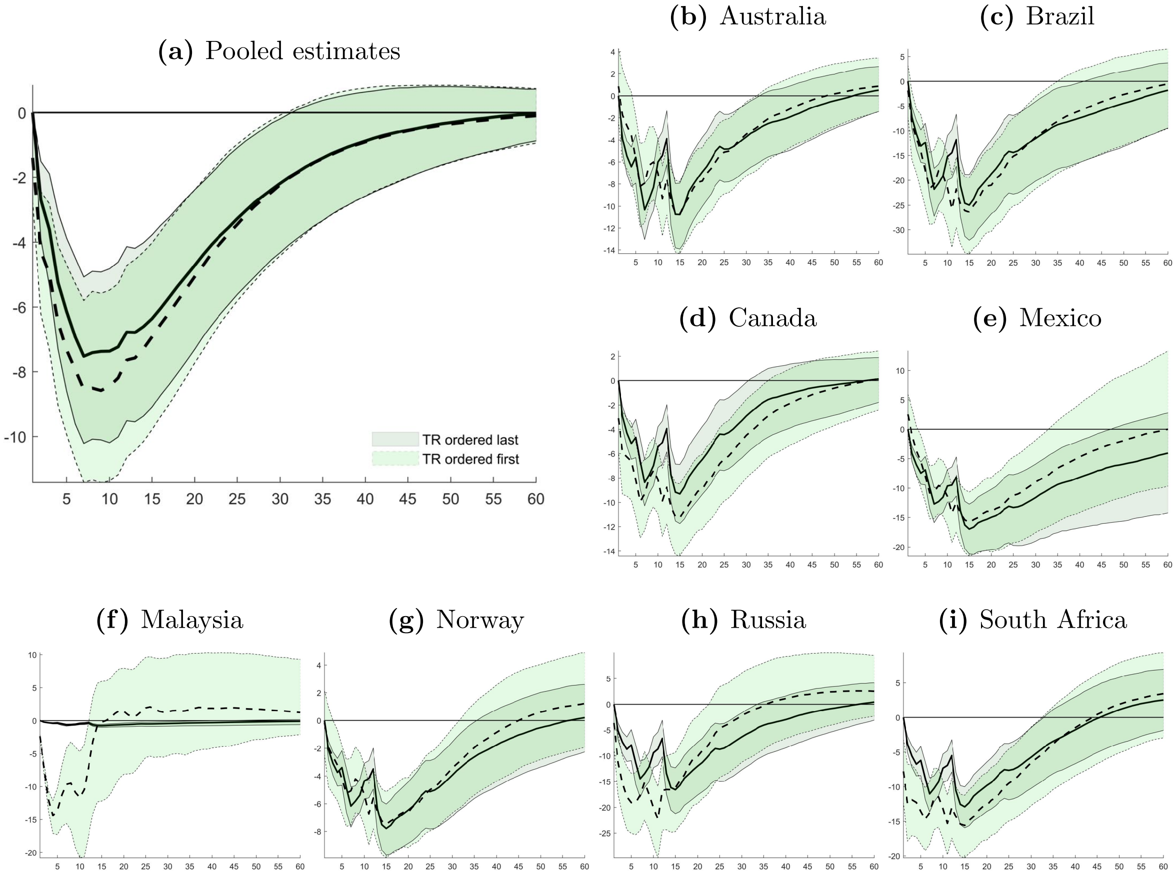

Figure 2. Pooled and partially pooled panel VAR results. Each graph reports the REER response following a one standard deviation exogenous innovation to the transition risk variable. The innovations are computed from two different recursive orderings, where the transition risk variable is ordered either first (dotted black) or last (solid black) in the system. The color shaded areas are 68% probability bands. All data is standardized prior to estimation. The pooled and partially pooled estimates are re-scaled using the average and country-specific standard deviation of the real exchange rates, respectively, and reflect percentage change.

for these conclusions. The REER responses are very similar irrespective of whether the transition risk variables are ordered first or last in the system. The exception to this general finding is Malaysia, where the random effects specification suggests a much stronger short-run depreciation when the transition risk variable is treated as contemporaneously exogenous to the other variables in the system.

For completeness, the impulse responses associated with the transition risk indexes themselves are reported in Figure B.3 in Appendix B. In short, they indicate fairly transitory response paths. Although we rightfully refrain from making strong structural claims, we also note that unexpected innovations to the other variables in the system give REER response paths reasonably in line with conventional economic theory, and that these variables respond as expected to transition risk innovations (Figures B.4 and B.5 in Appendix B): The business cycle indicators fall significantly together with negative interest rate dif-

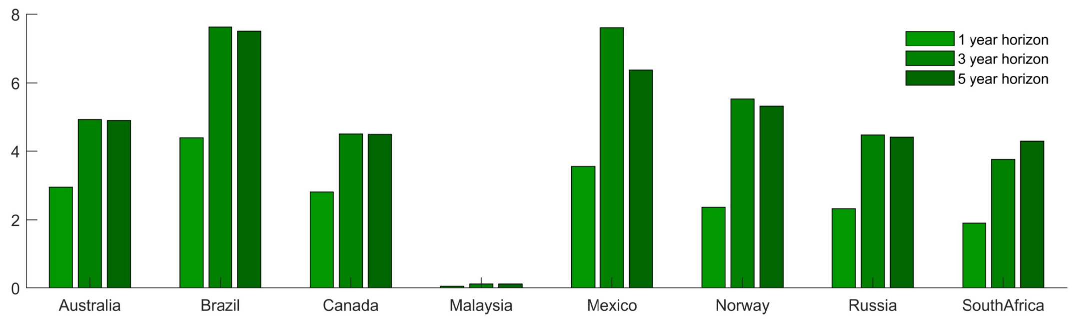

Figure 3. Partially pooled panel VAR results. REER and the variance explained by transition risk innovation. Each bar report the median estimate for a given horizon and country. Estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system.

ferentials, while the commodity price responses are insignificant. The latter result is likely because these countries are price takers, but the result can also potentially be affected by global commodity market dynamics. We discuss this further in Section 3.3. Still, it is unlikely that the proposed transition risk measures simply capture changing global demand or reflect general economic policy uncertainty. Augmenting the models with a global activity indicator (Baumeister and Hamilton, 2019), or news-based and country-specific economic policy uncertainty indexes (Baker et al., 2016), do not affect how transition risk affects real exchange rates (Figures B.6a and B.6b in Appendix B.).

3.2 Variance and historical shock decompositions

Figure 3 reports how much of the REER variance that can be explained by transition risk when this variable is ordered last in the VAR system. As such, these are conservative estimates. At the one year horizon, transition risk explains roughly 2.5% of the variation in the REERs. When the response horizon increases to three or five years, this number varies between 4% and 8% depending on which country one chooses to focus on. Again, the exception to this result is Malaysia, where transition risk does not seem to matter much. Over all, these numbers are small, but not negligible. For comparison, a large literature examining the effects of unexpected monetary policy innovations do not typically attribute more than 10% of the long-run REER fluctuations to such shocks (see, e.g., Kim et al. (2017) for a relatively recent example).

In line with the earlier finding about the insignificance of variable ordering, the transition risk indexes are also largely exogenous to the other variables in the system. As seen in Figure B.7, in Appendix B, most of variation in these indexes are explained by their own innovations.

To probe deeper into the timing of when transition risk historically has put upward or downward pressure on typical commodity currencies, Figure 4 reports the actual REER for each country (solid black line), the counterfactual REER without transition risk inno-

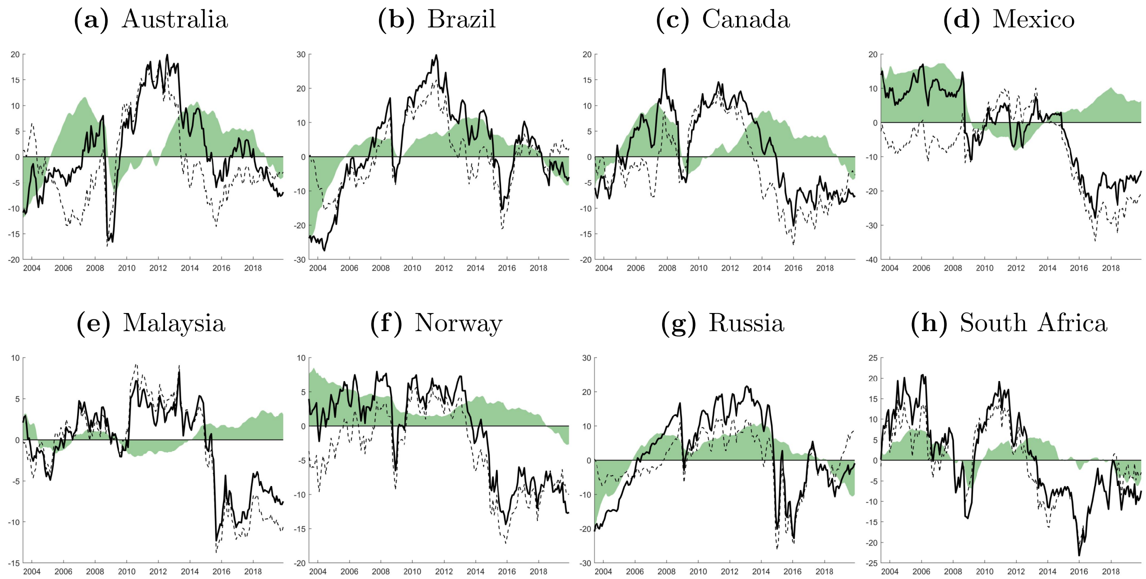

Figure 4. Partially pooled panel VAR results. REER and historical shock decompositions. Each graph reports the actual REER (solid black), the counterfactual REER without transition risk innovations (broken black line), and the difference between these two lines (green area). Estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system.

vations (broken black line), and the difference between these two lines (green area), i.e., the historical shock contribution from transition risk.

Given recent media coverage and REER developments, our prior view would have been consistent with negative transition risk contributions towards the latter part of the sample. We also observe this for six of the eight countries after roughly 2018. It is perhaps more surprising how the model interprets the mid 2000s for Australia, Brazil, Canada, and Russia, and the period 2012-2016 for most of the countries. For the former period and group of countries, transition risk was much higher than expected (within the model), and thus put strong negative pressure on the REERs. In this sense, events related to the introduction of the EU ETS and Kyoto enforcement are correlated with changes in how transition risk is written about in the press. In contrast, for the period 2012-2016 transition risk is generally interpreted as being lower than expected, which suggests that events associated with, e.g., the Paris Agreement, actually did not lead to unpredictable short-term increases in transition risk.

