Climate change and commodity currencies: Measuring transition risk with word embeddings*

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

This version: May 13, 2025

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 measure of climate change transition risk, based on neural network word embedding models for large-scale text analysis, and document that when it unexpectedly increases, major commodity currencies experience a persistent depreciation in line with economic theory. Expanding the analysis to a richer set of countries confirms a negative correlation between a country's carbon export dependency and exchange rate response to transition risk. Word embeddings have been crucial for scientific advances and improvements on down stream tasks in the Natural Language Processing literature over the last decade. Our study shows how they can be used to quantify an important but hard-to-measure concept in economics.

JEL-codes: C33, C82, E71, F41, Q54

Keywords: Exchange Rates, Climate, Risk, Commodities

*Comments from Jørgen J. Andersen, Hilde C. Bjørnland, Gunnar Bårdsen, Kasper Roszbach, 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, Espen R. Moen, Gisle Natvik, and Dagfinn Rime 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

‡Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School. Email: Vegard.H.Larsen@bi.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 the cost and opportunities following adjustments towards a lower-carbon economy. Perhaps due to ease of measurement, the majority of the 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. However, a growing literature is now starting to focus more on the transition risk component (Atanasova and Schwartz (2019), van der Ploeg and Rezai (2020), Sen and von Schickfus (2020), Gavriilidis (2021), Ardia et al. (2022)). The policy relevance of this focus is well illustrated by Norges Bank's reflections on recent exchange rate developments in 2019:

> *"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."(Norges Bank Monetary Policy Report 3/2019)*

The challenge is how to measure climate change transition risk to test conjectures of this type. In this article, we propose a novel measure of country-specific transition risk and analyze how its exogenous fluctuations affect carbon-intensive commodity currencies. In doing so, our contribution is foremost about measurement, but also about how climate risk affects valuations at a national level, and not only firm value, which has been the dominant research entity in existing research.

A simple economic definition, imposing a minimal set of assumptions, guides how we construct the transition risk measure and highlights why it should be particularly relevant in the current context. In particular, it follows from the present-value formulation of the real exchange rate that it embeds expectations about future macro fundamentals and currency risk premiums (Engel and West (2005), Froot and Ramadorai (2005)). In terms of the former, the well-known Balassa-Samuelson hypothesis (Balassa (1964), Samuelson (1964)) and the related Dutch disease theory (Bruno and Sachs (1982), Eastwood and Venables (1982), Corden (1984), van Wijnbergen (1984), Torvik (2001)) are particularly relevant for commodity exporters. Both theories predict a persistent appreciation of the real exchange rate following productivity shocks in the traded sector and windfall gains. However, news about potential future realizations of these factors affect the real exchange rate already today because it affects (long-run) macro expectations.[1] As such, when

[1] See, e.g., Charnavoki and Dolado (2014), Bjørnland and Thorsrud (2016), Arezki et al. (2017), Bjørnland

climate-related news perceived as unfavorable for future (commodity) production potential arrives, reflecting an increase in transition risk, the exchange rate should depreciate already today.[2]

In terms of the risk premium, assets that are particularly exposed to adjustment towards a lower-carbon economy should earn a premium because they are expected to deliver low returns in a low-carbon future. For high-carbon commodity exporters, this pulls the exchange rate in the same direction as the macro effect. However, as transition risk, by definition, is associated not only with the costs but also opportunities associated with change, hedging investments that pay off primarily in a low-carbon future should be discounted at a rate below the risk-free rate. Accordingly, in response to new information, and considering the effect stemming from the risk premium in isolation, the real exchange rate of a commodity exporter might appreciate on impact if investors believe that the content of the news provides a hedge against a low-carbon future state.[3]

While the risk premium and macroeconomic effects potentially pull in opposite directions, it seems reasonable to assume that the latter effect dominates. The costs of transitioning to a lower-carbon economy are arguably higher for countries whose economy is highly dependent on carbon-intensive income relative to countries where income is less carbon-intensive. Still, the distinction is useful because it illustrates the duality, or cost and opportunities perspective, embedded in the transition risk definition. 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 expectations change due to new information and that the media operate as "information intermediaries" between agents and the state of the world (Nimark and Pitschner (2019), Larsen and Thorsrud (2019), Larsen et al. (2021)).[4] However, as alluded to above, transition risk can change in response to various news events, ranging from discussion and implementation of actual policies to news that

et al. (2019), and (Harding et al., 2020) for recent empirical support for these theories.

[2]Classical examples of such news is the (potential) introduction of new carbon taxes or governmental laws restricting future commodity exploration (e.g., drilling licenses). More generally, it is by now well known that carbon budgets compatible with conventional temperature targets imply that new investments in high-carbon capital should be rapidly discontinued and existing production technologies scaled down or retrofitted at a cost (Campiglio and der Ploeg, 2021). However, the actual implementation and public support related to such policies are subject to time variation.

[3]Topical (Norwegian) examples are news reflecting an increased willingness among the public or government to allow for large (off-shore) windmill or solar investments. In terms of stocks, Pástor et al. (2022) makes a related argument, arguing that the discount rate effect incentivises companies to become greener because it increases their current market value (i.e., lowers their expected return and cost of capital).

[4]It is well 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)).

reflects changes in investor and consumer preferences and even more silent features related to systematic directional modification of ideas and narratives as they are spread in the public discourse (Shiller (2017), Hirshleifer (2020)). This makes filtering the news in search of specific events very challenging. Instead, we focus on a smaller set of more general keywords, such as "green", "uncertainty", "economy", and "petroleum", and quantify how these are used in context with each other across time.

Guided by the duality of the transition risk definition, we then let our news-based measure of transition risk consist of two components. The first component reflects the opportunity perspective of transition risk, capturing the extent to which a country is perceived as "green" rather than "brown". The second component reflects the cost perspective and captures the uncertainties related to the (influential) carbon-intensive sector of the economy. Going forward, we denote these as the "green transition dimension" and the "carbon-economy dimension" of transition risk.

We operationalize the above reasoning using a unique and large corpus 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 been crucial for scientific advances and improvements on down stream tasks in the Natural Language Processing literature (NLP) over the last decade.[5] The reason is that they capture well shared context of words in the corpus, 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 sum of word vectors representing the two transition risk dimensions and regress this aggregated word vector on word vectors for each country. To the extent that the aggregated word vector points in the intended transition risk dimension in vector space, the parameter estimates of these regressions measure how strong the association between a given country and climate change transition risk is and how it varies across time.

At an intuitive level, time variation in our suggested measures can be due to changes in how climate change-related words are used in relation to a country or because these words are used differently across time. Words such as "climate" and "green" might, for example, be used much more in relation to the words "risk" and "economy" today than a decade ago or when the flow of climate-related news is particularly high. The word embedding methodology we apply is designed to capture exactly these types of changes.

[5] Word embeddings are now important features in close to all Deep Learning algorithms performing tasks such as language modeling, text classification, translation, question answering and named entity recognition.

Performing an intrinsic evaluation of the estimated transition risk embedding itself suggests that it is associated with words we a priori would have expected it to be associated with, that these words together seem to form a cluster distinct from other related concepts, and that one can use the transition risk embedding to achieve reasonably good out-of-sample classification performance.

Turning to the down-stream task of analyzing the relationship between transition risk and real exchange rates, we focus on the commonly used commodity currencies of Australia, Brazil, Canada, Malaysia, Mexico, Norway, Russia, and South Africa and use Vector Autoregressive models to capture the unpredictable part of transition risk.[6]

Accounting for the dynamic interactions between commodity prices, interest rate differentials, business cycle developments, and various uncertainty measures, we show that exogenous transition risk innovations generally lead to a significant and persistent exchange rate depreciation, which explains roughly four percent of the long-run variation in the real exchange rates. Consistent with the underlying mechanism we build on, transition risk shocks also tend to cause persistently lower aggregate stock market valuations and reduced commodity supply. Likewise, including a rich set of other countries in the analysis shows a significant negative correlation between a country's carbon export dependency and the exchange rate response following transition risk shocks. When dissecting the two-part transition risk construction, we document that the two components of transition risk indeed have opposite effects on the real exchange rate. In line with the present-value formulation, shocks to the "green transition dimension" lead to an appreciation and a fall in expected returns. Finally, when comparing our suggested transition risk measures with commonly used climate risk alternatives, that are not explicitly aimed at capturing transition risk, we find that these indexes either do not matter much in the current context or that they likely capture only one dimension of transition risk.

Climate change is a significant global issue, and understanding how its associated risks affect economic and financial valuations is vital for market participants and policymakers. Naturally, there is extensive literature on the topic. Thus far, however, most of this literature has been concerned with firm value.[7] We contribute by proposing a methodology

[6]These countries produce a substantial amount of carbon-intensive commodities (Figure C.1 in Appendix C) 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.

[7]See, e.g., Bolton and Kacperczyk (2021), Hsu et al. (2022), 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). In relation to commodity producers, the 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)). The recent exploratory analysis by Cha et al. (2020) shares our focus on the foreign exchange market, but study physical climate risk shocks.

for measuring the transition risk component of climate risk and evaluating how it affects real exchange rate fluctuations. Such fluctuations are important from a macroeconomic perspective and the conduct of monetary policy in particular.

While our application is context specific, the methodology and intuition are general and potentially valuable for a broader set of applications. As such, this article speaks to a growing literature using text as data and tools from NLP and Machine Learning (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 Baker et al. (2016), Thorsrud (2018), Hansen et al. (2018), and Angelico et al. (2022) use topic models or Boolean search techniques to measure uncertainty, business cycle developments, monetary policy, and inflation expectations. In particular, this article relates to recent work by, e.g., Engle et al. (2020), Gavriilidis (2021), and Ardia et al. (2022), who all propose news-based climate risk measures using dictionary-based methods in combination with Boolean search or term frequency–inverse document frequency calculations. However, none of these indexes is designed to explicitly capture the country-specific transition risk component of climate risk. As discussed above, this component seems important for understanding 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. Section 4 expands the analysis by including a broad set of non-commodity currencies and using alternative climate risk proxies. Section 5 concludes.

2 Transition risk and measurement

Climate change transition risk is difficult to measure. First, changes in it can stem from many factors, ranging from new governmental regulations to shifts in preferences and demand among economic agents. Second, changes in expectations, that in turn affect economic outcomes, happen due to new information. However, new information about the variety of factors affecting transition risk is not available in commonly used databases. The second point above motivates why we use textual news data. By definition, the news industry operates as "information intermediaries" between agents and the state of the world and covers a variety of events and topics, including news relevant to transition risk. The first point above motivates why we propose to use a neural network word embedding model to quantify the high-dimensional textual news. Word embedding models utilize the co-occurrences of words in the corpus to predict neighboring words. As a by-product, words are represented as relatively small and dense vectors, with 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.

Second, when estimating word embedding models, running text can be used as implicit supervised training. This avoids the need for any hand-labeled supervision signal and makes the methodology flexible and user-friendly in many contexts. In contrast, popular NLP methods such as the Latent Dirichlet Allocation topic model, which has been applied in several recent economic studies (see, e.g., Larsen and Thorsrud (2019) and Hansen et al. (2018)), is an entirely unsupervised methodology where the user needs to define the meaning of the estimated topics ex-post. Similarly, applying commonly used word-count approaches would require that the researcher defines all the terms reflecting (country-specific) transition risk and that these terms are used more or less exclusively in the context of transition risk across time. By the nature of the problem, this would be a highly subjective and difficult task.

As an example, consider these two sentences sampled from news articles in August 2002 and April 2007, respectively: "...any worsening of the economic climate in Norway, particularly a further deterioration in the credit cycle..." and "...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, but only the latter is related to climate risk. A simple word-count-based method would easily count both as being about "climate" risk if using this keyword for search. In contrast, word embeddings capturing linguistic regularities and patterns, as well as the associated meaning between words, are better able to separate the two. In line with this, our results reported later also suggest that climate change transition risk was low in Norway in August 2002 and high in April 2007.

Finally, the word embedding logic, to some extent, also applies to the popular and widely used political uncertainty indexes first proposed by Baker et al. (2016). Like us, they focus on the shared context of words (such as "uncertainty", "political", and "economic"). However, whereas they use standard Boolean search techniques to count the co-occurrences of terms at the article level, we formulate the problem as predictive and use word embedding models to capture how words relate in a broader context. One reason for this choice is that it is unlikely that all the terms assumed to reflect country-specific transition risk occur in the same piece of text or article, making the number of counts potentially negatively biased when using Boolean search techniques. At the same time, widening the document definition to, e.g., one day, increases the risk of picking up counts that are unrelated to climate change transition risk. For example, there can much writing about "green" and "apples" that has nothing to do with transition risk, but that would still be counted when using one day as the document definition and standard Boolean search techniques.[8]

Below, in Sections 2.1 and 2.2, we describe in greater detail the DJ corpus and how we apply a word embedding model to construct quantitative and country-specific transition risk measures. Section 2.3 provides more intuition for the approach and an intrinsic evaluation of the estimated embedding space.

2.1 News data and word2vec

The DJ corpus consists of roughly 23 million news articles written in English, covering the period from 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 all news relevant to each country. Still, news stories relevant to 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*,s 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 vocabulary size 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 data blocks are used as input in a word embedding model. The famous

[8] Gavriilidis (2021) develops a climate policy uncertainty index (CPU) building on the method proposed in Baker et al. (2016) for measuring economic policy uncertainty. Since this is not country- or transition risk specific, the issue of picking up spurious counts is likely less pronounced for this measure. We compare our measure against it in Section 4.