3.3 Corroborative results

The theoretical mapping between traditional "Dutch disease economics" and transition risk gives rise to at least two additional testable hypotheses. First, since transition risk accommodates the future risk of unfavorable shifts in the production function of the commodity-producing sector, commodity supply should on average fall in response to positive transition risk innovations. Second, because natural resource income is an important part of aggregate income creation in major commodity exporters, the mechanisms that give rise to a persistent exchange rate depreciation might also affect forward-looking asset markets at the national level.[12]

To address these hypotheses we include either a country's commodity production or stock market index in the VAR and analyze how these variables respond to transition risk. Country-specific monthly data on coal and gas production is missing for most of the countries in our sample. We therefore restrict the analysis to seasonally adjusted oil production, and leave Australia and South Africa out of the analysis because they produce only small amounts of oil (see Figure B.1, in Appendix B). Furthermore, since the oil production series show very different trends across countries, and partly also the stock market indexes, we include a linear trend as an additional exogenous variable in these specifications. As before, the transition risk indexes are ordered last in the VAR systems.

The results are presented in Figure 5. In terms of the stock market responses, both the pooled and partially pooled estimates suggest a significant negative response in the short run. However, whereas the pooled estimates indicate that the effect eventually dies out, the partially pooled estimates are more persistent, especially in countries such as Australia, Brazil, and Canada.

The results regarding oil production are more model-dependent. For the pooled estimates, the oil production response is mostly negative, but highly uncertain. Given the differences in underlying oil production trends, this result is perhaps as expected when (unrealistically) pooling information across units. In contrast, for the partially pooled estimates we observe a significant reduction in oil production. For most countries this reduction is significant at the medium response horizons, but also rather persistent in countries such as Canada and Mexico. Interestingly, there are also signs that commodity production increases in the very short run, in line with the "green-paradox" originally coined by Sinn et al. (2008).[13]

[12] Even within major commodity exporters, some sectors might benefit at the expense of others when faced with transition risk. Indeed, the theoretical mechanism we build on predicts changes in sectoral capital allocations. In the case of Norway, we have explored this further, finding that changes in transition risk have an increasingly negative correlation with companies within the *Energy* portfolio on the Oslo Stock Exchange, whereas other sectors, such as *Telecom.*, experience an increasingly positive correlation. These additional results can be obtained on request.

[13] Since our proposed risk measures also accommodate failed discoveries of new resources, including the dynamic response of remaining commodity reserves is relevant. Such statistics, however, are only available at a yearly frequency. Still, yearly correlations do not indicate any consistent pattern between the two variables, ruling out transition risk as simply a proxy for changes in remaining reserves (Table B.2 in Appendix B).

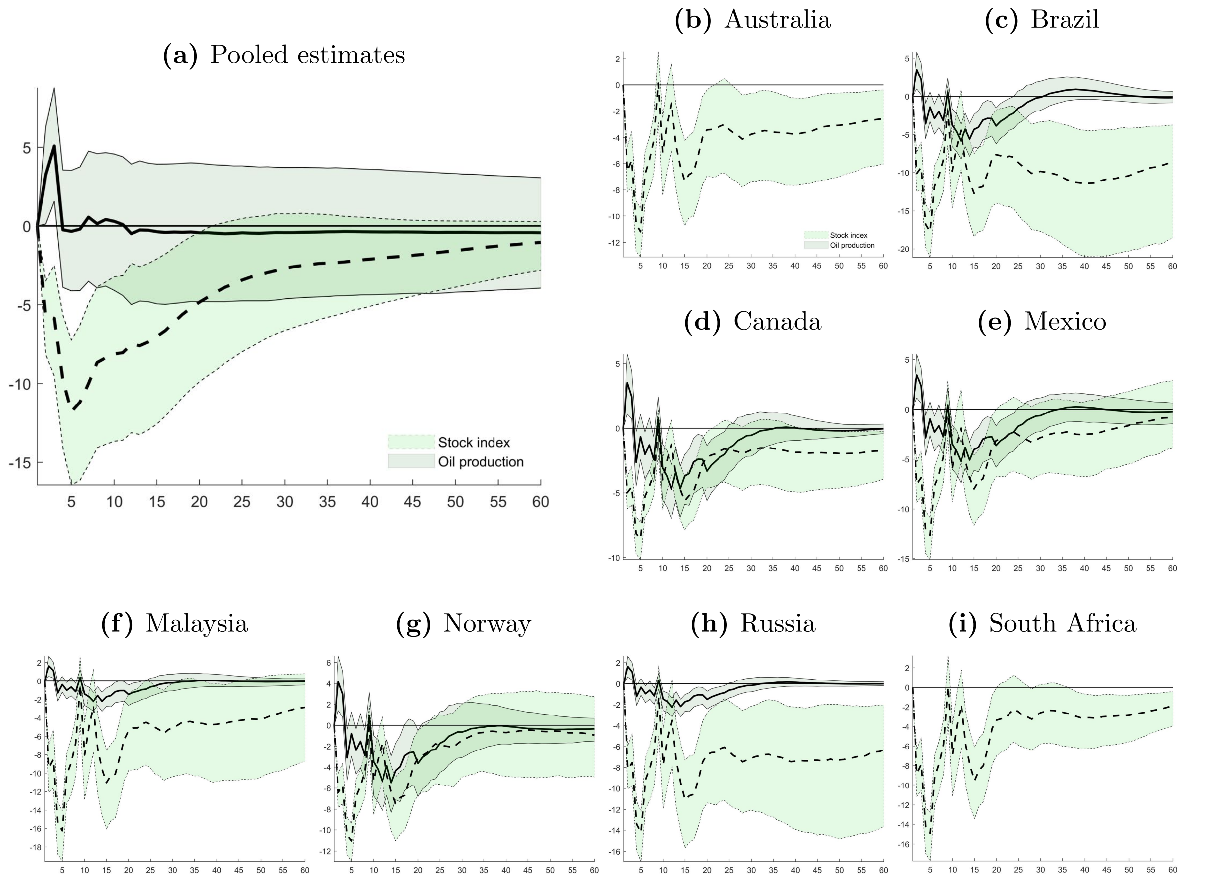

Figure 5. Pooled and partially pooled panel VAR results. Each graph reports the oil production and stock market responses following a one standard deviation exogenous innovation to transition risk. Estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system. The color shaded areas are 68% probability bands. All data is standardized prior to estimation. The pooled and partially pooled estimates are re-scaled using the average and country-specific standard deviation of either the production or stock market indexes, respectively, and reflect percentage change.

We interpret these results as largely consistent with our underlying theoretical motivation. The results for the stock market speak to a large literature in finance investigating the implications for firm value of increased climate risk, and in particular studies taking a "stranded assets" perspective (see, e.g., Ramelli et al. (2018), Atanasova and Schwartz (2019), van der Ploeg and Rezai (2020), Sen and von Schickfus (2020)). Empirical estimates of oil supply responses following climate risk innovations are more scarce. One exception is Barnett (2019), who constructs an (somewhat U.S.-centered) event-based climate policy index and finds that global oil supply increases in response to an increased likelihood of significant climate policies being introduced. In contrast, we focus on country-specific responses, and important oil suppliers such as Saudi Arabia and the United Arab Emirates are not part of our analysis because they do not have floating exchange rates. We leave it for future research to examine how transition risk affects global oil market dynamics.

4 Unit effects and falsification experiments

Below we fully relax the panel assumptions used in the previous sections, and estimate individual VAR models for each of the eight commodity countries analyzed in the previous sections. In addition, similar VAR models are estimated for all the other countries having floating exchange rates in the BIS real effective exchange rate database.[14] This allows us to not only analyze the sensitivity of the pooled (bayesian) estimates reported earlier, but also perform two types of "falsification" experiments.

First, although transition risk might be a risk all countries are exposed to, our theoretical motivation predicts that this type of climate risk should be particularly relevant for fossil fuel exporting countries. Thus, when analyzing a large number of countries we should expect to see a significant negative correlation between a country's commodity export dependency and the real exchange rate response following innovations in transition risk.

Second, because climate risk is not directly observed, the literature has used different approaches to approximate it. As a result, existing measures vary in the degree to which they capture *physical*, *liability*, or *transition* risk associated with climate change (Carney, 2015). In fact, it can be argued that the economic and financial literature have focused foremost on the former risk component, and that our contribution in terms of measurement is related to the latter component. To assess to what extent this innovation matters, we use three alternative existing proxies for climate risk, and compare the results to those obtained when using our proposed measure.

The scatter plots in Figure 6 largely confirm the first hypothesis. The y-axis reports the REER responses on either the one- or five-year horizon following a transition risk innovation, while the x-axis reports the net fossil fuel commodity export share relative to overall GDP. As clearly seen in the figure, there is a significant and negative relationship between these two variables. The box plots to the right in the figure further confirms this impression. The real exchange rate responses for commodity currency countries are on average more negative than in non-commodity currency countries. Indeed, for non-commodity currencies the responses are not significantly different from zero on average.

Furthermore, the one-year horizon REER responses for the commodity currencies are qualitatively in line with the pooled and partially pooled Panel VAR results reported earlier. The exception to this is Australia, which has the opposite sign relative to the

[14] Each VAR includes the same endogenous variables as in earlier sections. Because of the reduced degrees of freedom when estimating individual models compared to the panels, the lag length is decreased from 12 to 6, and parameter estimates are obtained by maximum likelihood estimation to relax the computational burden. Transition risk measures for all of the non-commodity currencies are estimated as described in Section 2.2.

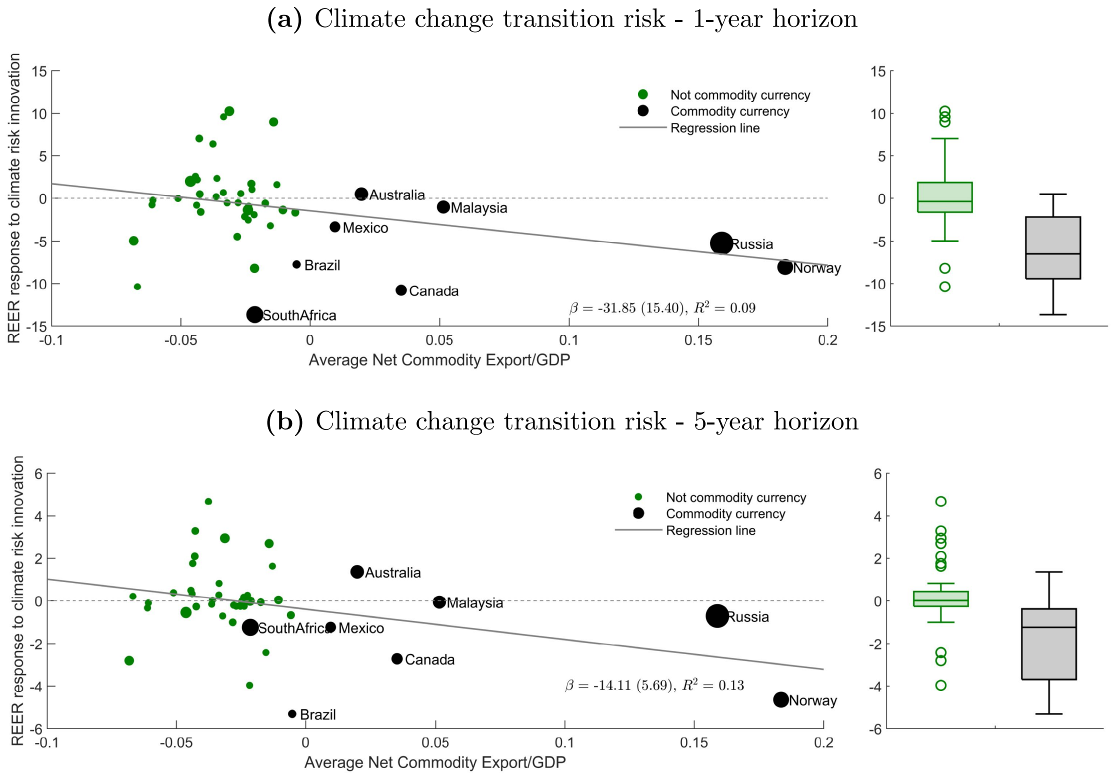

Figure 6. REER responses and commodity export shares. Each graph reports a country's REER response following a one standard deviation exogenous innovation to the transition risk variable (y-axis: in percentage change) together with net commodity exports relative to GDP (x-axis). The REER response estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system. The net commodity export relative to GDP statistic reflects the average across the period 2002-2019. Observations for (fossil fuel) commodity and non-commodity currencies are colored black and green, respectively. The size of the scatters reflects the country's CO2 emissions relative to GDP. The box plot to the right in each graph reports the median, interquartile range (IQR), and outliers ($1.5 \times IQR$ as circles) .

results reported in Figure 2. In contrast, at the five-year horizon the earlier Panel VAR point estimates tended to be positive, while the individual VAR results reported in Figure 6 suggest a more persistent depreciation for at least half of the commodity currencies.

Figure 7 summarizes the results from the second experiment. Here all the individual VAR models are re-estimated using either the recent news-based climate risk measures suggested by Engle et al. (2020) and Gavriilidis (2021), or the temperature anomalies described in Section 2.2, as alternative risk proxies.

The measure suggested by Engle et al. (2020) was developed for hedging (overall) climate risk in the asset market, but builds on a type of motivation similar to ours, where the news media implicitly operate as information intermediaries between agents and the state of the world. However, their index does not directly distinguish between the three different types of climate change risks, and essentially measures *how much* climate change is focused upon in the news using an inverse document frequency count-based approach.

In contrast, our word embedding approach measures in *what context* it is focused upon, and aims to capture country-specific transition risk. Figure 7a shows that this matters for describing the relationship between climate risk and commodity currencies. When using the Engle et al. (2020) measure, one observes that the estimated REER responses are theory-inconsistent regarding their sign, and that there is no significant relationship between a country's commodity export dependency and the real exchange rate response following innovations in climate risk.

As an alternative news-based index, Gavriilidis (2021) develops a climate policy uncertainty index (CPU) building on the method proposed in Baker et al. (2016) for measuring economic policy uncertainty. In particular, he searches for articles in eight leading US newspapers containing terms related to uncertainty, climate risk and regulation, and then scales the number of relevant articles per month with the total number of articles during the same month. By focusing on uncertainty and regulation, the intuition for this measure is related to ours, and using the CPU produces a negative correlation between a country's net commodity export dependency and the real exchange rate responses (Figure 7b). Still, this relationship is far from as strong as that produced by our proposed transition risk measures.

Figure 7c reports estimates from (4) when our measure of climate risk is replaced by the temperature anomaly statistics. As above, there is an insignificant correlation between a country's net commodity export dependency and the real exchange rate responses following innovations in temperature anomalies. And, there are very few signs that commodity currencies as a group have a lower REER response than non-commodity currencies.

One potential reason for these conflicting results might be that the other series, and perhaps temperature anomalies in particular, approximate physical climate risk rather than transition risk. Another reason might have to do with information diffusion and the role of media as information intermediaries. Most people follow the news, but only very few follow temperature anomaly statistics closely. A third reason might be that both the Engle et al. (2020) and Gavriilidis (2021) indexes are global measures (or somewhat U.S.-centered), while our measures are country-specific.

Another proxy for climate risk sometimes used in the literature is so-called Climate Change Performance Indexes (CCPI). A well-known set of measures in this respect are produced by the non-governmental organization *Germanwatch* since 2005. Their CCPIs track countries' efforts to combat climate change, and evaluates and compares their climate protection performance based on indicators covering categories such as GHG emissions, renewable energy, energy use, and climate policy. Unfortunately, these indexes are available only at a yearly frequency and are thus not appropriate for our VAR analysis. In

(a) Engle et. al. : Climate risk - 5-year horizon

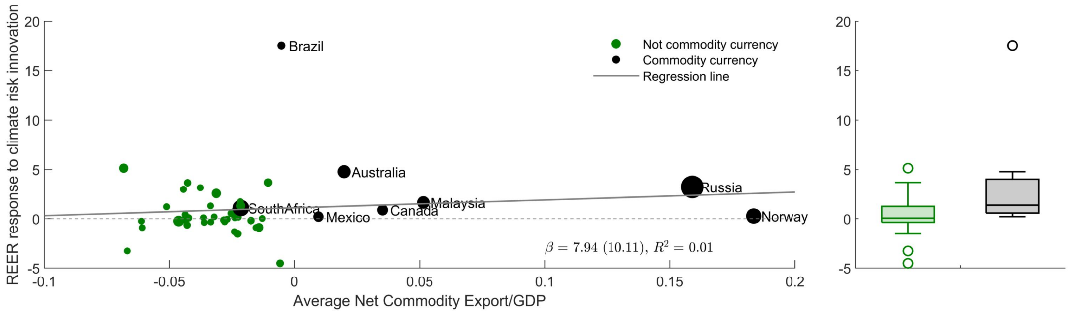

(b) Gavriilidis: Climate policy uncertainty - 5-year horizon

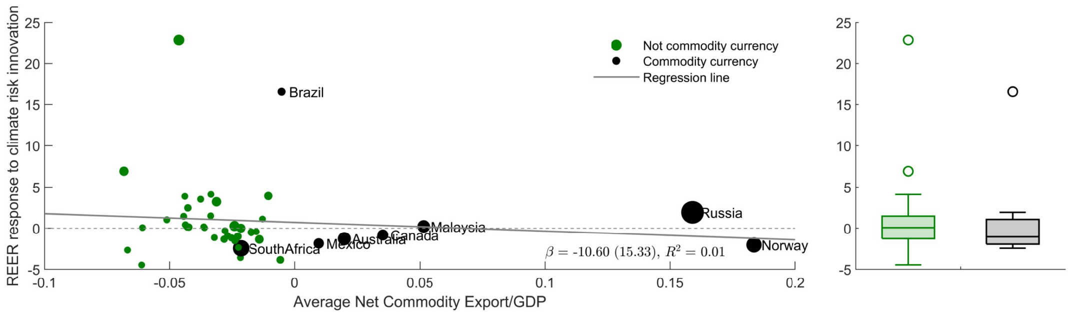

(c) GISS: Temperature anomalies - 5-year horizon

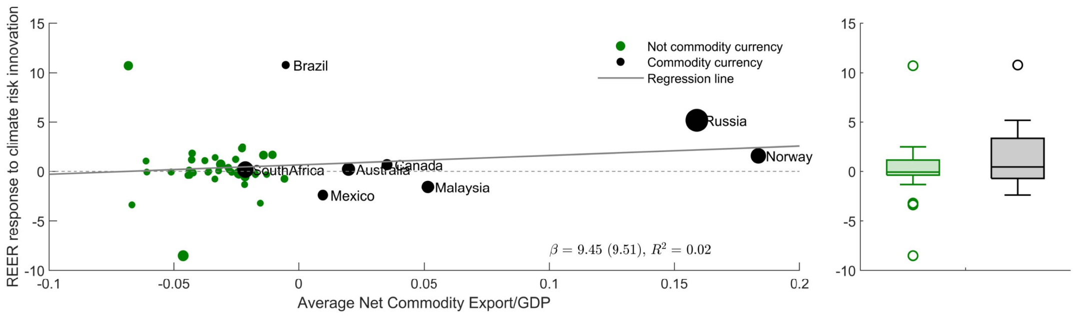

Figure 7. REER responses for alternative climate risk proxies and commodity export shares. Each graph reports a country's REER response following a one standard deviation exogenous innovation to either the Engle et al. (2020) climate risk index, the climate policy uncertainty index developed by Gavriilidis (2021), or temperature anomalies (y-axis: in percentage change) together with net commodity exports relative to GDP (x-axis). The REER response estimates are obtained assuming a recursive ordering with the climate risk variable ordered last in the VAR system. Results for the 1-year horizon are reported in Figure B.8 in Appendix B. See Figure 6 for additional details.

terms of correlations, however, Table B.3 in Appendix B shows that there is little evidence suggesting a significant yearly relationship between our proposed transition risk measures and the CCPIs. If anything, transition risk leads the CCPIs.

5 Conclusion

Economic theory on changes in natural resource income predicts an inverse relationship between real exchange rate developments and increases in climate change transition risk. In this article we propose a novel measure of such risk, constructed using media coverage and word embedding models to relate them to concerns about adjustments towards a lower-carbon economy voiced in the public discourse, and analyze how it affects the real exchange rates of eight major fossil fuel producers.

In line with theory we document that when transition risk increases, commodity currencies experience a persistent depreciation. According to our estimates, between 4% and 8% of the medium- to long-run fluctuations in the real exchange rate can be explained by unexpected transition risk innovations. Furthermore, when analyzing a rich set of countries, including both commodity currencies and non-commodity currencies, we find a significant negative correlation between a country's commodity export dependency and the real exchange rate response following innovations in transition risk. Finally, none of these findings apply when we use existing, and commonly used, climate risk proxies.

At a general level, our analysis contributes to the broader climate literature by proposing a methodology for measuring climate change transition risk. The vast majority of literature on the topic has concentrated mostly on the economic and financial effects arising from climate- and weather-related events, i.e., physical climate change risk. In contrast, our results show that decomposing climate risk into different components is relevant empirically, and that studies not distinguishing between different climate risk components might misinterpret the economic consequences of climate change.

More specifically, our study speaks to a large literature on changes in natural resource income and the pricing implications of climate risk. In terms of the former, corroborative results give further support for the "Dutch disease" mechanism, as we document that both commodity production and aggregate stock market valuations fall in response to transition risk innovations. In terms of the latter, most of this literature has been concerned with pricing of firms and firm value, while we document how transition risk also affects valuations at a national level.

References

Akram, Q. F. (2004). Oil prices and exchange rates: Norwegian evidence. *The Econometrics Journal 7*(2), 476–504.

Akram, Q. F. (2020). Oil price drivers, geopolitical uncertainty and oil exporters' currencies. *Energy Economics 89*, 104801.

Amano, R. A. and S. van Norden (1995). Terms of trade and real exchange rates: the Canadian evidence. *Journal of International Money and Finance 14*(1), 83 – 104.

Andersson, M., P. Bolton, and F. Samama (2016). Hedging climate risk. *Financial Analysts Journal 72*(3), 13–32.

Arezki, R., V. A. Ramey, and L. Sheng (2017). News Shocks in Open Economies: Evidence from Giant Oil Discoveries. *The Quarterly Journal of Economics 132*(1), 103–155.