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, evaluating if the center word w_c is likely to show up near the target word w_j to compute the classifier weights that will be the actual word embeddings.[9]

More formally, let w_i be a word from the vocabulary V, with size $|V|$, define a context window of size m, and assume a bigram model, where the probability of the sequence depends on the pairwise probability of a word in the sequence and the word next to it, as $P(w_{c-m}, w_{c-m+1}, \ldots, w_{c+m-1}, w_{c+m}) = \prod_{j=0,j\neq m}^{2m} P(w_{c-m+j}|w_c)$. The intuition for the skip-gram model is then to maximize this probability such 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 one target/center word pair (w_{c-m+j}, w_c), with vector representations $\boldsymbol{u}_{c-m+j}$ and $\boldsymbol{v}_c$, the likelihood is

$$L(\boldsymbol{\theta}) = log\frac{1}{1+e^{-\boldsymbol{u}'_{c-m+j}\cdot\boldsymbol{v}_c}} + \sum_{k=1}^{K} log\frac{1}{1+e^{\tilde{\boldsymbol{u}}'_k\cdot\boldsymbol{v}_c}}, \tag{1}$$

where $\boldsymbol{\theta}$ contains the latent word vectors, and the logistic (or sigmoid) function is used to turn the similarity measure between the word vectors for $\boldsymbol{v}_c$ and $\boldsymbol{u}_{c-m+j}$ into probabilities. The last term in (1) 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 must be generated and added to the data. These terms are commonly called noise terms ($\tilde{\boldsymbol{u}}_k$). For each target word, it is common to add K noise words.

Maximizing (1), and estimating the latent word vectors, is done using a two-layered neural network. $\boldsymbol{V} \in \mathbb{R}^{n\times|V|}$ is the parameter matrix in the first layer, with column $\boldsymbol{v}_i$ the input vector representation of word w_i (word embedding). $\boldsymbol{U} \in \mathbb{R}^{n\times|V|}$ is the parameter matrix in the second layer, with row $\boldsymbol{u}_i$ the output vector representation of word w_i (context embedding). Learning proceeds as follows: For a one hot input vector $\boldsymbol{x} \in \mathbb{R}^{|V|}$ of a center word, the first layer selects $\boldsymbol{v}_c = \boldsymbol{V}\boldsymbol{x}$ and the second layer uses the sigmoid activation function on the score $\boldsymbol{z} = \boldsymbol{U}\boldsymbol{v}_c$. The predicted values are compared to the one hot vectors of the actual output, and the unknown parameters ($\boldsymbol{V}$ and $\boldsymbol{U}$) are updated using Stochastic Gradient Descent. This method is fast, efficient to train, and available in many software packages. We set the context window $m = 5$ and restrict the word embedding length to $n = 100$. The network is trained for five epochs on every monthly data partition.

[9]The word2vec model can be trained using either the skip-gram or the continous bag-of-words (CBOW) algorithm. According to Mikolov et al. (2013) skip-gram works well with small datasets, and can better represent less frequent words. We therefore use this approach here. Other widely used word embedding models include GloVe (Pennington et al., 2014) and fastText Joulin et al. (2016).

2.2 Word embeddings and transition risk

Following Carney (2015), transition risk results from the cost and opportunities following adjustments towards a lower-carbon economy. In the context of commodity currencies, this motivates the two-part decomposition alluded to in the introduction. The first component reflects the opportunity perspective of transitioning, capturing the extent to which a country is perceived as "green" rather than "brown". The second component reflects the cost perspective, and captures the potentially worse outlook for the (influential) carbon-intensive sector of the economy. To provide a quantitative approximation to 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 the "green transition dimension" and the "carbon-economy dimension". This is illustrated in Table 1. The sum of the *concern*, *commodity*, and *economy* categories results in a vector intended to point in a direction encompassing economic risk 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 $climate^{-}$ capture the climate change dimension. Adding $climate^{+} - climate^{-}$ to the "carbon-economy dimension" therefore has the effect of starting at the latter dimension and taking one step in the more climate-friendly, or green, direction. Accordingly, the aggregated transition risk word vector reflects the economic risk dimension in the context of climate change.[10]

To capture the monthly association between countries and the transition risk word vector, we solve

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

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 data block, the subscript t is used to highlight that this relationship potentially changes through time.[11]

[10] Transformer-based models, such as BERT (Devlin et al., 2018), explicitly learn context specific embeddings, and might potentially deliver better representations of the the latent transition risk embedding. However, these models typically contain hundreds of millions parameters, making it unfeasible to estimate on a monthly basis. Transfer learning approaches are possible, but would easily involve some type of forward-looking bias in our setting.

[11] Kozlowski et al. (2019) apply a similar approach to uncover changes in cultural associations. E.g., to determine the gender association for the word "tennis", they project the embedding for this word onto the gender dimension of "man" − "woman", and document how the projection changes through time.

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

$$transition\ risk_t \approx \underbrace{(concern_t + commodity_t + economy_t)}_{\text{carbon-economy dimension}} + \underbrace{(climate_t^+ - climate_t^-)}_{\text{green transition dimension}}$$

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	=	$\frac{1}{n}(russia_t + russian_t)$

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 benefit the climate, we hypothesize that it will put downward pressure on commodity currencies.

We emphasize three points about this construction. First, because of differences in policies, public perception, and consumer and investor behavior across countries, the degree of transition risk is both time-varying and 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. 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 the country, the median correlation is never below 0.94 (Table C.1 in Appendix C). We have also experimented with actively excluding words with climate friendly (e.g., "oxygen" , "ecosystem", and "cooling") and less friendly associations (e.g., “dioxide”, “methane”, and “warming”) from the category lists in Table 1. Doing so we observe that the implied transition risk indexes are almost unaffected. However, removing one or more of the key dimensions from the construction, results in very different indexes. Finally, although the transition risk measures constructed here are motivated by the commodity currency setting, the methodology and intuition are general and potentially valuable in a broader 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 is randomly partitioned into 20 subcorpora, and the word2vec algorithm is run to produce the word embedding matrix for each data partition. Then, the error of the projection statistic TR_t for each subsample s is $e^s = \sqrt{\tau_s}(TR_t^s - TR_t)$, where τ_s and TR_t^s are the number of texts and the solution to (2), 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 bounds of the confidence interval.

Figure 1 reports the country-specific transition risk measures together with the estimated uncertainty. The transition risk measures are precisely estimated and display considerable 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, the risk estimates peak sometime after 2013 for most countries.

2.3 Understanding transition risk embeddings

In studies using (news) text as data to measure the flow of information or latent concepts, annotating estimates like those in Figure 1 with historical events is common to informally validate how plausible they 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. Whereas events likely affect *how much* different topics are discussed in the public discourse, the events might not change *how and in which context* these topics are discussed. 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.

Figures 2 and 3 are better suited to build intuition and evaluate the word embedding approach. Figure 2a reports the historical pairwise cosine similarity for some of the five categories in the transition risk definition in Table 1. Consistent with the conjecture that the flow of economic news has become more climate oriented, the results indicate that “economy” related words are now more associated with “climate” related words than

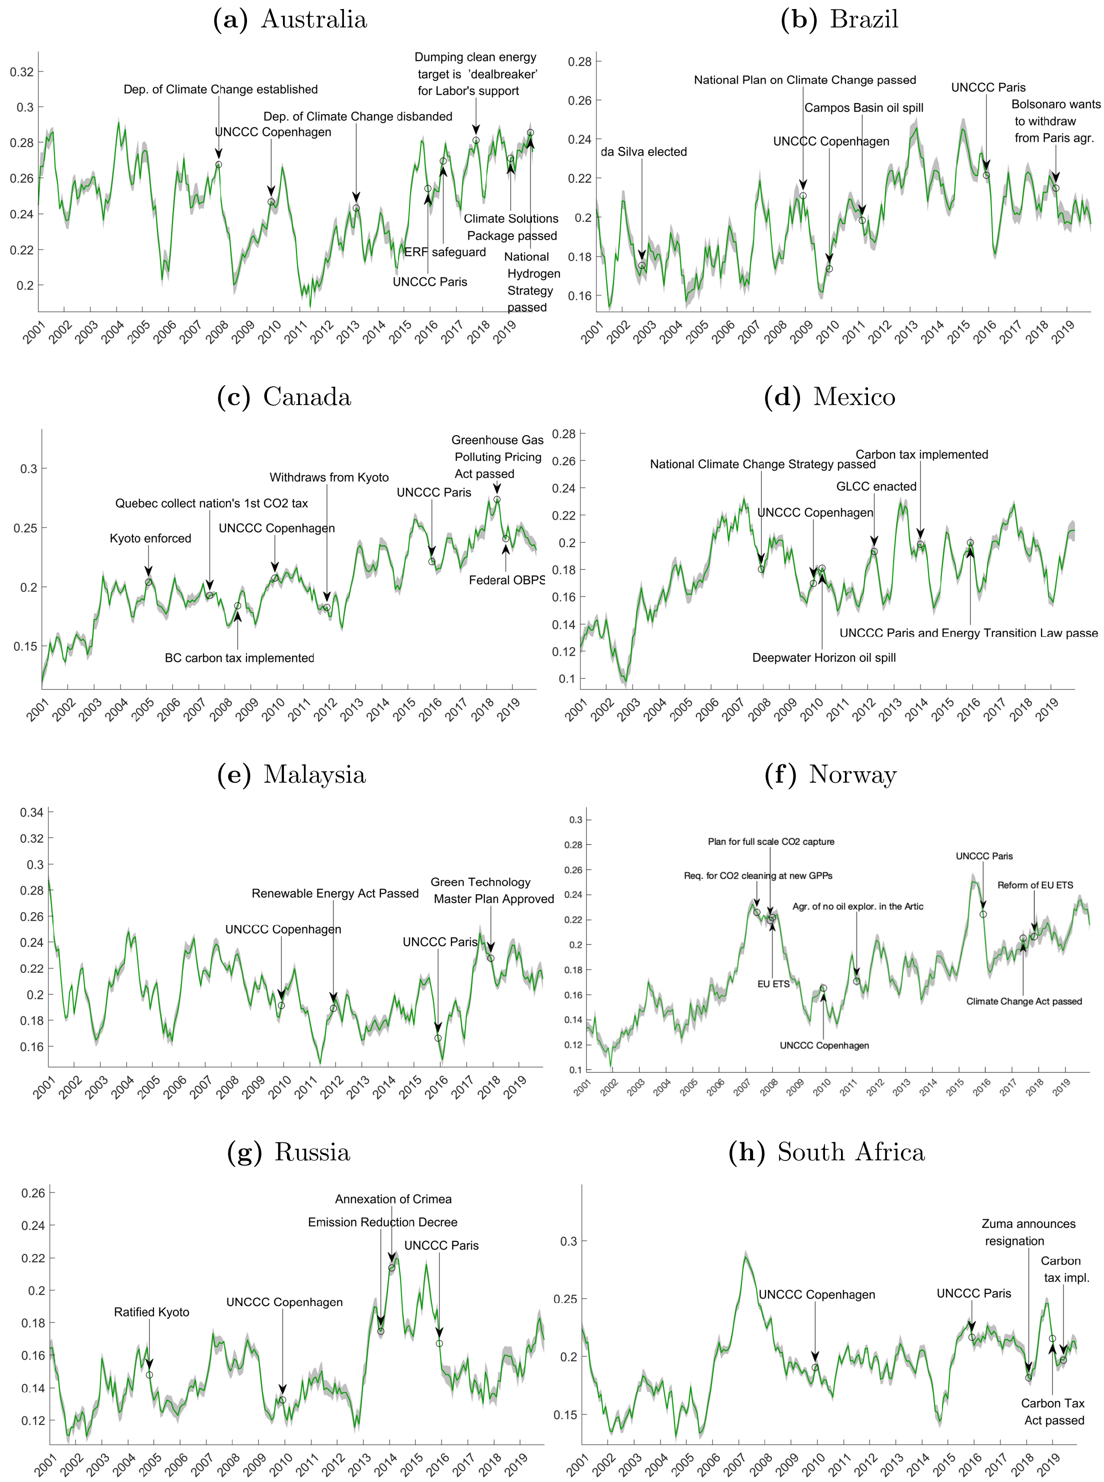

Figure 1. Climate change transition risk. The green lines show the mean estimates. The color shadings cover the 90% confidence intervals. The annotations report some important climate change events. For visual clarity, the raw TR_t series are smoothed using moving averages with a window size of seven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

before. A similar pattern is observed for how these "climate" terms relate to words in the "concern" category. In contrast, it does not seem to be a trending pattern between "commodity" and "climate" related words. Still, the degree of association between these

(a) Cosine similarity **(b)** Word Embedding t-SNE plot, 2019

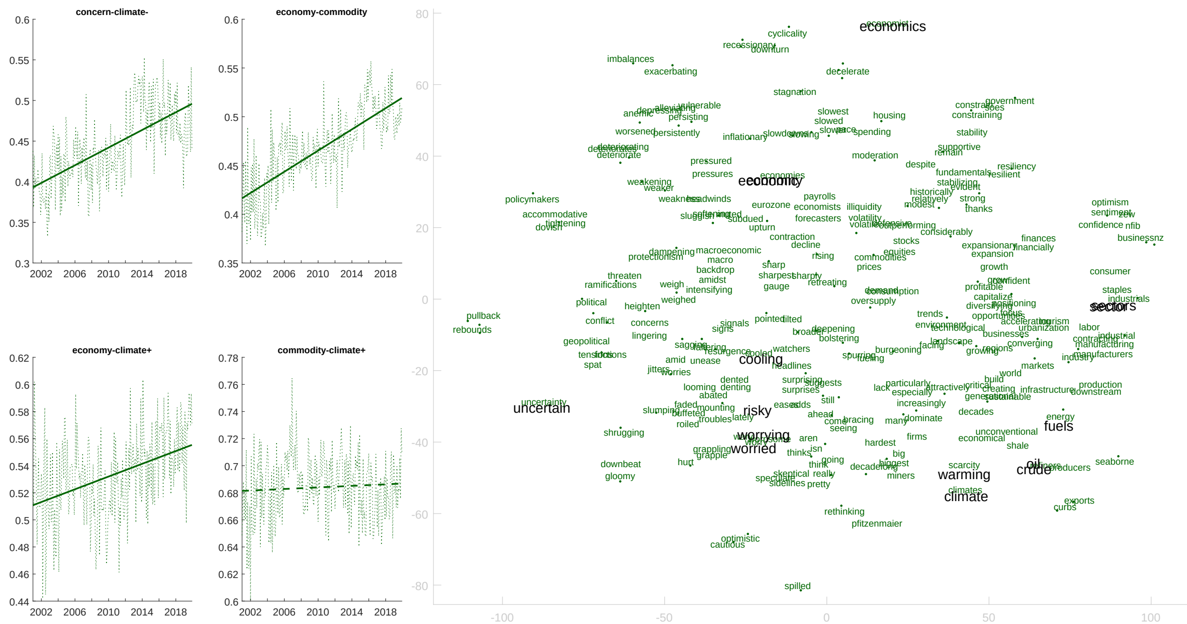

Figure 2. Figure 2a reports the historical pairwise cosine similarity between some of the transition risk categorizes defined in Table 1. The thin broken lines report the raw estimates. The tick lines report the linear trend estimated from $s(w_{it}, w_{jt}) = \alpha + \beta t + e_t$, where $s(w_{it}, w_{jt})$ is the cosine similarity between category i and j at time t. If the estimated trend is not significant at the 5% level, a broken line is reported. Figure C.2, in Appendix C, reports the historical pairwise cosine similarities for all five categories in the transition risk definition in Table 1. Figure 2b reports a t-SNE plot of the word embeddings most closely related to the aggregate transition risk embedding. For each month in the year 2019 the 3000 most similar words to the transition risk vector are extracted (based on cosine similarity scores). Then, we focus on the intersection of this set and compute the average word embedding for each word (across months) in the retained set. The words associated with black color are the words defined in Table 1 that also appear in the retained set. Since independent applications of the word2vec algorithm might result in arbitrary orthogonal transformations, we follow, e.g., Hamilton et al. (2016), and use orthogonal Procrustes to align the word embeddings before averaging. The two-dimensional visualization of the high-dimensional embeddings relies on the t-SNE algorithm (Van der Maaten and Hinton, 2008). We implement this method setting the perplexity to 10, reduce the original dimension of the embedding space to 50 using PCA prior to estimation, and allow for up to 5000 optimization iterations. These choices are common in the literature.

terms is high throughout the sample, indicating that writing about commodities is often done in the context of climate change.

Figure 2b illustrates in two-dimensional space, using the commonly used t-SNE algorithm (Van der Maaten and Hinton, 2008), how the aggregated risk vector encompasses the transition risk concept. It shows which individual words in the corpus are most related to the aggregated transition risk vector. For ease of interpretability, we focus on one particular year, 2019, and the average results across months for this year. While one should be cautious about drawing strong conclusions regarding potential cluster sizes and distances in this type of plot, it delivers at least two relevant conclusions.

First, words that are closely related to the words defined in Table 1 tend to cluster together rather intuitively. For example, the words in the immediate proximity to "crude" and "oil" relate to commodities. Similarly, in the region close to "warming" and "climate", we find words such as "environment", "sustainable", and "generational". For this reason, the plot also illustrates well how the exact words chosen in Table 1 are unlikely to drive our results. For example, whether we had included the term "worrisome" or not in the "concern" set in Table 1 likely matters very little since the "worrisome" embedding is almost identical to the "worried" embedding, which we do include.

Second, transition risk is covered by a relatively wide range of news, and different texts relate to it by using different words. This makes it a high-dimensional object to analyze empirically. As alluded to above, however, words that share contextual information but are seldom used in the same texts, can have similar word embeddings which facilitate arithmetic operations to capture associative meaning, as illustrated in the figure.

To further illustrate this contextual aspect, we construct aggregate embeddings for three other concepts, defined by summing individual words such that "Monetary policy" = "*monetary*" + "*policy*", "Physical climate risk" = "*extreme*" + "*weather*" + "*risk*", and "Policy uncertainty" = "*economic*" + "*policy*" + "*uncertainty*". These concepts range from being loosely to highly related to climate change transition risk.[12] Next, we follow the same procedure as above to find the most related words for each of these embeddings and then compute the unique terms associated with each concept as well as the common terms shared by two or more of them.

Figure 3a reports the low-dimensional t-SNE representation of how the word embeddings in each of these sets relate to each other. The common terms do not form a particular cluster, as expected, while the unique "Monetary policy" and "Policy uncertainty" terms overlap to some extent. "Policy uncertainty" terms do also, to some extent, overlap with transition risk terms, which is intuitive, but both the unique "Physical climate risk" and transition risk terms form distinct clusters. The latter terms are also well-encompassed by the terms used to define the transition risk vector. Roughly 50% percent of the transition risk terms listed in Figure 2b is retained as transition risk unique in Figure 3a.[13] In other words, although the concepts above share many words, they form distinct clusters in Figure 3a because the aggregated embeddings capture associative meaning.

[12] The cosine similarity between these vectors and the transition risk embedding are roughly 0.50, 0.60, and 0.75, respectively. We do not claim that these alternative constructions are the best possible approximations to the different concepts. They are mainly used for illustrative purposes.

[13] The "purity" of the clusters visualized in Figure 3a are not a result of the t-SNE algorithm. Using K-means clustering to estimate four clusters from the combined embedding matrix shows that the unique concept terms in Figure 3a typically are allocated to distinct clusters (Table C.2 in Appendix C).

(a) Clustering (b) Classification performance

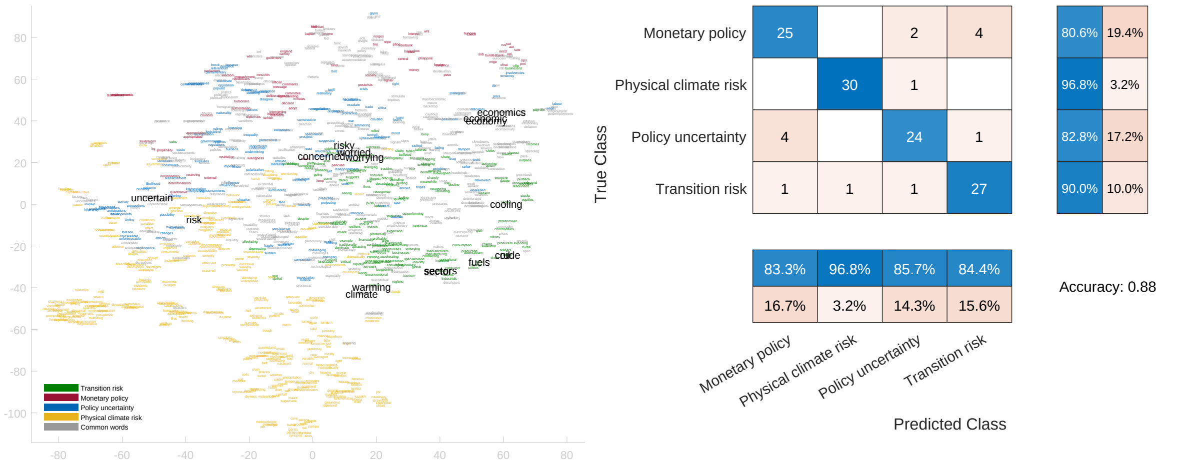

Figure 3. Figure 3a reports a t-SNE plot of the words and concepts related to transition risk. The results are produced following the same procedure as in Figure 2b for the alternative concepts described in the text. The t-SNE algorithm is applied on the embedding matrix containing the unique terms associated with each concept as well as the common terms shared by two or more of them. The colors reflect terms unique to one concept. Common terms are gray. Figure C.2, in Appendix C, zoom in on the different areas in the graph (making words easier to read). Figure 3b reports the confusion matrix chart from the AI-based classification experiment where the estimated word embeddings are used to classify sentences generated by ChatGPT. The center square reports the number of predicted and true class labels. The right column and bottom row reports the recall and precision rate, respectively.

To better formalize that the clusters observed in Figure 3a are distinct and relevant, we perform a small AI-based auditing experiment. We first ask the new ChatGPT model (OpenAI, 2021) to generate roughly 150 sentences about each of the four concepts described above. This data is then used to train a standard multinomial logistic regression using as input how well the words in these sentences explain the concept embeddings. Appendix B describes how we generate the training and validation data. Figure 3b reports a confusion matrix chart summarizing the classification performance on the (out-of-sample) validation data. Taking into account the simplicity of this experiment, the overall 88% accuracy score is reasonable. For the sentences belonging to the transition risk category, we achieve 90% recall and a precision above 80%. I.e., some sentences about the other (related) concepts are wrongly categorized as being about transition risk. In sum, however, the classification results do, to a large extent, echo the information in Figure 3a.[14]

[14] In the next section, when modeling how transition risk affect commodity currencies, we control for both interest rates and other measures of uncertainty in the regressions to alleviate potential concerns about capturing the wrong type of shocks.

3 Transition risk and commodity currencies

Can climate change transition risk explain commodity currency developments? To take into account the dynamic interaction between a set of potentially endogenous variables and identify exogenous disturbances we address this question using a Vector Autoregressive (VAR) modeling framework

$$\boldsymbol{y}_{c,t} = \boldsymbol{A}_{c,1}\boldsymbol{y}_{c,t-1} + \ldots + \boldsymbol{A}_{c,p}\boldsymbol{y}_{c,t-p} + \boldsymbol{D}_c\boldsymbol{x}_t + \boldsymbol{e}_{c,t} \quad \boldsymbol{e}_{c,t} \sim i.i.d.N(0, \boldsymbol{\Sigma}_c), \tag{3}$$

where c and t denote 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 a 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 measures of uncertainty to capture "flight-to-quality" effects in times of (financial) crisis, war, and conflict (Forbes and Warnock (2012), Rey (2015), Goldberg and Krogstrup (2018), Caldara and Iacoviello (2022), Akram (2020)).

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 obtained from OECD's panel of leading indicators and their business tendency survey. For countries where this variable is unavailable, we use the year-on-year growth in industrial production. $r_{c,t}$ is the short-run real interest rate differential, computed using trade weights, while $ComX_{c,t}$ is the real commodity price index obtained from Gruss and Kebhaj (2019).[15] The real effective exchange rates, $REER_{c,t}$, are obtained from BIS. $\boldsymbol{x}_t = \begin{bmatrix} 1 & VIX_t & GPR_t \end{bmatrix}'$, where VIX_t is a measure of financial uncertainty and GPR_t is a measure of geopolitical risk obtained from Caldara and Iacoviello (2022). In the interest of conserving space, a more detailed description of the economic variables is relegated to Appendix A.

While the model in (3) allows for a rich description of the dynamic relationship between the variables, our focus is on how transition risk innovations affect the real exchange rate. For this purpose, we identify exogenous innovations as $\boldsymbol{\varepsilon}_{c,t} = \boldsymbol{P}_c\boldsymbol{e}_{c,t}$, 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

[15] $ComX_{c,t}$ takes into account the basket of commodities produced by country c, and is constructed using time-varying net-export shares. Our main results are robust to using the alternative commodity price indexes derived by Gruss and Kebhaj (2019).

transition risk is contemporaneously unaffected by shocks to the other variables in the system and therefore identify transition risk shocks by ordering climate risk either first or last in the system. These two alternative 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 results are qualitatively unaffected by the particular ordering, suggesting that transition risk is fairly exogenous to the other economic indicators in the 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 have an inflation-targeting monetary policy regime.[16]

3.1 Pooled and partially pooled estimates

We begin by considering two panel VAR versions of (3), 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. 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). Both of these specifications are fairly standard in the literature. A more detailed description of the models is relegated to Appendix D.

Figure 4 reports the response of the REER to a one standard deviation transition risk innovation. Figure 4a shows the pooled responses, while Figures 4b-4i 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 across model 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, with the exception of Malaysia, treating transition risk as either completely contemporaneously

[16] A battery of tests give inconsistent results across countries, regarding both the existence of variable unit roots and the degree of cointegration. An earlier working paper version of this paper documents 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). In unreported results we also confirm that our main qualitative conclusions are robust to using VAR models with REERs in difference form.