Arrhenius, S. (1896). On the influence of carbonic acid in the air upon the temperature of the ground. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 41*(251), 237–276.

Atanasova, C. and E. S. Schwartz (2019). Stranded fossil fuel reserves and firm value. Technical report, National Bureau of Economic Research.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Barnett, M. D. (2019). *A run on oil: Climate policy, stranded assets, and asset prices.* Ph. D. thesis, The University of Chicago.

Batten, S., R. Sowerbutts, and M. Tanaka (2016, May). Let's talk about the weather: the impact of climate change on central banks. Bank of England working papers 603, Bank of England.

Baumeister, C. and J. D. Hamilton (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review 109*(5), 1873–1910.

Bjørnland, H. C. and L. A. Thorsrud (2016). Boom or gloom? Examining the Dutch disease in two-speed economies. *The Economic Journal 126*(598), 2219–2256.

Bjørnland, H. C., L. A. Thorsrud, and R. Torvik (2019). Dutch disease dynamics reconsidered. *European Economic Review 119*, 411 – 433.

Bodart, V., B. Candelon, and J.-F. Carpantier (2012). Real exchanges rates in commodity producing countries: A reappraisal. *Journal of International Money and Finance 31*(6), 1482 – 1502.

Bolton, P. and M. Kacperczyk (2020, April). Do investors care about carbon risk? Working Paper 26968, National Bureau of Economic Research.

Boykoff, M. T. and J. M. Boykoff (2007). Climate change and journalistic norms: A case-study of us mass-media coverage. *Geoforum 38*(6), 1190–1204. Theme Issue: Geographies of Generosity.

British Petroleum Company (2020). BP statistical review of world energy. Technical report, British Petroleum Co.

Bruno, M. and J. Sachs (1982). Energy and resource allocation: a dynamic model of the "Dutch Disease". *The Review of Economic Studies 49*(5), 845–859.

Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature 527*, 235–239.

Caldara, D. and M. Iacoviello (2018). Measuring geopolitical risk. *FRB International Finance Discussion Paper* (1222).

Campiglio, E. and R. v. der Ploeg (2021, June). Macro-financial transition risks in the fight against global warming. Working paper, SSRN.

Canova, F. and M. Ciccarelli (2013). Panel vector autoregressive models: A survey. In *VAR Models in Macroeconomics–New Developments and Applications: Essays in Honor of Christopher A. Sims*. Emerald Group Publishing Limited.

Carattini, S., G. Heutel, and G. Melkadze (2021, March). Climate policy, financial frictions, and transition risk. Working Paper 28525, National Bureau of Economic Research.

Carney, M. (2015). Breaking the tragedy of the horizon – climate change and financial stability. *Speech given at Lloyd's of London*.

Cha, M., S. Lee, E. N. Mark, J. Nauerz, J. Rawls, and Z. Wei (2020). Climate Change and Real Exchange Rates. Mimeo, University of Notre Dame.

Charnavoki, V. and J. Dolado (2014). The Effects of Global Shocks on Small Commodity-Exporting Economies: Lessons from Canada. *American Economic Journal: Macroeconomics 6*(2), 207–237.

Chen, Y. and K. Rogoff (2003). Commodity currencies. *Journal of International Economics 60*, 133–160.

Choi, D., Z. Gao, and W. Jiang (2020, 02). Attention to Global Warming. *The Review of Financial Studies 33*(3), 1112–1145.

Corden, W. M. (1984). Booming sector and Dutch disease economics: Survey and consolidation. *Oxford Economic Papers 36*(3), 359–380.

Daniel, K. D., R. B. Litterman, and G. Wagner (2019). Declining Co2 price paths. *Proceedings of the National Academy of Sciences 116*(42), 20886–20891.

Dell, M., B. F. Jones, and B. A. Olken (2014). What do we learn from the weather? the new climate-economy literature. *Journal of Economic Literature 52*(3), 740–98.

Deschenes, O. and M. Greenstone (2007, March). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review 97*(1), 354–385.

Eastwood, R. and A. J. Venables (1982). The macroeconomic implications of a resource discovery in an open economy. *The Economic Journal 92*(366), 285–299.

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel (2020). Hedging Climate Change News. *The Review of Financial Studies 33*(3), 1184–1216.

Ferraro, D., K. Rogoff, and B. Rossi (2015). Can oil prices forecast exchange rates? An empirical analysis of the relationship between commodity prices and exchange rates. *Journal of International Money and Finance 54*, 116 – 141.

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

Freeman, M. C., G. Wagner, and R. J. Zeckhauser (2015). Climate sensitivity uncertainty: when is good news bad? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 373*(2055), 20150092.

Gavriilidis, K. (2021). Measuring climate policy uncertainty. *Available at SSRN 3847388*.

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

Goldberg, L. S. and S. Krogstrup (2018). International capital flow pressures. Technical report, National Bureau of Economic Research.

Gruss, B. and S. Kebhaj (2019). Commodity terms of trade. Technical report, USA.

Hale, B. (2010). Using newspaper coverage analysis to evaluate public perception of management in river-floodplain systems. *Environmental Management 45*(5), 1155–1163.

Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the fomc: A computational linguistics approach. *The Quarterly Journal of Economics 133*(2), 801–870.

Hirshleifer, D. (2020). Presidential address: Social transmission bias in economics and finance. *The Journal of Finance 75*(4), 1779–1831.

Hodrick, R. J. and E. C. Prescott (1997). Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking 29*(1), 1–16.

Hong, H., F. W. Li, and J. Xu (2019). Climate risks and market efficiency. *Journal of econometrics 208*(1), 265–281.

Hsu, P.-H., K. Li, and C.-Y. Tsou (2020). The pollution premium. *Available at SSRN 3578215*.

In, S. Y., K. Y. Park, and A. Monk (2017). Is 'Being Green' Rewarded in the Market? An Empirical Investigation of Decarbonization Risk and Stock Returns. *International Association for Energy Economics (Singapore Issue) 46*, 48.

Jarociński, M. (2010). Responses to monetary policy shocks in the east and the west of europe: a comparison. *Journal of Applied Econometrics 25*(5), 833–868.

Kadiyala, K. R. and S. Karlsson (1997). Numerical methods for estimation and inference in bayesian var-models. *Journal of Applied Econometrics 12*(2), 99–132.

Karlsson, S. (2013). Forecasting with bayesian vector autoregression. *Handbook of economic forecasting 2*, 791–897.

Keeling, C. D. (1970). Is carbon dioxide from fossil fuel changing man's environment? *Proceedings of the American Philosophical Society 114*(1), 10–17.

Kim, S.-H., S. Moon, and C. Velasco (2017). Delayed overshooting: Is it an '80s puzzle? *Journal of Political Economy 125*(5), 1570–1598.

Kohlscheen, E., F. Avalos, and A. Schrimpf (2017). When the Walk Is Not Random: Commodity Prices and Exchange Rates. *International Journal of Central Banking 13*(2), 121–158.

Kozlowski, A. C., M. Taddy, and J. A. Evans (2019). The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review 84*(5), 905–949.

Krueger, P., Z. Sautner, and L. T. Starks (2020). The Importance of Climate Risks for Institutional Investors. *The Review of Financial Studies 33*(3), 1067–1111.

Kumar, A., W. Xin, and C. Zhang (2019). Climate sensitivity and predictable returns. *Available at SSRN 3331872*.

Larsen, V. H. and L. A. Thorsrud (2019). The value of news for economic developments. *Journal of Econometrics 210*(1), 203–218.

Larsen, V. H., L. A. Thorsrud, and J. Zhulanova (2020). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics* (Forthcoming).

Lenssen, N., G. Schmidt, J. Hansen, M. Menne, A. Persin, R. Ruedy, and D. Zyss (2019). Improvements in the GISTEMP uncertainty model. *J. Geophys. Res. Atmos. 124*(12), 6307–6326.

Litterman, R. B. (1986). Forecasting with bayesian vector autoregressions - five years of experience. *Journal of Business and Economic Statistics 4*(1), 25–38.

Mikolov, T., K. Chen, G. Corrado, and J. Dean (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119.

Nimark, K. P. and S. Pitschner (2019). News media and delegated information choice. *Journal of Economic Theory 181*, 160 – 196.

Nordhaus, W. D. (1977). Economic growth and climate: the carbon dioxide problem. *The American Economic Review 67*(1), 341–346.

P., S. (1986). Informing and educating the public about risk. *Risk Anal. 6*(4), 403–415.

Pesaran, M. H. and Y. Shin (1998). An autoregressive distributed-lag modelling approach to cointegration analysis. *Econometric Society Monographs 31*, 371–413.

Politis, D. N. and J. P. Romano (1994, 12). Large sample confidence regions based on subsamples under minimal assumptions. *Annals of Statistics 22*(4), 2031–2050.

Ramelli, S., A. F. Wagner, R. J. Zeckhauser, and A. Ziegler (2018, November). Investor rewards to climate responsibility: Evidence from the 2016 climate policy shock. Working Paper 25310, National Bureau of Economic Research.

Rey, H. (2015). Dilemma not trilemma: the global financial cycle and monetary policy independence. Technical report, National Bureau of Economic Research.

Rogoff, K. and M. Obstfeld (1996). *Foundations of International Macroeconomics*. Cambridge, Massachusetts: MIT Press.

Rossi, B. (2013). Exchange rate predictability. *Journal of Economic Literature 51*(4), 1063–1119.

Sampei, Y. and M. Aoyagi-Usui (2009). Mass-media coverage, its influence on public awareness of climate-change issues, and implications for japan's national campaign to reduce greenhouse gas emissions. *Global Environmental Change 19*(2), 203–212.

Schoenfeld, A. C., R. F. Meier, and R. J. Griffin (1979). Constructing a social problem: The press and the environment. *Social Problems 27*(1), 38–61.

Sen, S. and M.-T. von Schickfus (2020). Climate policy, stranded assets, and investors' expectations. *Journal of Environmental Economics and Management 100*, 102277.

Shiller, R. J. (2017). Narrative economics. *American Economic Review 107*(4), 967–1004.

Sinn, H.-W. et al. (2008). Das grüne paradoxon; plädoyer für eine illusionsfreie klimapolitik. *Monographs in Economics*.

Stock, J. H. and M. W. Watson (1993). A simple estimator of cointegrating vectors in higher order integrated systems. *Econometrica*, 783–820.

ter Ellen, S., V. H. Larsen, and L. A. Thorsrud (2020). Narrative monetary policy shocks and the media. *Journal of Money, Credit and Banking*. Forthcoming.

Thorsrud, L. A. (2018). Words are the new numbers: A newsy coincident index of the business cycle. *Journal of Business & Economic Statistics*, 1–17.

Torvik, R. (2001). Learning by doing and the Dutch disease. *European Economic Review 45*(2), 285–306.

van der Ploeg, F. and A. Rezai (2020). Stranded assets in the transition to a carbon-free economy. *Annual Review of Resource Economics 12*(1), 281–298.

van Wijnbergen, S. (1984). The "Dutch Disease": a disease after all? *The Economic Journal 94*(373), 41–55.