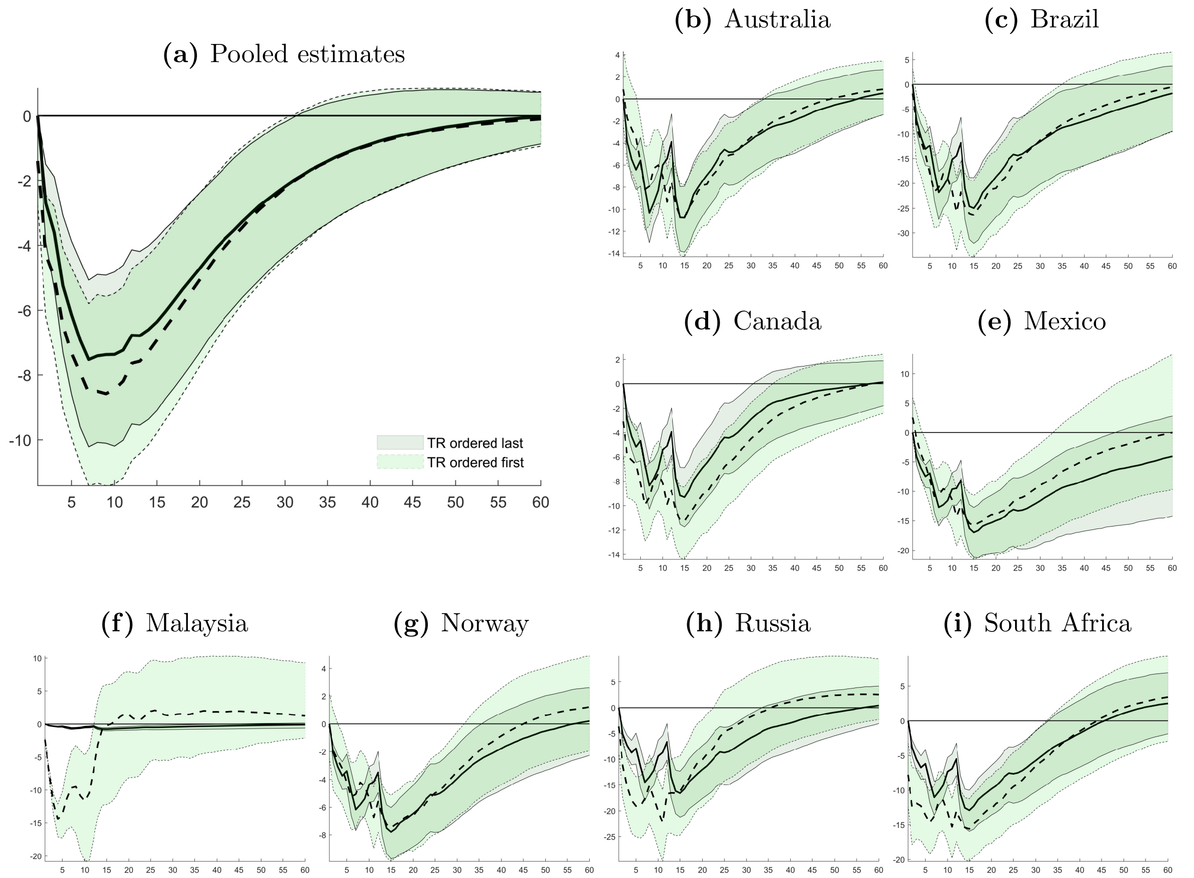

Figure 4. Pooled and partially pooled panel VAR results. Each graph reports the REER response following a one standard deviation exogenous innovation to transition risk. The innovations are computed from two different recursive orderings, where the transition risk variable is ordered either first (dotted black) or last (solid black) in the system. The color shaded areas are 68% probability bands. All data is standardized prior to estimation. The estimates are re-scaled using either the average or country-specific standard deviation of the real exchange rates and reflect percentage change.

exogenous or endogenous to the other variables in the VAR system does not matter qualitatively for these conclusions.

For completeness, the impulse responses associated with the transition risk indexes themselves are reported in Figure C.4 in Appendix C. In short, they indicate fairly transitory response paths. Although we rightfully refrain from making strong structural claims, unexpected innovations to the other variables in the system give REER response paths reasonably in line with conventional economic theory, and these variables respond as we would expect to transition risk innovations (Figures C.5 and C.6 in Appendix C): The business cycle indicators and interest rate differentials both decline significantly, 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) does not affect how transition risk affects real exchange rates (Figures C.7a and C.7b in Appendix C.).

The bar plots in Figure C.9, in Appendix C, report the share of the REER variance explained by transition risk. 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. Overall, these numbers are small, but not negligible. For comparison, a large literature examining the effects of monetary policy shocks does 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 terms of the transition risk indexes, Figure C.8 in Appendix C documents that they are driven mainly by their own shocks and thus largely exogenous to innovations in the other variables in the system.

To probe deeper into the timing of when transition risk historically has affected commodity currencies, Figure C.9 also reports 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. Interestingly, for the period 2012–2016, transition risk is generally interpreted as being lower than expected, suggesting that events associated with, e.g., the Paris Agreement, did not lead to unpredictable short-term increases in transition risk. However, consistent with our assumption that transition risk is country-specific, perhaps the most striking feature of Figure C.9 is its large heterogeneity. For example, depending on the country one chooses to focus on, transition risk shocks put either a strong positive or negative pressure on the REERs in the early part of the sample.

3.2 Dissecting the two-part transition risk construction

Section 2.2 defines transition risk as the association between a given country and the sum of two components reflecting a "green transition dimension" and a "carbon-economy dimension". Considering the two components in isolation, it follows by construction from Table 1 and equation (2) that positive innovations in the "green transition dimension" associates a country with becoming greener while innovations in the "carbon-economy dimension" are associated with potentially worse prospects for the commodity-producing sector of the economy. The combined effect is transition risk because it reflects the costs and opportunities following adjustments towards a lower-carbon economy and considers the economic risk dimension in the context of climate change.

Still, considering the two components in isolation increases our economic understanding of transition risk and its construction. Indeed, the present-value formulation of the real exchange rate suggests that unexpected innovations to the two components of transition risk in isolation might give opposite effects on the exchange rate. To see this more clearly, let a star denote foreign variables, and consider the formulation:

$$q_t = \sum_{h=1}^{\infty} E_t(ir^*_{t+h-1} - ir_{t+h-1}) - \sum_{h=1}^{\infty} E_t(rx_{t+h}) + \gamma_t \tag{4}$$

which writes the real exchange rate $q_t = s_t + p^*_t - p_t$ as a function of expected excess returns ($E_t(rx_{t+h})$), expected macroeconomic fundamentals associated with real interest rate differentials ($E_t(ir^*_{t+h-1} - ir_{t+h-1})$) and the long-run exchange rate ($\gamma_t = lim_{h\to\infty} E_t(q_{t+h})$).[17] Linking (4) to present-value models of stocks, it is common to refer to the interest rate differentials as cash flows and the expected returns as discount rates or the risk premium.

For an investor wanting to hedge against a low-carbon future, and if being more associated with becoming greener is interpreted as being better prepared for a low-carbon future, equation (4) implies that the risk premium falls and the real exchange rate appreciates following unexpected innovations to the "green transition dimension". In contrast, unexpected innovations in the "carbon-economy dimension" are assumed to foremost affect the macro fundamentals of the economy via traditional Dutch disease effects or the Balassa-Samuelson hypothesis. According to (4), this should lead to a depreciation.[18]

To test these predictions, we re-estimate the pooled panel VAR model exchanging the transition risk measure by either of its two components. Both components are constructed from (2). Figure C.10, in Appendix C, reports the series. To allow interest rate differentials and the other variables in the VAR to respond on impact to exogenous innovations in the alternative indexes, we order them first in the system. Figure 5a reports the results.

[17] To derive at this expression, consider a strategy that borrows in home currency and invests in foreign currency. The excess return on this strategy is $rx_{t+1} = s_{t+1} - s_t + i^*_t - i_t$, where s_t is the (log) nominal exchange rate in home currency per unit of foreign currency at date t, and i_t and i^*_t are the home and foreign nominal interest rates between dates t and $t+1$. Expressing the excess return in terms of the real depreciation rate and the real interest rate differential, iterating forward, and taking conditional expectations one arrives at (4). Engel and West (2005), Froot and Ramadorai (2005), Menkhoff et al. (2017) and Dahlquist and Penasse (2022) are prominent and recent examples using the present-value formulation to analyze exchange rate fluctuations.

[18] In theory such innovations can affect both $\sum_{h=1}^{\infty} E_t(ir^*_{t+h-1} - ir_{t+h-1})$ and γ_t. Imposing $\gamma_t = 0$ makes the present-value formulation consistent with PPP. However, it is well documented that real exchange rates can display large and persistent deviations from PPP. Harding et al. (2020) is a recent example focusing on commodity currencies. Thus, we do not make any attempt at separating between the two. Moreover, equation (4) makes no assumptions regarding the processes that drive expectations, making fluctuations in discount rates and cash flows potentially correlated.

(a) Real exchange rate | (b) (Ex-post) Excess returns

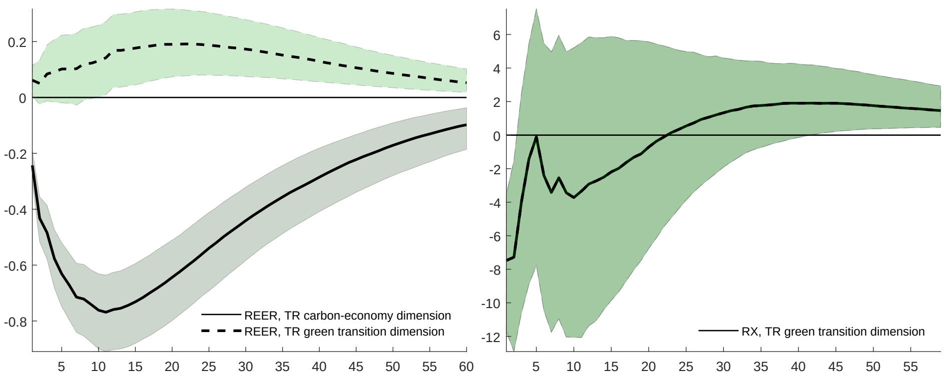

Figure 5. Pooled panel VAR results. Figure 5a reports the REER response following a one standard deviation exogenous innovation to either the "green transition dimension" or the "carbon-economy dimension" of transition risk. Figure 5b reports the ex-post excess returns implied by a a one standard deviation exogenous innovation to the "green transition dimension" component of transition risk.

A one standard deviation innovation in the "green transition dimension" of transition risk leads to a persistent appreciation of the REER, while the opposite holds for unexpected innovations in the "carbon-economy dimension". The sum of the two responses is very close to the aggregate results reported in Figure 4a, suggesting that the dominant effect comes from the costs associated with macroeconomic fundamentals, as expected.

While in line with the simple intuition embedded in equation (4), the results reported in Figure 5a also highlight the duality of the transition risk concept. It is costly for a commodity-producing country to transition to a low-carbon future because the transition has adverse economic effects on the influential commodity-producing sector of the economy (and, thus, the aggregate macroeconomy). At the same time, *transition* defines moving from one state to a potentially different state. Making a green transition therefore potentially makes the same country better prepared for a low-carbon future. Figure 5b provides further evidence for this risk premium argument and reports the ex-post excess returns, i.e., $rx_{t+h} = \Delta q_{t+h} + (ir^*_{t+h-1} - ir_{t+h-1})$, using the estimated response functions of the REER and the real interest rate differential following a shock to the "green transition dimension".[19] Although the long-run responses show some signs of being positive and significant, such shocks lead to significantly lower excess returns in the short run.[20]

[19] Since interest rates are recorded in annual terms we multiply Δq_{t+h} by 12. See, e.g., Bjørnland (2009) and Evans and Rime (2016) for similar VAR-based estimates of ex-post excess returns.

[20] We have also computed the response path of excess returns following unexpected innovations to the "carbon-economy dimension" of transition risk. The theoretical predictions for this component's relationship with discount rates are much weaker than for the "transition dimension" case. In the literature violations of UIP are more common than the alternative, and when UIP does not hold one can find arguments for both positive and negative discount rate effects following innovations to fundamentals. We relegate the results for these innovations to Figure C.12 in Appendix C.

Figure C.11c, in Appendix C, reports how shocks to the "green transition dimension" historically have affected ex-post cumulative excess returns. On average across the sample, these shocks tend to have increased excess returns. Despite substantial variation, the most positive innovations for many countries stem from the latter part of the sample. While this result might seem to contradict the one above, it is actually fully consistent with (4) and efficient markets. It also echoes the findings in recent asset pricing studies, such as Pástor et al. (2022), arguing that the strong performance of green assets in recent years reflects unexpected news related to environmental concerns, and not high expected returns. In other words, the flow of climate change-related news increases the price and thus realized returns, as in 5a, but pushes down expected returns, as in 5b.

3.3 Corroborative results

Because natural resource income is an integral 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. Since transition risk accommodates the future risk of unfavorable shifts in the production function of the commodity-producing sector, commodity supply might also fall in response to positive transition risk innovations.

To address these additional hypotheses, we include 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. Therefore, we restrict the analysis to seasonally adjusted oil production, leaving Australia and South Africa out of the analysis because they produce only small amounts of oil (Figure C.1, in Appendix C). 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 and focus on the partially pooled estimates.

The results are presented in Figure C.13 in Appendix C. Following transition risk innovations, our estimates suggest significant negative stock market responses that also tend to be rather persistent. The results further suggest a temporary reduction in oil production. For most countries, this reduction is significant at the medium response horizons. Interestingly, there are also signs that commodity production increases in the short run, in line with the "green paradox" originally coined by Sinn et al. (2008).[21]

[21] Changes in remaining commodity reserves not due to climate risk might potentially be picked up by our proposed transition risk measures. Reserve 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 C.3 in Appendix C).

We interpret these results as largely consistent with our underlying motivation. The results for the stock market speak to a large literature in finance investigating the implications of increased climate risk for firm value, 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 rather scarce. One exception is Barnett (2019), who finds that global oil supply increases in response to an increased likelihood of significant climate policies being introduced. 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.[22] This allows us to analyze the sensitivity of the pooled (Bayesian) estimates reported earlier and perform two "falsification" experiments.

First, although transition risk is a risk that all countries are exposed to, our theoretical motivation predicts that this type of climate risk should be particularly relevant for countries in which income is highly carbon-dependent. Thus, when analyzing a large number of countries, we expect a significant negative correlation between a country's (carbon) commodity export dependency and the REER response following transition risk shocks.