Zhang, H. J., J.-M. Dufour, and J. W. Galbraith (2016). Exchange rates and commodity prices: Measuring causality at multiple horizons. *Journal of Empirical Finance 36*, 100 – 120.

Appendices for online publication

Appendix A Data Description

Exchange rates and trade weights. The real effective exchange rate indices $REER_{c,t}$ are obtained from the Bank for International Settlements (BIS). The $REER_{c,t}$ is based on trade weights, where 40 of the most important trading partners for country c are considered. The trade weights $w_{c,i,t}$ of country c, trading partner i, and time t are also used to construct interest rate differentials. See below. The weights are available for three-year periods: 1999-2002, 2003-2005, ..., and 2014-2016. As trade weights for the period 2017-2019 were not yet available, we use the last available trade weights for this latter period.

Interest rate differentials. Due to data availability issues the short-term interest rates are obtained from different sources. The majority of interest rate differentials are computed using 3-month Treasury bill yields obtained from the Global Financial Data (GFD) database (available for 36 out of 60 countries). For the remaining countries we use 3-month interbank interest rates obtained from the GFD database (available for 7 out of 60 countries), Treasury bill yields and interbank rates with 3- month maturity from OECD's MEI database (available for 11 out of 60 countries), or short-term interest rates collected from Macrobond (4 out of 60 countries). For Argentina and Turkey, we could not obtain any representative short-term interest rates for the whole sample period. Year-on-year inflation for most countries is obtained from BIS. Inflation for Taiwan and Colombia is obtained from the GFD. The CPI for Russia is obtained from FRED. Missing values for monthly inflation of the United Arab Emirates from Jan 2001 to Dec 2008 are replaced by the annual inflation obtained from the FRED database. Real short-term interest rates $r^*_{c,t}$ for country c are created by subtracting year-on-year inflation from nominal short-term interest rates. The real short-term interest rate differential is then created by taking the difference between the real short-term interest rate and the trade-weighted real short-term interest rates of its 38 available trading partners: $r_{c,t} = r^*_{c,t} - \sum_{i=1}^{38} w_{c,i,t} * r^*_{t,i}$

Commodity price indexes. The country-specific commodity price indexes are obtained from Gruss and Kebhaj (2019). Their preferred measure is obtained by multiplying commodity-specific price indexes with the time-varying weights of each country's net export shares relative to the GDP. The alternative indexes derived in Gruss and Kebhaj (2019) use either fixed weights or time-varying weights based on each country's export (not net) shares relative to the GDP.

Business cycle indicators. Our preferred business cycle indicator is the forward-looking (amplitude-adjusted) business confidence indicators provided by OECD in their

MEI database. However, this measure is not available for all countries (available for 41 out of 60 countries). In cases where the business confidence indicator is missing, we instead use year-on-year changes in industrial production obtained from OECD or Macrobond.

Uncertainty measures. We obtained three different uncertainty measures. The volatility index for financial markets UNC_t is obtained from the Chicago Board Options Exchange, which retrieves the constant 30-day expected volatility from call and put options on the S&P500. The (global) geopolitical risk index GPR_t is obtained from Caldara and Iacoviello (2018), while the news-based country-specific economic policy uncertainty measures $EPU_{c,t}$ are obtained from Baker et al. (2016). Both GPR_t and $EPU_{c,t}$ are based on counting the occurrence of words related to geopolitical tensions or economic policy uncertainty in leading international newspapers.

Fuel net export as a share of GDP. Fuel exports and imports for each country on an annual frequency are obtained from the World Integrated Trade Solution (WITS). The term 'fuel' describes all products classified in section 27 of the HS1996 code list "Mineral fuels, oils & product of their distilliation; etc". GDP at an annual frequency is obtained from the World Bank's World Development Indicators database.

Reserves of fossil fuels. Reserves of oil and coal at an annual frequency are obtained from the BP Statistical Review of World Energy.

World industrial production index. This measure is constructed by Baumeister and Hamilton (2019) and combines industrial production of OECD countries plus the world's six largest non-OECD economies.

Oil production. Crude oil production including lease condensates is obtained at a monthly frequency from U.S. Energy Information Administration. The series are seasonally adjusted using the X12-ARIMA filter from the U.S. Census Bureau.

Stock market indices. The MSCI IMI total return indexes in local currency are sourced from *Macrobond*.

Alternative climate risk proxies. The news-based (general) climate risk measure is obtained from Engle et al. (2020), while the news-based climate policy uncertainty is obtained from Gavriilidis (2021). Both series are available at a monthly frequency. The Climate Change Performance Index (CCPI) reports for the years 2005-2019 are obtained from Germanwatch. We focused on the ranking of countries since Germanwatch's methodology to calculate the score of the CCPI changed over time. The CCPI is only available at an annual frequency.

Temperature Anomalies. The temperature anomalies are obtained from the GISTEMP Team, 2020: GISS Surface Temperature Analysis (GISTEMP), version 4, NASA Goddard Institute for Space Studies. The dataset was accessed on 18 October 2020 at https://data.giss.nasa.gov/gistemp/. See Lenssen et al. (2019) for details and the most recent description of the data. By definition, these time series measure deviations from the corresponding 1951-1980 means.

Appendix B Additional results

Table B.1. Transition risk and word selection robustness. The table reports the correlation (median and percentiles) between the benchmark transition risk measure and the ones based on 30000 random leave-one-word-out simulations.

Percentile	*Australia*	*Brazil*	*Canada*	*Malaysia*	*Mexico*	*Norway*	*Russia*	*SouthAfrica*
5%	0.93	0.86	0.94	0.90	0.94	0.95	0.91	0.93
Median	0.96	0.95	0.98	0.96	0.97	0.98	0.96	0.97
95%	0.98	0.97	0.99	0.98	0.98	0.99	0.98	0.99

Table B.2. Correlation of leads and lags of yearly transition risk (ΔTR) and remaining commodity reserves (ΔRR). Transition risk is converted to yearly numbers using monthly means. Commodity reserves are strongly trending. Correlations are computed using the first difference of the variables. For all countries, except South Africa, we use remaining oil reserves. For South Africa, which produces very little oil, remaining coal reserves is used. *, **, and *** denote the 10%, 5%, and 1% significance level, respectively.

Percentile	*Australia*	*Brazil*	*Canada*	*Malaysia*	*Mexico*	*Norway*	*Russia*	*SouthAfrica*
$\Delta TR_{t-1}, \Delta RR_t$	-0.33	0.38	0.13	-0.05	-0.01	-0.46*	0.10	0.27
$\Delta TR_t, \Delta RR_t$	-0.08	0.00	0.13	-0.29	-0.20	0.36	-0.22	0.27
$\Delta TR_t, \Delta RR_{t-1}$	0.05	0.05	-0.28	0.34	0.03	-0.14	-0.29	-0.55**

Table B.3. Correlation of leads and lags of yearly transition risk (TR) and CCPI ($CCPI$). Transition risk is converted to yearly numbers using monthly means. *, **, and *** denote the 10%, 5%, and 1% significance level, respectively.

Percentile	*Australia*	*Brazil*	*Canada*	*Malaysia*	*Mexico*	*Norway*	*Russia*	*SouthAfrica*
$TR_{t-1}, CCPI_t$	0.18	0.59**	-0.08	0.14	0.33	0.21	0.38	0.58**
$TR_t, CCPI_t$	0.43	0.76***	-0.17	0.36	0.25	0.10	0.54**	0.26
$TR_t, CCPI_{t-1}$	0.29	0.74***	-0.29	-0.09	0.02	0.00	0.15	-0.26

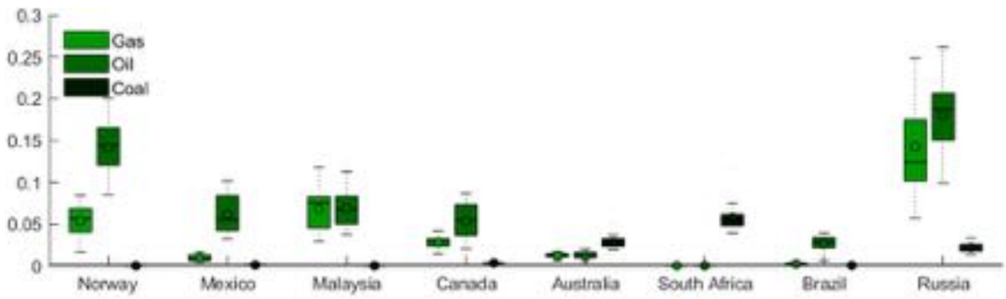

Figure B.1. Gas, oil, and coal production relative to GDP. For each country, the figure reports a standard box plot of the production shares for the period 2002 to 2019. The underlying data is sourced from British Petroleum Company (2020).

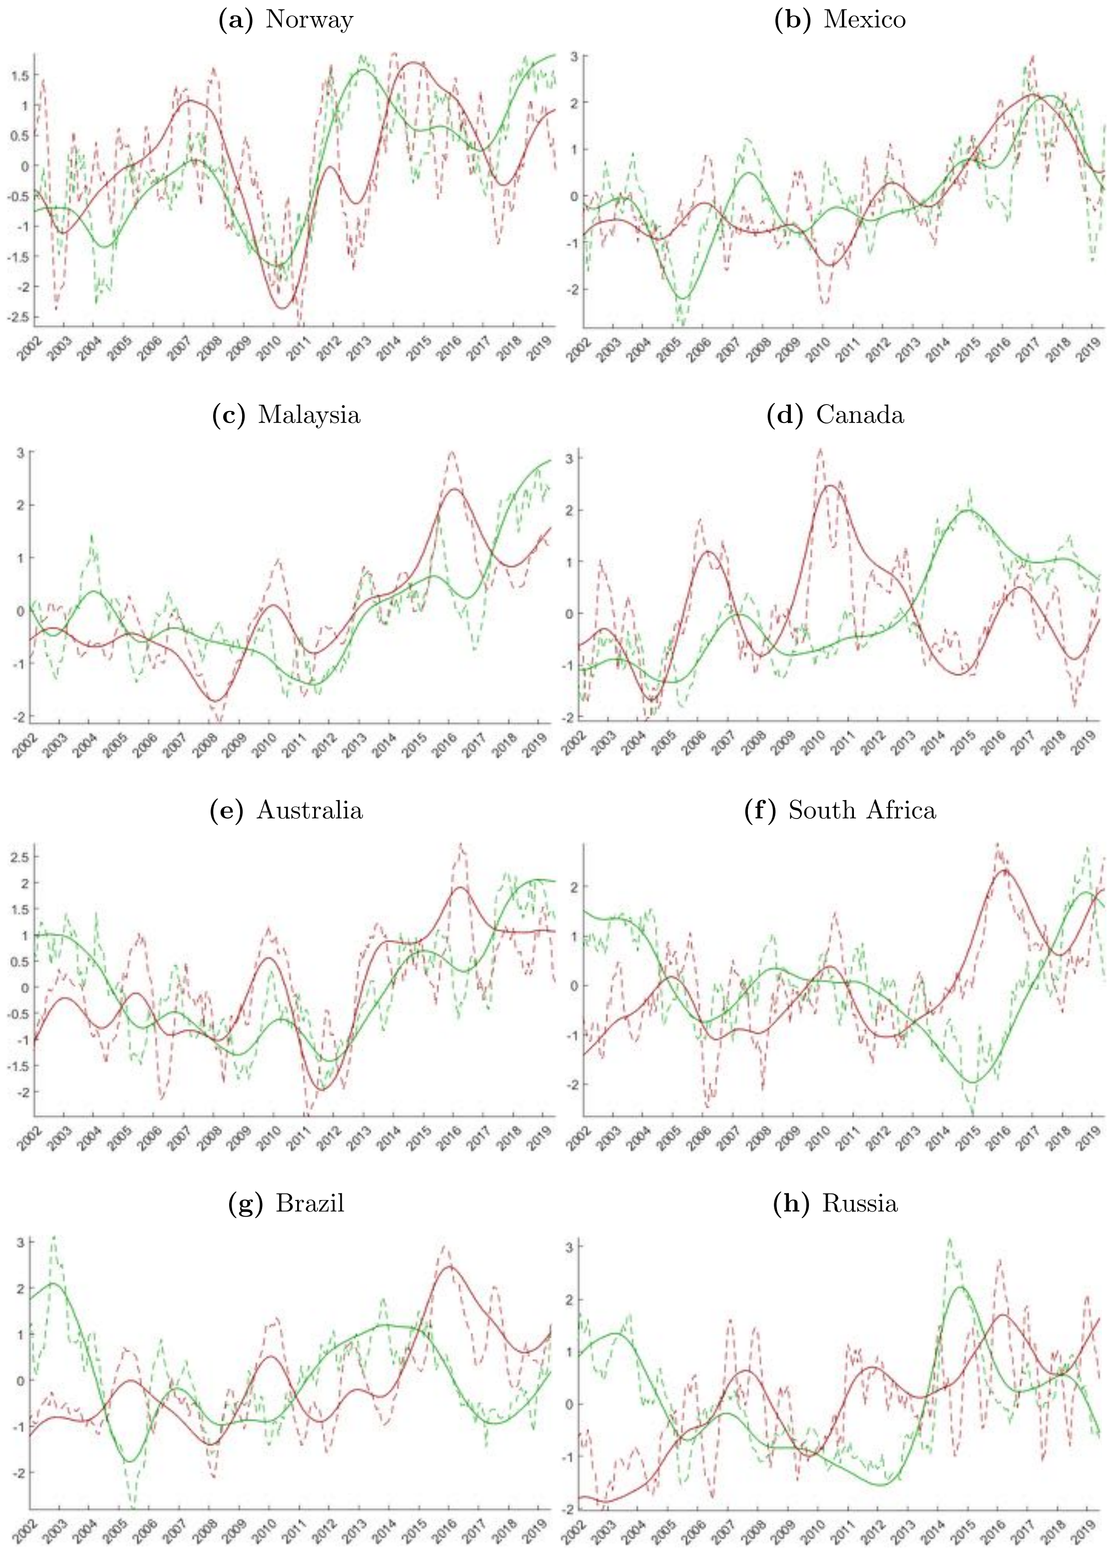

Figure B.2. Transition risk (green) and temperature anomalies (red). The dotted lines report the raw series. The solid lines report the data when a Hodrick–Prescott filter (Hodrick and Prescott (1997)), with a smoothing parameter set to 1600, is used to extract the low-frequency fluctuations from the series.

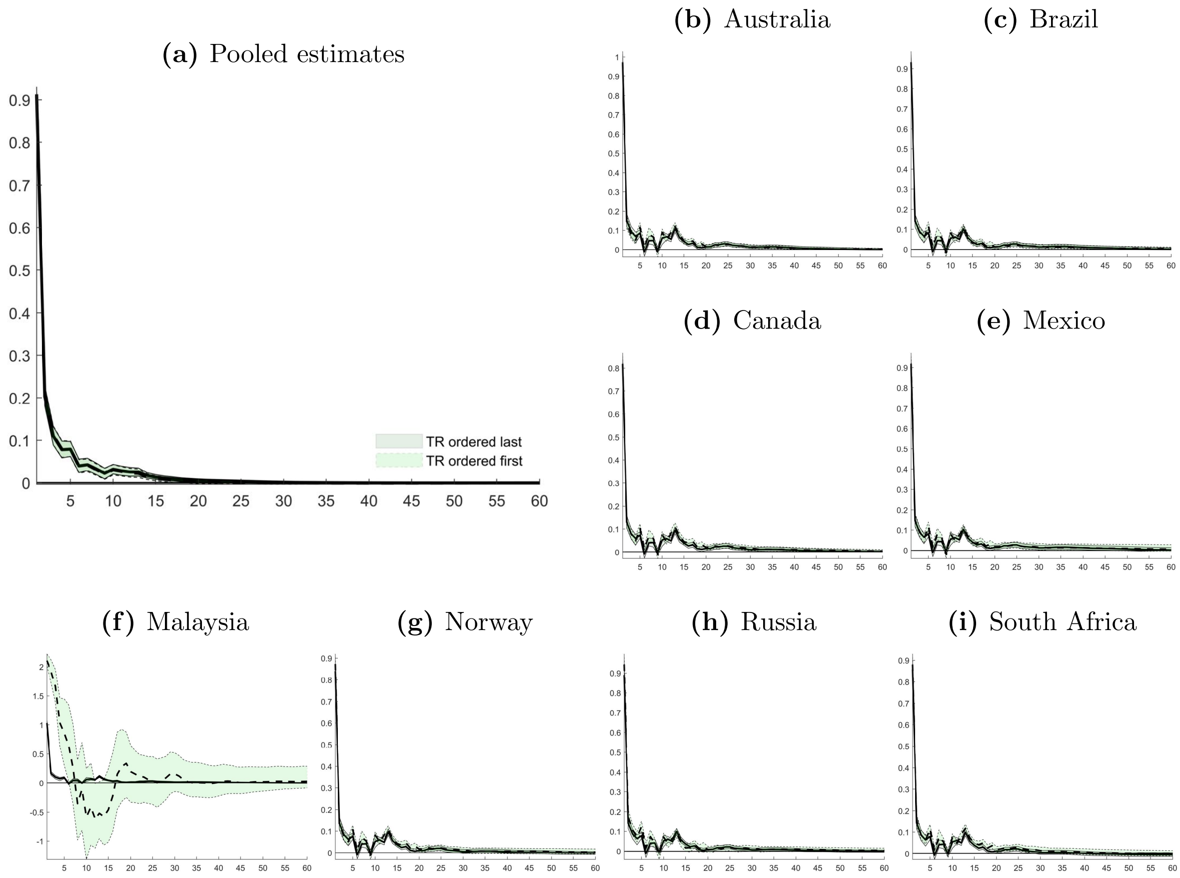

Figure B.3. Pooled and partially pooled panel VAR results. Each graph reports the transition risk response following a one standard deviation exogenous innovation to the transition risk variable. The innovations are computed from two different recursive orderings, where the transition risk variable is ordered either first (dotted black) or last (solid black) in the system. The color shaded areas are 68% probability bands.

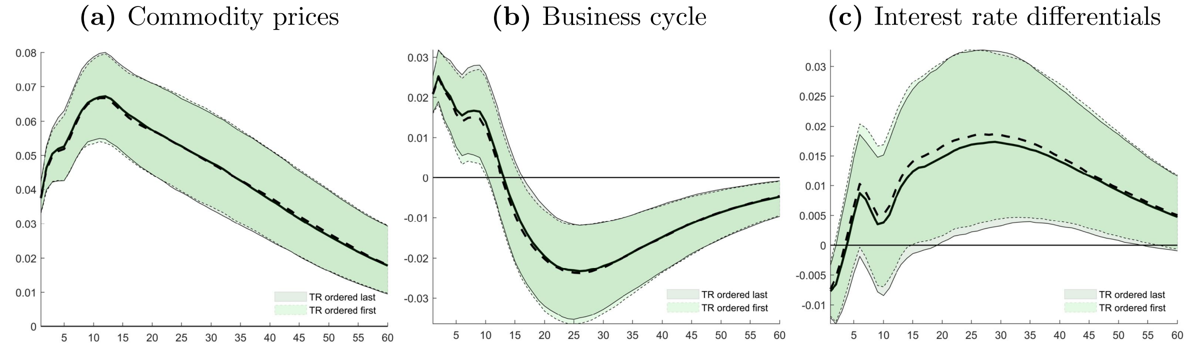

Figure B.4. Pooled VAR and REER responses. Each graph reports the REER response following a one standard deviation exogenous innovation to commodity prices, the business cycle index, or interest rate differentials. The innovations are computed from two different recursive orderings, where the transition risk variable is ordered either first (solid black) or last (dotted black) in the system. The color shaded areas are 68% probability bands.

(a) Commodity prices (b) Business cycle (c) Interest rate differentials

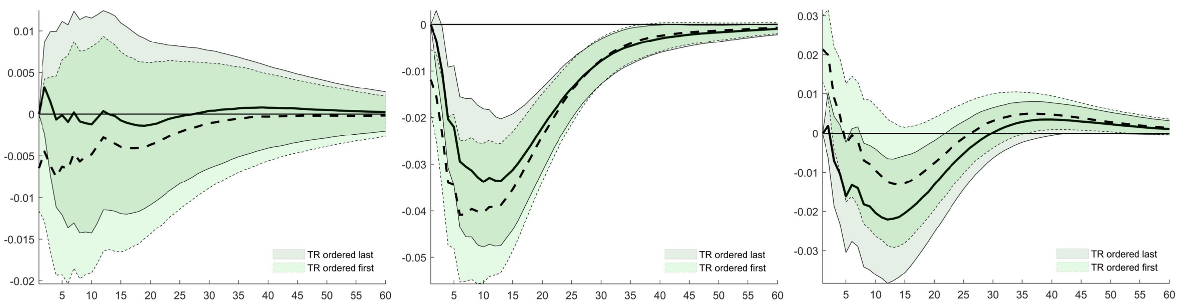

Figure B.5. Pooled VAR and macroeconomic responses. The graphs report the responses of commodity prices, the business cycle index, and interest rate differentials, following a one standard deviation exogenous innovation to transition risk. The innovations are computed from two different recursive orderings, where the transition risk variable is ordered either first (solid black) or last (dotted black) in the system. The color shaded areas are 68% probability bands.