Figure 6 largely confirms this 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. There is a significant and negative relationship between these two variables. The box plots to the right in the figure further confirm this impression. The real exchange rate responses for non-commodity currencies are not significantly different from zero on average, while they are clearly negative for the commodity currencies. Furthermore, the REER responses for the commodity currencies are, with some exceptions, qualitatively in line with the panel VAR results reported earlier at the one-year horizon, but tend to suggest a more persistent depreciation for at least half of the commodity currencies at the five-year horizon.

[22] Each VAR includes the same endogenous variables as in earlier sections. Because of 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 to relax the computational burden. Transition risk measures for all of the non-commodity currencies are obtained as described in Section 2.2.

(a) Climate change transition risk - 1-year horizon

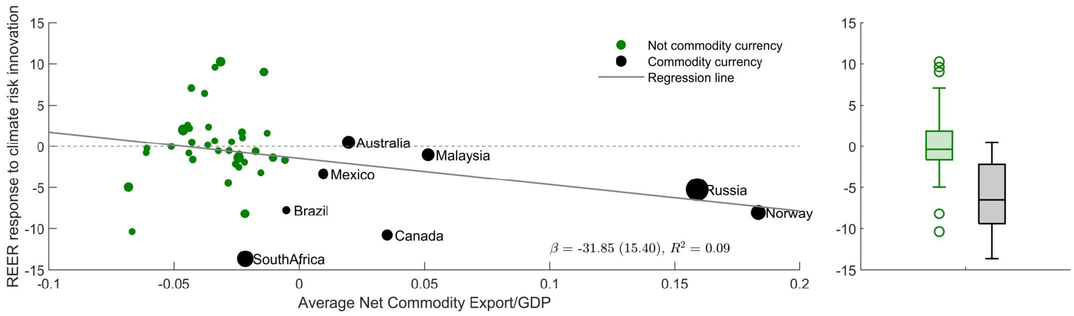

(b) Climate change transition risk - 5-year horizon

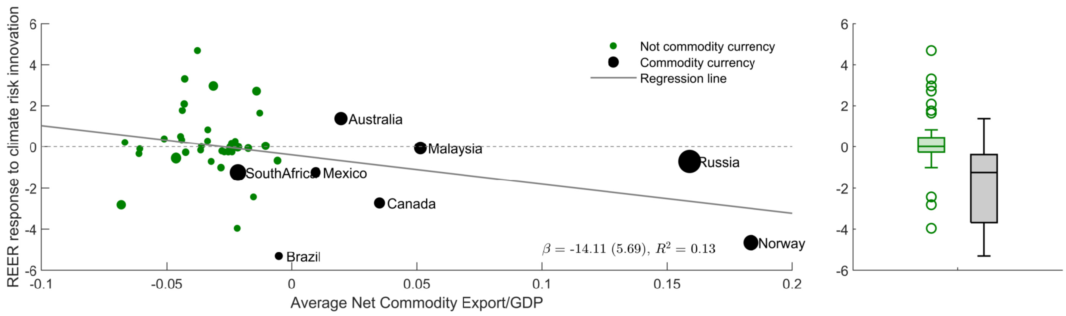

Figure 6. REER responses and commodity export shares. Each graph reports a country's REER response following a one standard deviation exogenous innovation to transition risk (y-axis: in percentage change) together with net commodity exports relative to GDP (x-axis). The REER responses are obtained assuming 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) .

Second, because climate risk is not directly observed, existing approximations used in the literature vary in the degree to which they capture *physical*, *liability*, and *transition* risks associated with climate change (Carney, 2015). In fact, it can be argued that the economic and financial literature has 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 four existing alternative proxies for climate risk and compare the results to those obtained when using our proposed measure.

The four alternative climate risk approximations we consider are the recent news-based climate risk measures suggested by Engle et al. (2020), Gavriilidis (2021), and Ardia et al. (2022), and temperature anomalies.[23] Temperature anomalies are perhaps one of the most direct and widely used measures of climate change (see, e.g., Deschenes and Greenstone

[23] For this purpose we collect statistics from the GISS Surface Temperature Analysis and construct country-specific monthly time series of abnormal temperature fluctuations (see Appendix A). The three news-based indexes were gratefully provided to us by the authors.

(2007) and Kumar et al. (2019)), but are likely a better approximation for *physical climate risk* than *transition risk*. The three news-based indexes build 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, but differ in terms of their construction. The measure suggested by Engle et al. (2020) was developed for hedging climate risk in the asset market, but, as they explicitly state in their article, the index does not distinguish between the 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. Ardia et al. (2022) is inspired by the Engle et al. (2020) approach and construct a climate change concerns score (CCC) measuring and combining the levels of negativity and risk discussed in each news article, while Gavriilidis (2021) develops a climate policy uncertainty index (CPU) building on the method proposed in Baker et al. (2016) for measuring economic policy uncertainty.[24] By focusing on uncertainty, climate risk and regulation, the CPU and CCC measures are conceptually more related to our transition risk index than the Engle et al. (2020) index.

Figure 7a reports the alternative news-based indexes. Since neither of these is country-specific, we plot them together with the common component of our country-specific transition risk measures. Figure C.3 and Table C.4 in Appendix C reports comparable statistics for temperature anomalies. As seen in the figures and tables, the correlation between the alternative climate risk approximations and our suggested indexes is, at times, very high, but overall not perfect.

Figure 7b summarizes the results from the second experiment. Here all the individual VAR models are re-estimated using the alternative climate risk approximations, and the t-statistics from regressing the REER responses on the net commodity export shares are reported across response horizons. I.e., similar to the (scaled) regression slope parameters in Figure 6. See Figures C.14 and C.15, in Appendix C, for directly comparable graphs. Two main findings stand out. First, using the CPU and CCC tends to produce a negative correlation between a country's net commodity export dependency and the REER responses. Still, this relationship is far from as strong as that produced by our proposed transition risk measures. Second, when using temperature anomalies or the Engle et al. (2020) measure, one observes a positive, although not significant, relationship between REER responses and net export shares for all response horizons. To the extent that the measure proposed by Engle et al. (2020) captures the risk premium effect in

[24] Gavriilidis (2021) 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. Ardia et al. (2022) collects news from 10 major U.S. newspapers, selects only articles tagged with a climate change categorization, and then count and aggregate terms related to negativity and risk using two lexicons.

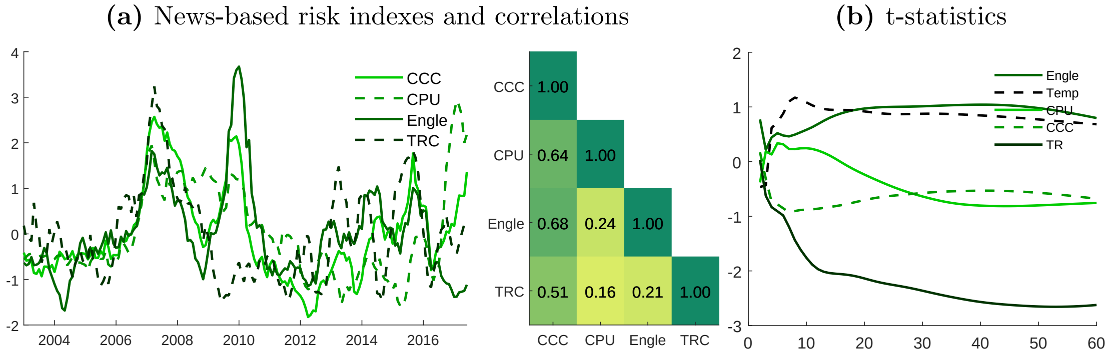

Figure 7. Figure 7a reports the alternative news-based climate risk indexes produced by Engle et al. (2020), Gavriilidis (2021), and Ardia et al. (2022), denoted *Engle*, *CPU* and *CCC*, respectively, together with the common component of the country-specific transition risk indexes (*TRC*). The common component is computed using PCA, and explains roughly 40% of the variation in the data. For visual clarity, the series are smoothed using moving averages with a window size of seven months and standardized. The correlation statistics are computed using the non-smoothed data. Figure 7b reports the t-statistics from regressing the REER responses on the net commodity export shares, where the REER responses are those obtained from the individual VAR models using the different risk approximations. The x-axis denote the response horizons.

the present-value formula, as alluded to by its hedging purpose, this result speaks to the discussion in Section 3.2. However, the result also raises the question of why countries with a high net commodity export share should be more exposed to this effect than less commodity depended economies. One potential explanation for this might be that the flow of climate-related news often is presented in the same context as commodity-relevant news, as indicated by Figure 2a.

In sum, our results suggest that the proposed index adds value because it is designed to explicitly capture the transition risk component of climate change risk and because it captures the duality of the transition risk definition. I.e., transitioning from “brown” to “green” might, in isolation, have a positive effect on the price. However, for a country highly dependent on commodity income, the transition also involves substantial costs and potential loss of revenue, which put downward pressure on the price. The word embedding approach facilitates our construction of country-specific transition risk measures. None of the alternative new-based indexes analyzed here is country-specific. This discrepancy might also play a role in explaining the differences in results.

5 Conclusion

The economic risks posed by climate change have at least three components; *physical risk*, *liability risk*, and *transition risk*. Perhaps due to ease of measurement, the economic literature has focused foremost on the former of these components. We propose a measure of climate change transition risk based on neural network word embedding models for large-scale text analysis.

Textual news data is used as input to the model because new information affects expectations about the future and, thus, economic agent's actions today. The word embedding framework is proposed because word embeddings densely encode linguistic regularities and patterns and allow for arithmetic operations capturing associative meaning. That is, although transition risk is country-specific and can change in response to various news events, arithmetic operations on a few relevant terms yield aggregate embeddings pointing in a direction associated with transition risk in vector space.

While the intuition for our proposed methodology is general and can be applied in many different settings, the application is on commodity currencies. The present-value formulation of exchange rates structures how we construct the transition risk measures in the current context. It also highlights why this type of climate risk is particularly relevant for countries where income is highly dependent on high-carbon exports.

In line with theory, we document that when transition risk increases, commodity currencies experience a persistent depreciation while non-commodity currencies are unaffected on average. According to our estimates, roughly 4% of the long-run fluctuations in the real exchange rate can be explained by unexpected transition risk innovations. When we use existing climate risk proxies not designed to explicitly capture the transition risk component of climate change, these findings do not apply. Hence, the conclusions drawn about the economic consequences of climate change are highly dependent on the considered climate risk component.

Apart from speaking to a growing literature using tools from NLP and ML to improve and facilitate measurement in economics, our study speaks to a large literature concerned with the pricing of firms and firm value. By focusing on exchange rates, we document how transition risk also affects valuations at a national level. This obviously has implications for policymaking, and monetary policy in particular.

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.

Angelico, C., J. Marcucci, M. Miccoli, and F. Quarta (2022). Can we measure inflation expectations using twitter? *Journal of Econometrics 228*(2), 259–277.

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2022). Climate change concerns and the performance of green vs. brown stocks. *Management Science, forthcoming*.

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.

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.

Balassa, B. (1964). The purchasing-power parity doctrine: A reappraisal. *Journal of Political Economy 72*(6), 584–596.

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. (2009). Monetary policy and exchange rate overshooting: Dornbusch was right after all. *Journal of International Economics 79*(1), 64–77.

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 (2021). Do investors care about carbon risk? *Journal of Financial Economics 142*(2), 517–549.

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 (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225.

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.

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.

Dahlquist, M. and J. Penasse (2022). The missing risk premium in exchange rates. *Journal of Financial Economics 143*(2), 697–715.

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.

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

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.

Engel, C. and K. West (2005). Exchange rates and fundamentals. *Journal of Political Economy 113*(3), 485–517.

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.

Evans, M. D. and D. Rime (2016). Order flow information and spot rate dynamics. *Journal of International Money and Finance 69*, 45–68.

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.

Froot, K. A. and T. Ramadorai (2005). Currency returns, intrinsic value, and institutional-investor flows. *The Journal of Finance 60*(3), 1535–1566.

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

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.

Hamilton, W. L., J. Leskovec, and D. Jurafsky (2016). Diachronic word embeddings reveal statistical laws of semantic change. *arXiv preprint arXiv:1605.09096*.

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.

Harding, T., R. Stefanski, and G. Toews (2020, 02). Boom Goes the Price: Giant Resource Discoveries and Real Exchange Rate Appreciation. *The Economic Journal 130*(630), 1715–1728.

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 (2022). The pollution premium. *Journal of Finance, forthcoming*.

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.

Joulin, A., E. Grave, P. Bojanowski, and T. Mikolov (2016). Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*.

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.

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 (2021). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics 117*, 507–520.

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.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2017). Currency value. *The Review of Financial Studies 30*(2), 416–441.

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.

OpenAI (2021). Chatgpt.

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

Pástor, L., R. F. Stambaugh, and L. A. Taylor (2022). Dissecting green returns. *Journal of Financial Economics 146*(2), 403–424.

Pennington, J., R. Socher, and C. Manning (2014, October). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, pp. 1532–1543. Association for Computational Linguistics.

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.

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.

Samuelson, P. A. (1964). Theoretical notes on trade problems. *The Review of Economics and Statistics 46*(2), 145–154.

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.

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 Maaten, L. and G. Hinton (2008). Visualizing data using t-sne. *Journal of machine learning research 9*(11).

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 (2022), 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. The longitude and latitude resolution provided in the database are used to construct country-specific monthly time series of abnormal temperature fluctuations.