(a) Global activity (b) Economic policy uncertainty (c) Oil production and the stock market

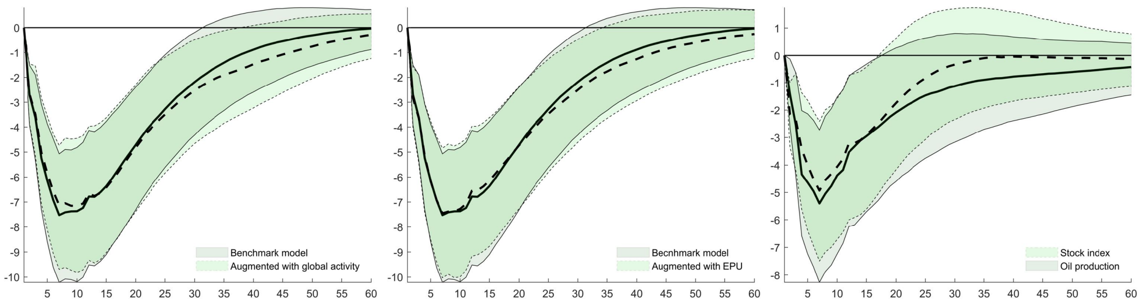

Figure B.6. Pooled VAR and REER responses. Each graph reports the real exchange rate response following a one standard deviation exogenous innovation to the transition risk variable. The transition risk variable is ordered last in the system. Each VAR is augmented with the global activity measure proposed by Baumeister and Hamilton (2019), the economic uncertainty indexes (EPU) developed by Baker et al. (2016), oil production or the aggregated stock market indexes. The color shaded areas are 68% probability bands.

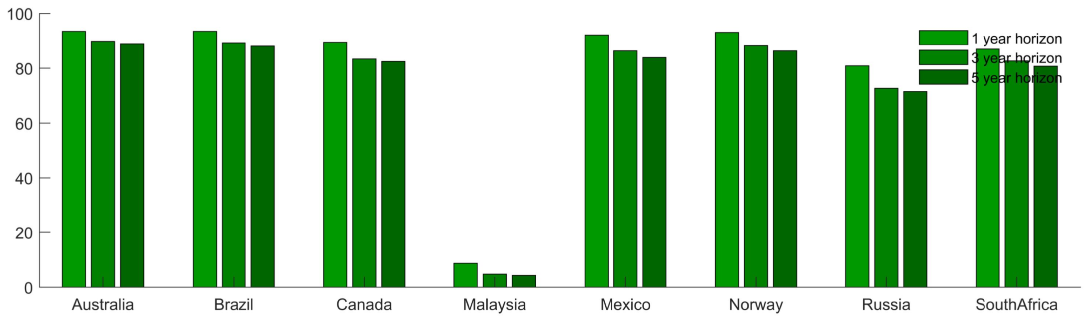

Figure B.7. Partially pooled panel VAR results. Transition risk and the variance explained by transition risk innovations. Each bar reports the median estimate for a given horizon and country. Estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system.

(a) Engle et. al. - 1-year horizon

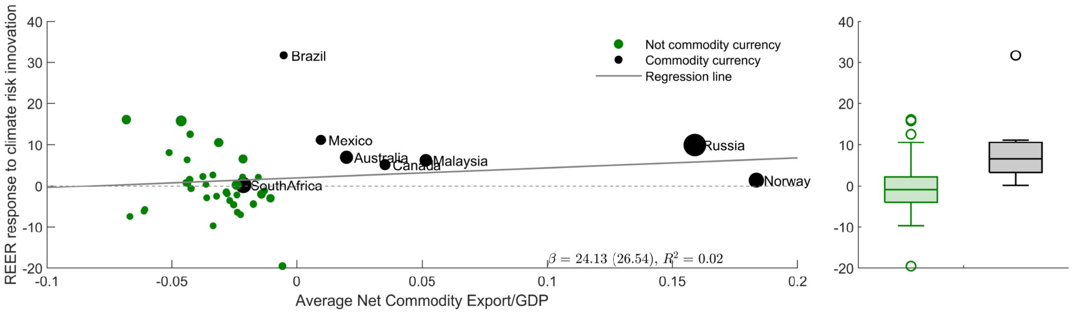

(b) Gavriilidis - 1-year horizon

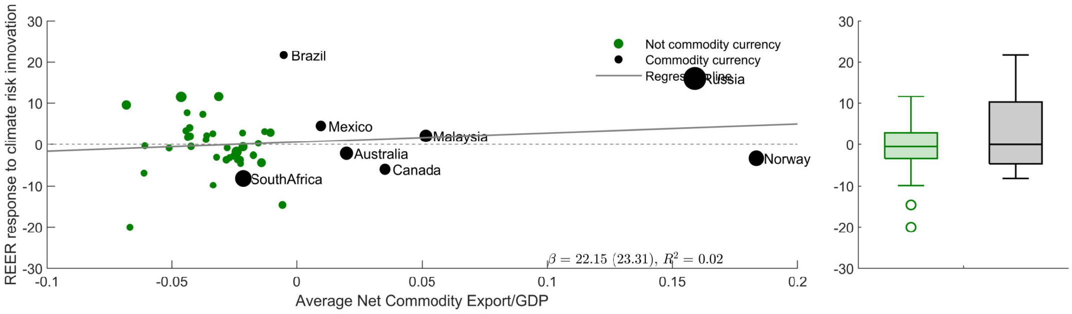

(c) Temperature anomalies - 1-year horizon

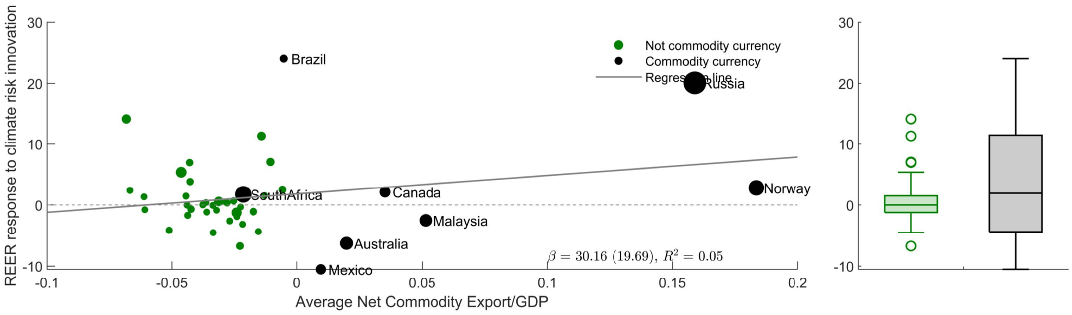

Figure B.8. REER responses for alternative climate risk proxies and commodity export shares. Each graph reports a country's REER response following a one standard deviation exogenous innovation to the Engle et al. (2020) climate risk index, the climate policy uncertainty index developed by Gavriilidis (2021), or temperature anomalies (y-axis: in percentage change) together with net commodity exports relative to GDP (x-axis). The REER response estimates are obtained assuming a recursive ordering with the climate risk variable ordered last in the VAR system. See Figure 6 for additional details.

Appendix C Panel VAR details

Below we provide a short technical description of the pooled panel VAR estimation routines. We start by describing the random effects specification, and then turn to the fully pooled panel VAR specification.

C.1 Partially pooled panel VAR model

First, rewriting (4) as a SUR system in vectorized form allowing for cross-sectional heterogeneity:

$$\boldsymbol{y}_c = \bar{\boldsymbol{X}}_c \boldsymbol{\beta}_c + \boldsymbol{\varepsilon_c} \quad \boldsymbol{\varepsilon}_c \sim N(0, \bar{\boldsymbol{\Sigma}}_c) \text{ with } \bar{\boldsymbol{\Sigma}}_c = \boldsymbol{\Sigma}_c \otimes \boldsymbol{I}_T \tag{5}$$

with

$$\boldsymbol{y}_c = \underbrace{vec(\boldsymbol{Y}_c)}_{nT\times 1}, \quad \bar{\boldsymbol{X}}_c = \underbrace{(\boldsymbol{I}_n \otimes \boldsymbol{X}_c)}_{nT\times q}, \quad \boldsymbol{\beta}_c = \underbrace{vec(\boldsymbol{B}_c)}_{q\times 1}, \quad \boldsymbol{\varepsilon}_c = \underbrace{vec(\boldsymbol{\epsilon}_c)}_{nT\times 1} \tag{6}$$

where n is the number of endogenous variables, T the sample size, $q = nk = n(np + m)$, m is the number of exogenous variables, and

$$\boldsymbol{Y}_c = \underbrace{\begin{pmatrix} \boldsymbol{y}'_{c,1} \\ \boldsymbol{y}'_{c,2} \\ \vdots \\ \boldsymbol{y}'_{c,T} \end{pmatrix}}_{T\times n}, \quad \boldsymbol{X}_c = \underbrace{\begin{pmatrix} \boldsymbol{y}'_{c,0} & \dots & \boldsymbol{y}'_{c,1-p} & \boldsymbol{x}'_o \\ \boldsymbol{y}'_{c,1} & \dots & \boldsymbol{y}'_{c,2-p} & \boldsymbol{x}'_1 \\ \vdots & \ddots & \vdots & \vdots \\ \boldsymbol{y}'_{c,T-1} & \dots & \boldsymbol{y}'_{c,T-p} & \boldsymbol{x}'_T \end{pmatrix}}_{T\times k}, \quad \boldsymbol{B}_c = \underbrace{\begin{pmatrix} (\boldsymbol{A}_c^1)' \\ \vdots \\ (\boldsymbol{A}_c^p)' \\ \boldsymbol{C}'_c \end{pmatrix}}_{k\times n}, \quad \boldsymbol{\epsilon}_c = \underbrace{\begin{pmatrix} \boldsymbol{\epsilon}'_{c,1} \\ \boldsymbol{\epsilon}'_{c,2} \\ \vdots \\ \boldsymbol{\epsilon}'_{c,T} \end{pmatrix}}_{T\times n} \tag{7}$$

In total the model specification in (5) implies that each unit comprises q coefficients to estimate. With N units in total, Nq coefficients have to be estimated for the whole model. Thus, to take advantage of the cross sectional information in the data we assume a random effects specification where for each unit c, $\boldsymbol{\beta}_c$ can be expressed as $\boldsymbol{\beta}_c = \boldsymbol{b} + \boldsymbol{b}_c$ and $\boldsymbol{b}_c \sim N(\boldsymbol{b}, \boldsymbol{\Sigma}_b)$. It then follows that:

$$\boldsymbol{\beta}_c \sim N(\boldsymbol{b}, \boldsymbol{\Sigma}_b) \tag{8}$$

i.e., VAR coefficients differ across units, but are drawn from a distribution with similar mean and variance. We implement this using the hierarchical prior approach developed by Jarociński (2010).

For $\boldsymbol{b}$ the selected functional form is simply a diffuse (improper) prior $\pi(\boldsymbol{b}) \propto 1$. For $\boldsymbol{\Sigma}_b$ the functional form is designed to replicate the Minnesota coefficient covariance matrix prior. This specification relies on a diagonal $q \times q$ covariance matrix $\boldsymbol{\Omega}_b$ with elements:

$$\sigma^2_{a_{ii}} = (\frac{1}{l^{\lambda_3}})^2, \quad \sigma^2_{a_{ij}} = (\frac{\sigma^2_i}{\sigma^2_j})(\frac{\lambda_2}{l^{\lambda_3}})^2, \quad \sigma^2_{d_i} = \sigma^2_i(\lambda_4)^2 \tag{9}$$

relating the variance of $\boldsymbol{\beta}_c$ to the own lags of endogenous variables (a_{ii}), cross-lag coefficients (a_{ij}), and exogenous variables (d_i). $\sigma_\cdot^2$ are scaling parameters obtained by fitting autoregressive models by OLS for the n endogenous variables of the model, and computing their standard deviations, while the λ's are set to values typically found in the literature, i.e., $\lambda_2 = 0.5$, $\lambda_3 = 1$, and $\lambda_4 = 10^2$. The full covariance matrix $\boldsymbol{\Sigma}_b$ is then defined as:

$$\boldsymbol{\Sigma}_b = (\lambda_1 \otimes \boldsymbol{I}_q)\boldsymbol{\Omega}_b \tag{10}$$

where $\boldsymbol{\Omega}_b$ is treated as fixed and known, and the role of λ_1 is discussed below. Finally, the prior distribution for $\boldsymbol{\Sigma}_c$ is simply the classical diffuse prior given by $\pi(\boldsymbol{\Sigma}_c) \propto |\boldsymbol{\Sigma}_c|^{-(n+1)/2}$.

Conceptually, the difference between pooled and random effects estimation is determined by λ_1. Setting $\lambda_1 = 0$ in (10) implies that all the $\boldsymbol{\beta}_c$'s take the identical value $\boldsymbol{b}$, i.e., data is fully pooled. In contrast, treating λ_1 as a random variable allows for cross-sectional heterogeneity. In this case we use the inverse Gamma distribution as a prior distribution for λ_1, implying $\pi(\lambda_1|s_0/2, v_0/2) \propto \lambda^{\frac{-s_0}{2}-1} exp(\frac{-v_o}{2\lambda_2})$, with shape $s_0/2$ and scale $v_0/2$, and set $s_0 = v_0 = 0.002$, which we experience gives a reasonable balance between individual (large λ_1) and pooled (small λ_1) estimates.

In the case of a (fixed) $\lambda_1 = 0$, draws from the posterior distributions can be obtained from its analytical solution. When the random effects specification is adopted, the posterior distributions do not allow for any analytical derivations, and a Gibbs sampler framework is used to draw from the appropriate conditional posterior distributions. Details about each of these cases are well documented in, e.g., Kadiyala and Karlsson (1997), Jarociński (2010), and Canova and Ciccarelli (2013), and also shortly described in below. Here we note that we obtain 100000 draws from the posterior, use the last 2000 for further inference, and ensure that the systems are invertible by disregarding draws implying non-stationarity.

C.2 Gibbs sampler for the partially pooled panel VAR

The model's unknown parameters are $\boldsymbol{b}$, λ_1, $\boldsymbol{\beta}_c$, and $\boldsymbol{\Sigma}_c$. The posterior is approximated by making draws from the following sequence of conditional posterior distributions, where d denote the d^{th} draw:

1. Draw $\boldsymbol{b}^d$ from a multivariate normal distribution:

$$\boldsymbol{b}^d \sim N(\boldsymbol{\beta}_m^{d-1}, N^{-1}\boldsymbol{\Sigma}_b^{d-1}) \text{ with } \boldsymbol{\beta}_m = N^{-1}\sum \boldsymbol{\beta}_c^{d-1}$$

2. Draw λ_1^d from an inverse Gamma distribution:

$$\lambda_1^d \sim IG(\frac{\bar{s}}{2}, \frac{\bar{v}}{2}) \text{ with } \bar{s} = h + s_0 \text{ and } \bar{v} = v_0 + \sum((\boldsymbol{\beta}_c^{d-1} - \boldsymbol{b}^d)'(\boldsymbol{\Omega}_b^{-1})(\boldsymbol{\beta}_c^{d-1} - \boldsymbol{b}^d))$$

and obtain $\boldsymbol{\Sigma}_b^d = (\lambda_1^d \otimes \boldsymbol{I}_q)\boldsymbol{\Omega}_b$

3. Draw $\boldsymbol{\beta}_c^d$ for each unit from a multivariate normal distribution:

$$\boldsymbol{\beta}_c^d \sim N(\bar{\boldsymbol{\beta}}_c, \bar{\boldsymbol{\Omega}}_c)$$

with

$$\bar{\boldsymbol{\Omega}}_c = \left[(\boldsymbol{\Sigma}_c^{d-1})^{-1} \otimes \boldsymbol{X}_c'\boldsymbol{X}_c + (\boldsymbol{\Sigma}_b^d)^{-1}\right]^{-1}$$
$$\bar{\boldsymbol{\beta}}_c = \bar{\boldsymbol{\Omega}}_c \left[(\boldsymbol{\Sigma}_c^{d-1})^{-1} \otimes \boldsymbol{X}_c')\boldsymbol{y}_c + (\boldsymbol{\Sigma}_b^d)^{-1}\boldsymbol{b}^d\right]$$

4. Draw $\boldsymbol{\Sigma}_c^d$ for each unit the inverse Wishart ditribution:

$$\boldsymbol{\Sigma}_c^d \sim IW(\tilde{\boldsymbol{S}}_c, T) \text{ with } \tilde{\boldsymbol{S}}_c = (\boldsymbol{Y}_c - \boldsymbol{X}_c\boldsymbol{B}_c^d)'(\boldsymbol{Y}_c - \boldsymbol{X}_c\boldsymbol{B}_c^d)$$

As starting values, i.e., for $d = 1$, we set $\boldsymbol{\beta}_c^0$ and $\boldsymbol{\Sigma}_c^0$ equal to the implied OLS values, and $\lambda_1^0 = 0.01$.

C.3 Pooled panel VAR model

For the fully pooled panel VAR model a natural conjugate normal-Wishart prior is used when estimating the model. First, define:

$$\boldsymbol{Y}_t = \begin{pmatrix} \boldsymbol{y}_{1,t}' \\ \boldsymbol{y}_{2,t}' \\ \vdots \\ \boldsymbol{y}_{N,t}' \end{pmatrix}, \quad \boldsymbol{X}_t = \begin{pmatrix} \boldsymbol{y}_{1,t-1}' & \cdots & \boldsymbol{y}_{1-p,t}' & \boldsymbol{x}_t' \\ \boldsymbol{y}_{2,t-1}' & \cdots & \boldsymbol{y}_{2,t-p}' & \boldsymbol{x}_t' \\ \vdots & \ddots & \vdots & \vdots \\ \boldsymbol{y}_{N,t-1}' & \cdots & \boldsymbol{y}_{N,t-p}' & \boldsymbol{x}_t' \end{pmatrix}, \quad \boldsymbol{B} = \begin{pmatrix} (\boldsymbol{A}^1)' \\ \vdots \\ (\boldsymbol{A}^p)' \\ \boldsymbol{D}' \end{pmatrix}, \quad \boldsymbol{\epsilon}_c = \begin{pmatrix} \boldsymbol{\xi}_{1,t}' \\ \boldsymbol{\xi}_{2,t}' \\ \vdots \\ \boldsymbol{\xi}_{N,t}' \end{pmatrix} \tag{11}$$

Then, stacking (11) over T time periods one gets $\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{\xi}$, and writing this expression in vectorised form gives:

$$\boldsymbol{y} = \bar{\boldsymbol{X}}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \quad \boldsymbol{\varepsilon} \sim N(0, \bar{\boldsymbol{\Sigma}}) \text{ with } \bar{\boldsymbol{\Sigma}} = \boldsymbol{\Sigma} \otimes \boldsymbol{I}_{NT} \tag{12}$$

with

$$\boldsymbol{y} = \underbrace{vec(\boldsymbol{Y})}_{NnT\times 1}, \quad \bar{\boldsymbol{X}} = \underbrace{(\boldsymbol{I}_n \otimes \boldsymbol{X})}_{NnT\times q}, \quad \boldsymbol{\beta} = \underbrace{vec(\boldsymbol{B})}_{q\times 1}, \quad \boldsymbol{\varepsilon} = \underbrace{vec(\boldsymbol{\xi})}_{NnT\times 1} \tag{13}$$

For the normal-Wishart prior specification, the prior for $\boldsymbol{\beta}$ is assumed to be multivariate normal:

$$\boldsymbol{\beta} \sim N(\boldsymbol{\beta}_0, \boldsymbol{\Sigma} \otimes \Phi_0) \tag{14}$$

where the elements of $\boldsymbol{\beta}_0$ are set to 0.8 for the first lag of own endogenous variables and zero otherwise, and Φ_0 is as a $k \times k$ diagonal matrix with entries defined as in Karlsson (2013):

$$\sigma_{a_{ij}}^2 = \left(\frac{1}{\sigma_j^2}\right)\left(\frac{\lambda_1}{l^{\lambda_3}}\right)^2, \quad \sigma_d^2 = (\lambda_1\lambda_4)^2 \tag{15}$$

where the residual variance terms are defined by estimating a pooled autoregressive model over the each of the n endogenous variables. For the fully pooled VAR we follow the usual convention and set $\lambda_1 = 0.1$, $\lambda_3 = 1$, and $\lambda_4 = 10^2$ (i.e., λ_1 is treated very differently here than in the partially pooled Panel VAR model).

The prior for $\boldsymbol{\Sigma}$ is inverse Wishart:

$$\boldsymbol{\Sigma} \sim IW(\boldsymbol{S}_0, \alpha_0) \text{ with } \boldsymbol{S}_0 = (\alpha_0 - n - 1)\boldsymbol{\Sigma_0} \tag{16}$$

where $\alpha_0 = n + 2$ and $\boldsymbol{\Sigma_0}$ is a diagonal matrix with variance terms obtained as above. As such, the covariance matrix of one equation is now proportional to the covariance matrix of the other equations, which is not a restriction in the partially pooled specification.

Because these priors are conjugate, draws from the posterior distribution can be obtained from analytical solutions. In particular:

$$\begin{aligned} &\pi(\boldsymbol{\Sigma}|\boldsymbol{y}) \sim IW(\bar{\alpha}, \bar{\boldsymbol{S}}) \\ &\pi(\boldsymbol{\beta}|\boldsymbol{y}) \sim MT(\bar{\boldsymbol{B}}, \bar{\boldsymbol{S}}, \bar{\boldsymbol{\Phi}}, \tilde{\alpha}) \end{aligned} \tag{17}$$

with

$$\begin{aligned} \bar{\boldsymbol{\Phi}} &= \left[\boldsymbol{\Phi}_0^{-1} + \boldsymbol{X}'\boldsymbol{X}\right]^{-1} \\ \bar{\boldsymbol{B}} &= \bar{\boldsymbol{\Phi}} \left[\boldsymbol{\Phi}_0^{-1}\boldsymbol{B}_0 + \boldsymbol{X}'\boldsymbol{Y}\right]^{-1} \\ \bar{\boldsymbol{S}} &= \boldsymbol{Y}'\boldsymbol{Y} + \boldsymbol{S}_0 + \boldsymbol{B}_0'\boldsymbol{\Phi}_0^{-1}\boldsymbol{B}_0 - \boldsymbol{B}'\bar{\boldsymbol{\Phi}}^{-1}\boldsymbol{B} \end{aligned} \tag{18}$$

and $\bar{\alpha} = NT + \alpha_0$ and $\tilde{\alpha} = \bar{\alpha} - n + 1$.