Appendix B AI-audit

To better formalize that the clusters observed in Figure 3a are distinct and relevant, we design a simple classification experiment. As an alternative to manually and subjectively classifying sentences in our corpus, we exploit the potential in the newly developed ChatGPT technology (OpenAI, 2021) to construct training and testing data for this experiment. We thus asked ChatGPT to answer the questions listed in Table B.1. To not make the sampling of data strictly dependent on how we phrase the instructions, four different, although similar, types of questions were asked for each concept.

All output from the machine was then collected as individual sentences and categorized according to the setting in which ChatGPT generated them. Some machine calls failed, and the number of sentences generated was not always as asked for. In total we obtained 605 sentences, where roughly 25% of them belong to each of the four categories.

In the next step we used a standard multinomial logistic regression to classify the text based on the four word embeddings constructed from our dataset. To map concept embeddings to sentences and construct predictors we regress each concept embedding on the word vectors for the words in a given sentence. If a sentence word is not in our corpus it is dropped. The R-squared from these regressions are then used as explanatory variables in the classification task.

The multinomial logistic regression is trained on 80% of the data. The remaining data is used for out-of-sample predictive validation.

Table B.1. Instructions given to ChatGPT and examples of output.

Input	
For transition risk:	Can you generate 80 sentences about climate change transition risk in an oil exporting country context and without necessarily using the phrase "climate change transition risk"
	Can you generate 20 sentences about climate change transition risk in an oil exporting country context
	Can you write 20 sentences about to cost and opportunities associated with a transition to a low-carbon economy
	Can you write a short news story about how climate change transition risk might affect the economy going forward
For monetary policy:	Can you generate 80 sentences about monetary policy without necessarily using the phrase "monetary policy"
	Can you generate 20 sentences about monetary policy
	Can you generate 20 sentences about the current monetary policy situation
	Can you write a short news story about how monetary policy might affect the economy going forward
For policy uncertainty:	Can you generate 50 sentences about economic policy uncertainty without necessarily using the phrase "economic policy uncertainty"
	Can you generate 20 sentences about economic policy uncertainty
	Can you generate 20 sentences reflecting various aspects of economic policy uncertainty
	Can you write a short news story about how economic policy uncertainty might affect the economy going forward
For physical climate risk:	Can you generate 50 sentences about physical climate change risk without necessarily using the phrase "physical climate change risk"
	Can you generate 20 sentences about physical climate change risk
	Can you generate 20 sentences about how extreme weather might affect the economy
	Can you write a short news story about how physical climate risk might affect the economy going forward
Examples of output	
Transition risk	The oil-dependent economy is facing decreased demand for fossil fuels.
Transition risk	The shift towards renewable energy is affecting the country's main source of revenue.
Transition risk	Oil reserves are becoming less valuable as the world moves away from fossil fuels.
Monetary policy	The central bank influences the supply of money in the economy.
Monetary policy	Interest rates play a key role in shaping economic activity.
Monetary policy	The central bank aims to maintain price stability and full employment.
Economic policy uncertainty	The government's fiscal and monetary policies play a crucial role in shaping the economy.
Economic policy uncertainty	Uncertainty about future economic conditions can lead to a lack of investment and economic slowdown.
Economic policy uncertainty	The government's actions can impact consumer and business confidence.
Physical climate risk	Rising temperatures are causing more frequent and severe weather events.
Physical climate risk	The melting of polar ice caps is leading to rising sea levels.
Physical climate risk	Extreme weather conditions are damaging infrastructure and causing power outages.

Appendix C Additional results

Table C.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 C.2. K-means clustering. The K-means algorithm is used to estimate four clusters based on the combined embedding matrix visualized in Figure 3a. The table reports how the fraction of unique terms for the four concepts listed in the first row are allocated across the four estimated clusters.

Cluster	"Transition risk"	"Monetary policy"	"Policy uncertainty"	"Physical climate risk"
1	0.03	0.93	0.26	0.00
2	0.00	0.01	0.00	0.46
3	0.04	0.04	0.32	0.50
4	0.94	0.02	0.42	0.04

Table C.3. 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 C.4. 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 C.3, in Appendix C, 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***

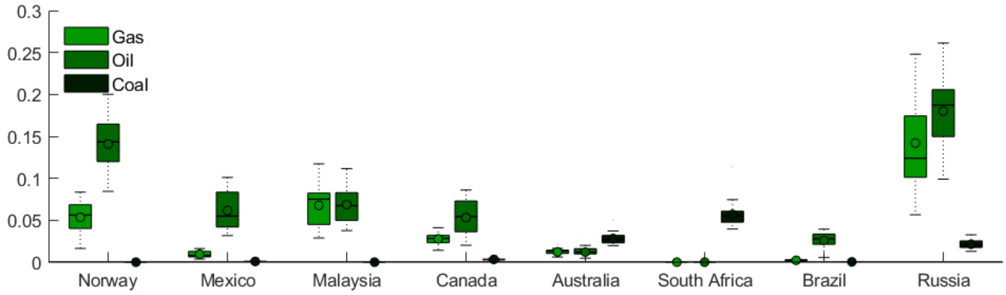

Figure C.1. Gas, oil, and coal production relative to GDP. For each country, the figure reports a standard box plot of the production shares for the period 2002 to 2019. The underlying data is sourced from British Petroleum Company (2020).

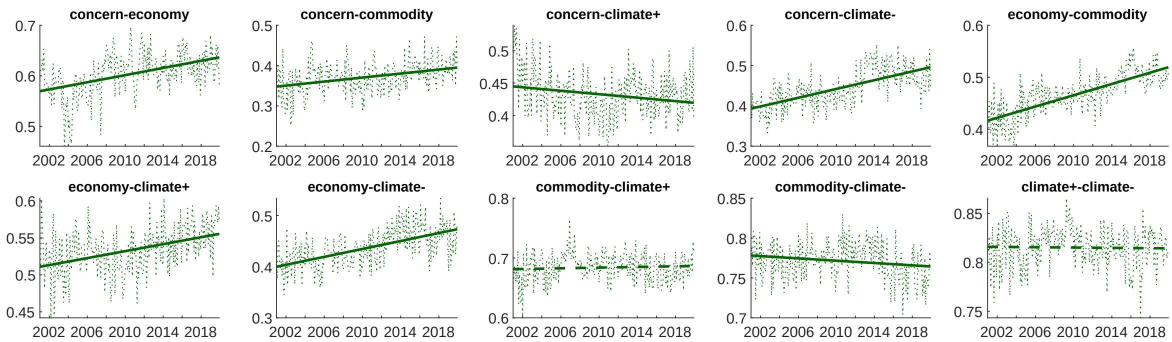

Figure C.2. The figure reports the historical pairwise cosine similarity between all of the transition risk categorizes defined in Table 1. The thin broken lines report the raw monthly estimates. The tick lines report the linear trend estimated from $s(w_{it}, w_{jt}) = \alpha + \beta t + e_t$, where $s(w_{it}, w_{jt})$ is the cosine similarity between category i and j at time t. If the estimated trend is not significant at the 5% level, a broken line is reported.

(a) Middle zoom

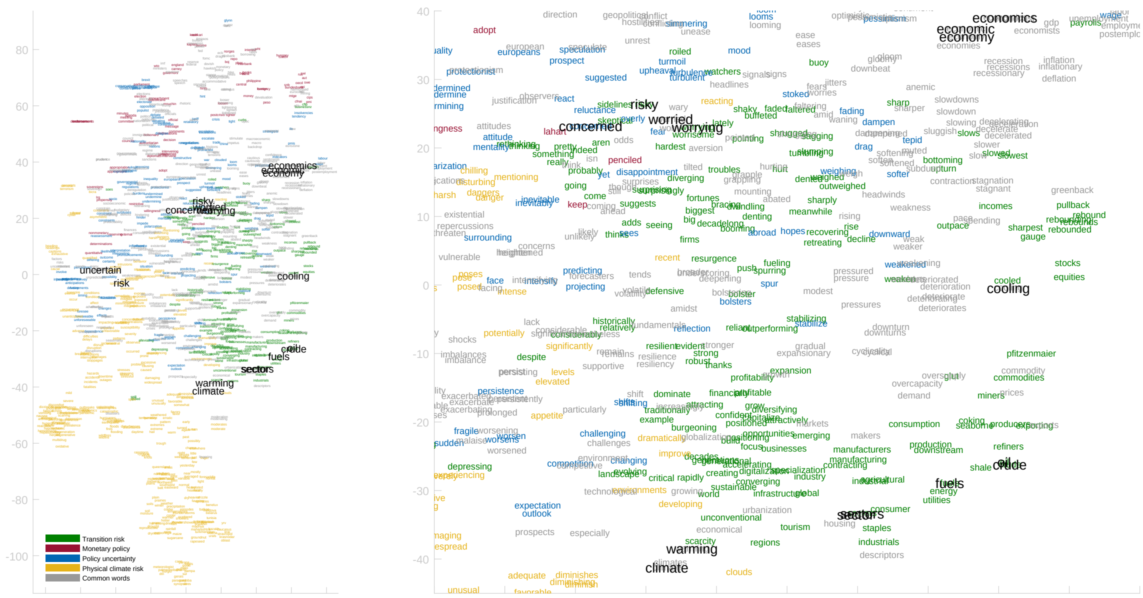

(b) Upper zoom

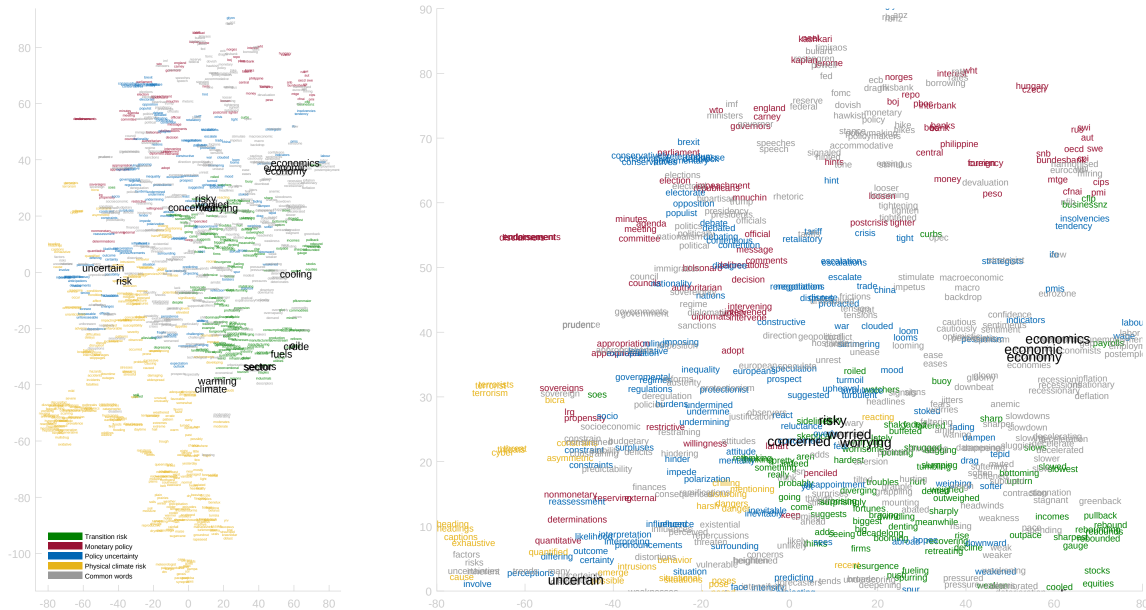

(c) Lower zoom

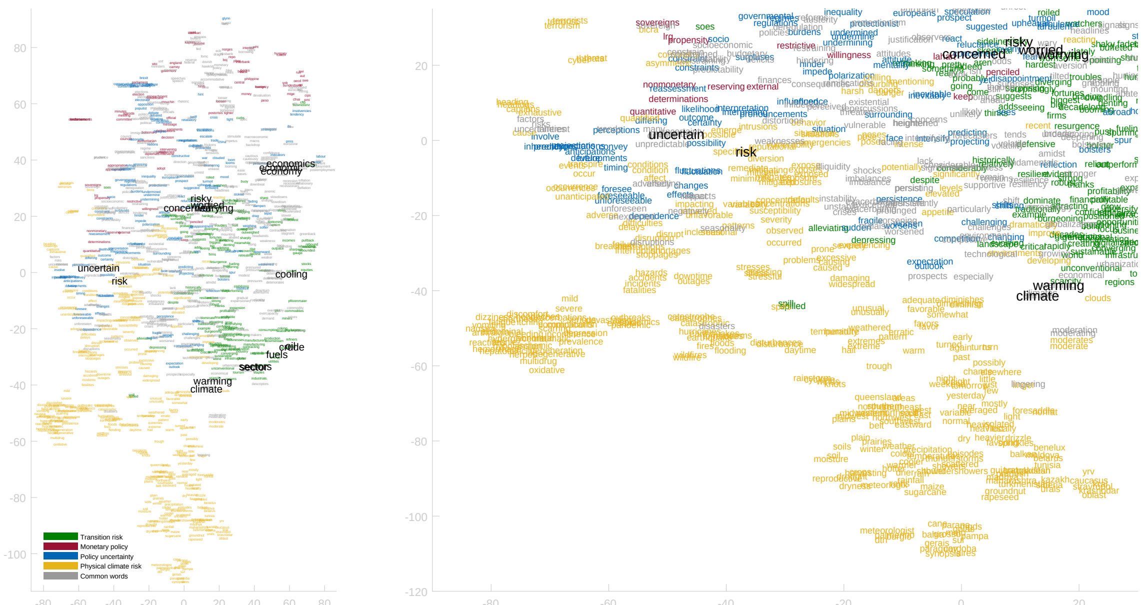

Figure C.2. Word Embedding t-SNE plots, 2019, of transition risk related concepts. The results are produced following the same procedure as in Figure 2b for the alternative concepts "Monetary policy", "Policy uncertainty", and "Physical climate risk". The t-SNE algorithm is applied on the embedding matrix containing the unique terms associated with each concept as well as the common terms shared by two or more of them. The colors reflect terms that are unique to one concept. Common terms are gray. The graphs to the left reports all terms and concepts. The graphs to the right zoom in on the middle, upper and lower areas of the left graphs, respectively.

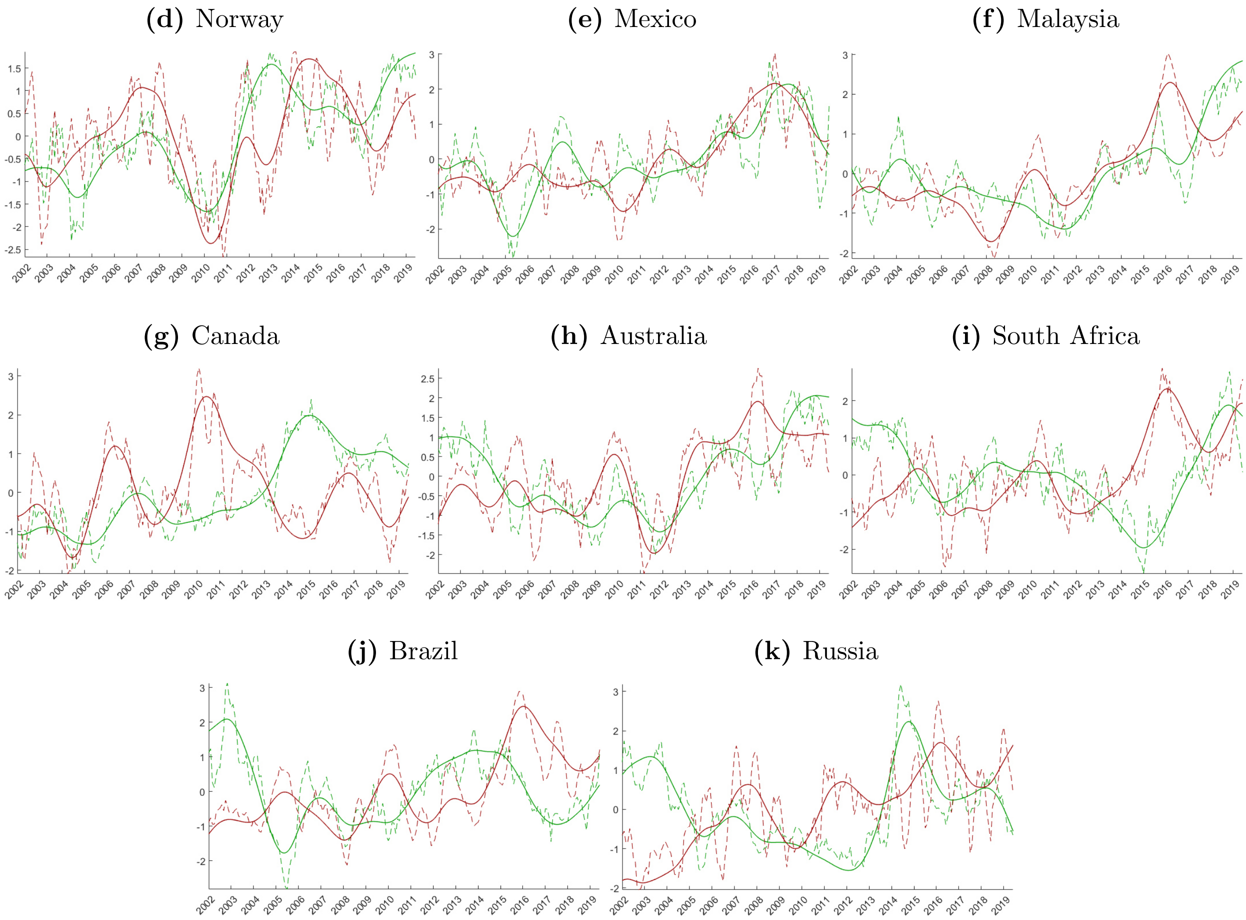

Figure C.3. Transition risk (green) and temperature anomalies (red). The dotted lines report the raw series. The solid lines report the data when a Hodrick–Prescott filter (Hodrick and Prescott (1997)), with a smoothing parameter set to 1600, is used to extract the low-frequency fluctuations from the series.

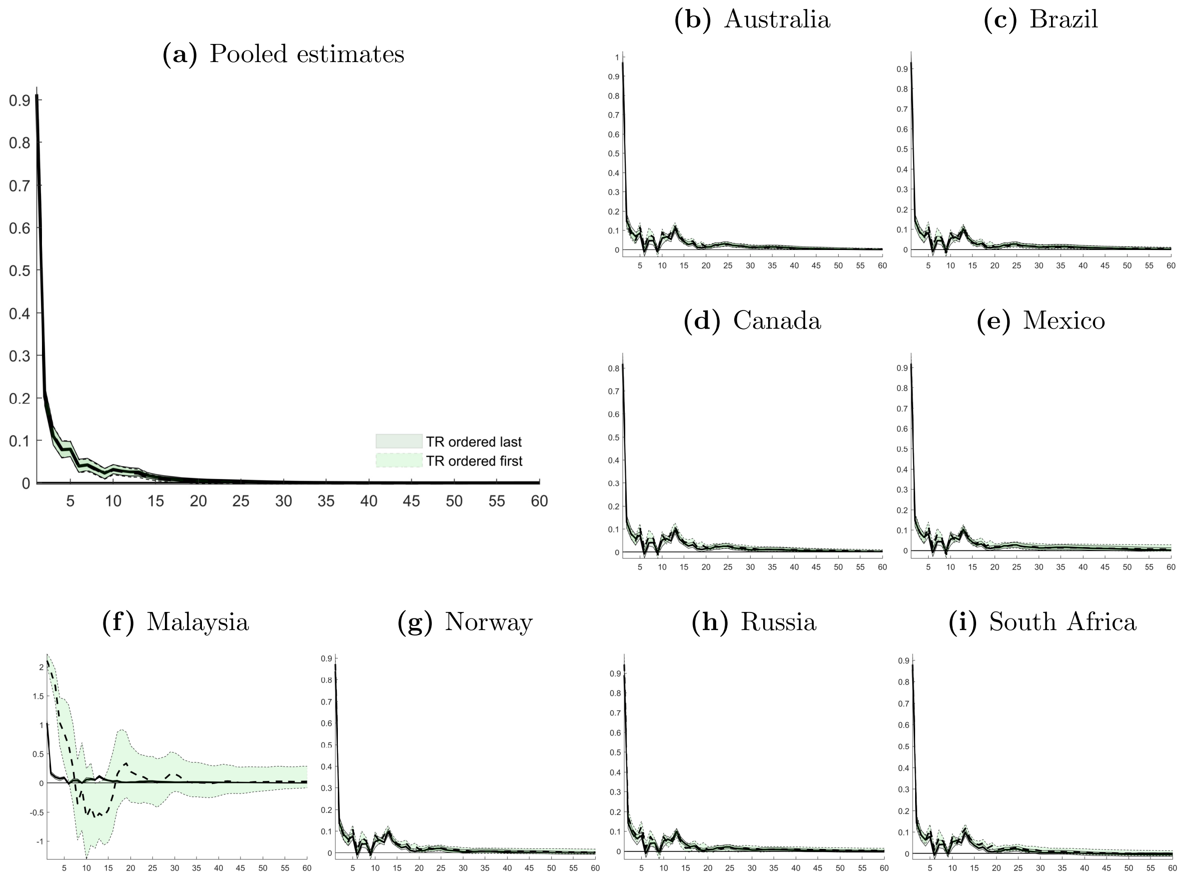

Figure C.4. Pooled and partially pooled panel VAR results. Each graph reports the transition risk response following a one standard deviation exogenous innovation to the transition risk variable. The innovations are computed from two different recursive orderings, where the transition risk variable is ordered either first (dotted black) or last (solid black) in the system. The color shaded areas are 68% probability bands.

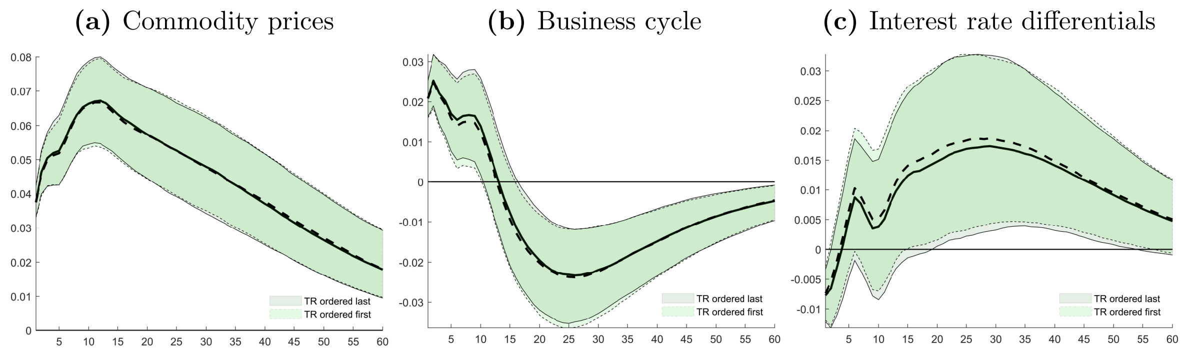

Figure C.5. Pooled VAR and REER responses. Each graph reports the REER response following a one standard deviation exogenous innovation to commodity prices, the business cycle index, or interest rate differentials. The innovations are computed from two different recursive orderings, where the transition 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

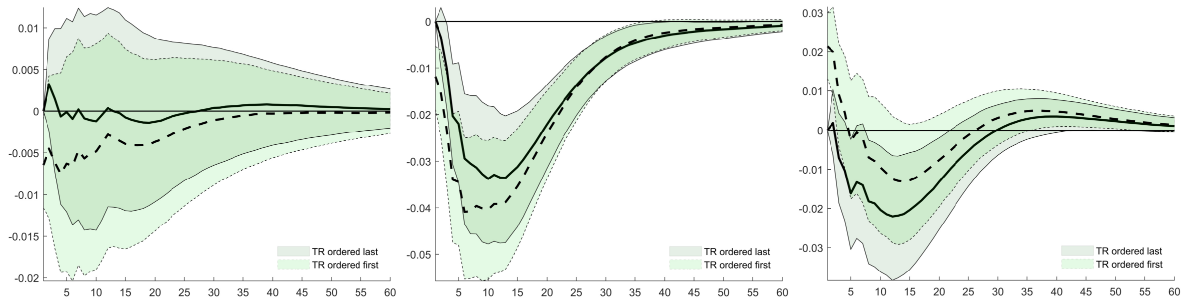

Figure C.6. Pooled VAR and macroeconomic responses. The graphs report the responses of commodity prices, the business cycle index, and interest rate differentials, following a one standard deviation exogenous innovation to transition risk. The innovations are computed from two different recursive orderings, 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

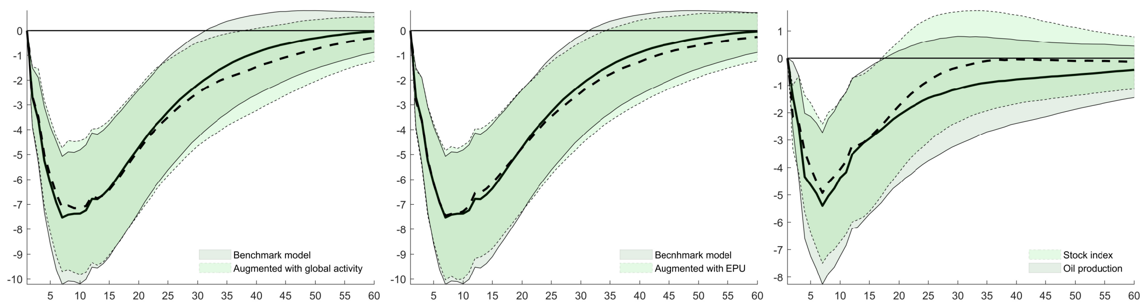

Figure C.7. Pooled VAR and REER responses. Each graph reports the real exchange rate response following a one standard deviation exogenous innovation to the transition risk variable. The transition risk variable is ordered last in the system. Each VAR is augmented with the global activity measure proposed 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.

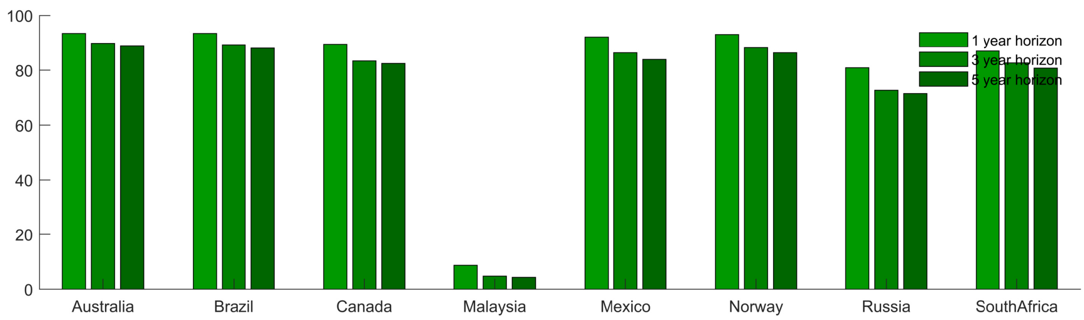

Figure C.8. Partially pooled panel VAR results. Transition risk and the variance explained by transition risk innovations. Each bar reports the median estimate for a given horizon and country. Estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system.

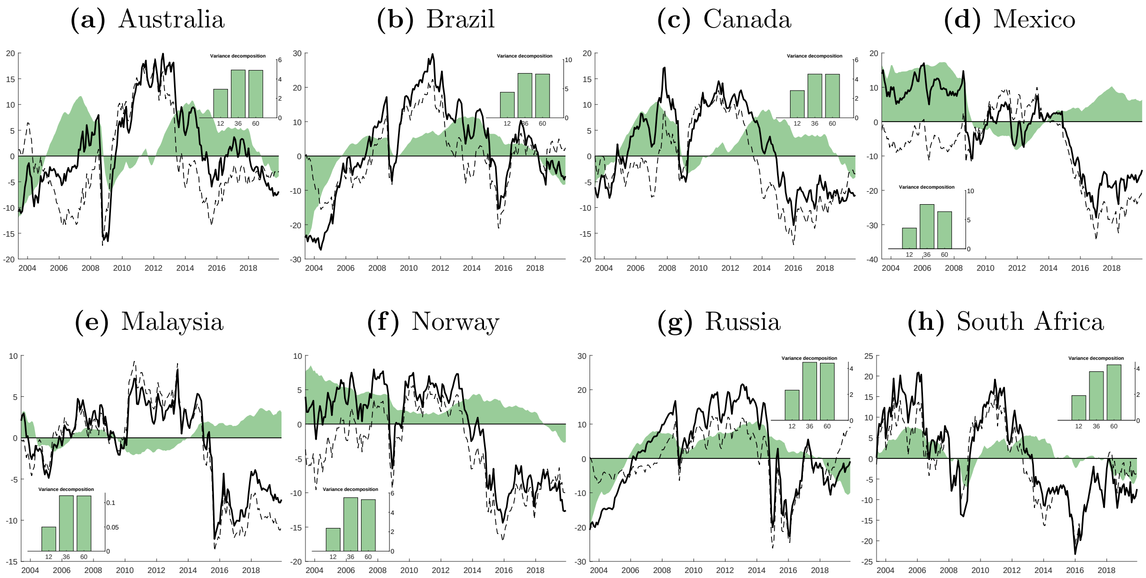

Figure C.9. Partially pooled panel VAR results. The lines report the actual REER (solid line), the counterfactual REER without transition risk innovations (broken line), and the difference between these two lines (green area). The bar plots report the (median) variance decomposition at three different horizons. Estimates are obtained ordering the transition risk variable last in the VAR system.

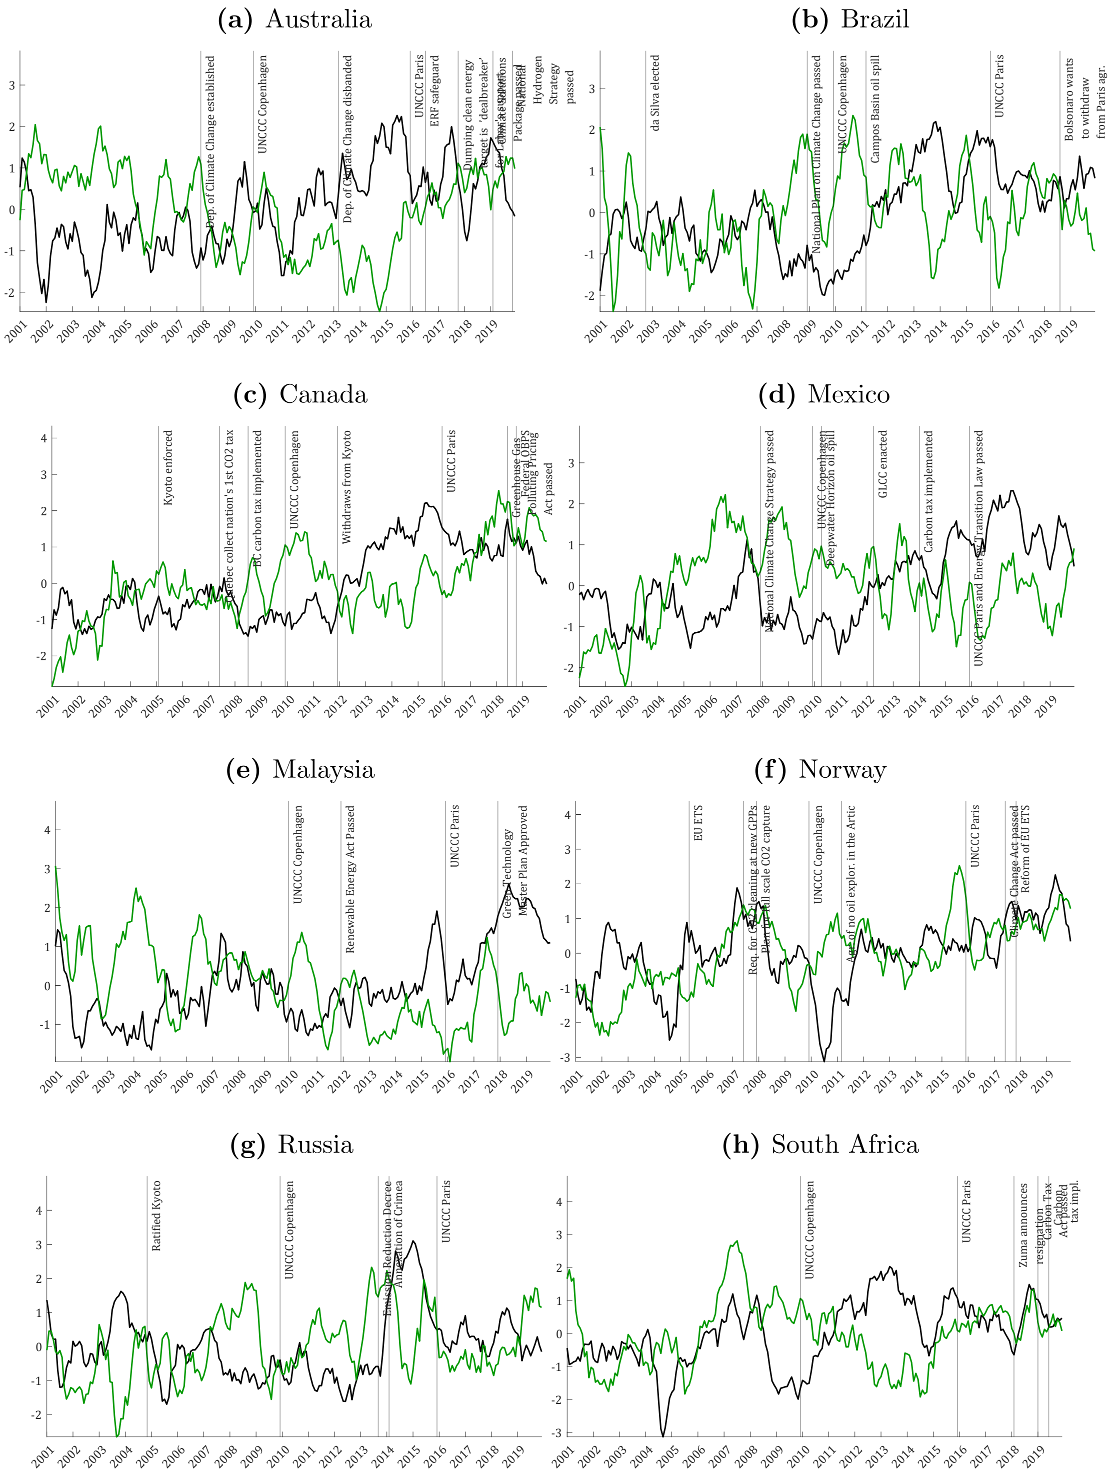

Figure C.10. Climate change transition risk components. The green line report the "green transition dimension" of transition risk and the black line reports the "carbon-economy dimension". The annotations report some important international and domestic political climate change events. For visual clarity, the series are smoothed using moving averages with a window size of seven months and standardized.

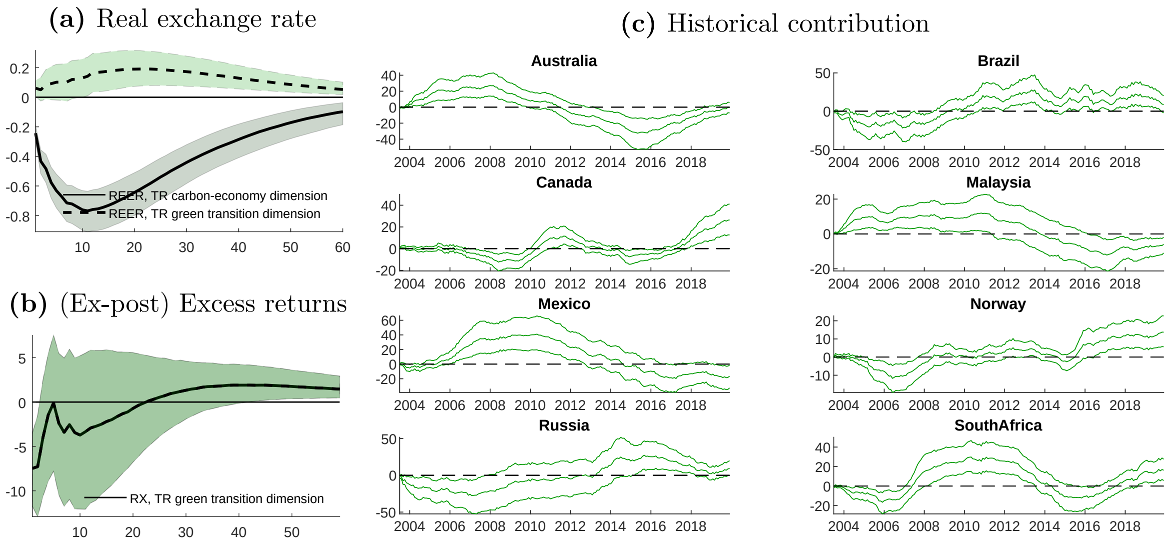

Figure C.11. Pooled panel VAR results. Figure C.11a reports the REER response following a one standard deviation exogenous innovation to either the "green transition dimension" or the "carbon-economy dimension" of transition risk. Figure C.11b reports the ex-post excess returns implied by a a one standard deviation exogenous innovation to the "green transition dimension" component of transition risk. Figure C.11c reports how shocks to the "green transition dimension" of transition risk historically have affected ex-post cumulative excess returns. Because the estimated time series of structural shocks are different for each country in the panel, the historical shock decomposition are country specific. The color shaded areas are 68% probability bands. Estimates are obtained assuming a recursive ordering with the transition risk components ordered first in the VAR system.

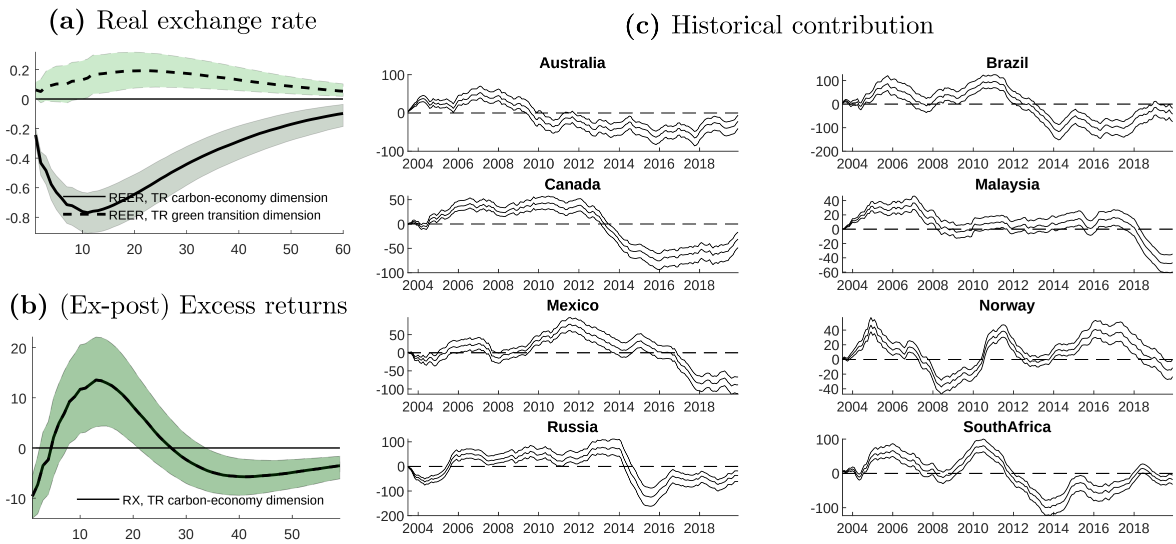

Figure C.12. Pooled panel VAR results. Figure C.12a reports the REER response following a one standard deviation exogenous innovation to either the "green transition dimension" or the "carbon-economy dimension" of transition risk. Figure C.12b reports the ex-post excess returns implied by a a one standard deviation exogenous innovation to the "carbon-economy dimension" component of transition risk. Figure C.12c reports how shocks to the "carbon-economy dimension" of transition risk historically have affected ex-post cumulative excess returns. Because the estimated time series of structural shocks are different for each country in the panel, the historical shock decomposition are country specific. The color shaded areas are 68% probability bands. Estimates are obtained assuming a recursive ordering with the transition risk components ordered first in the VAR system.

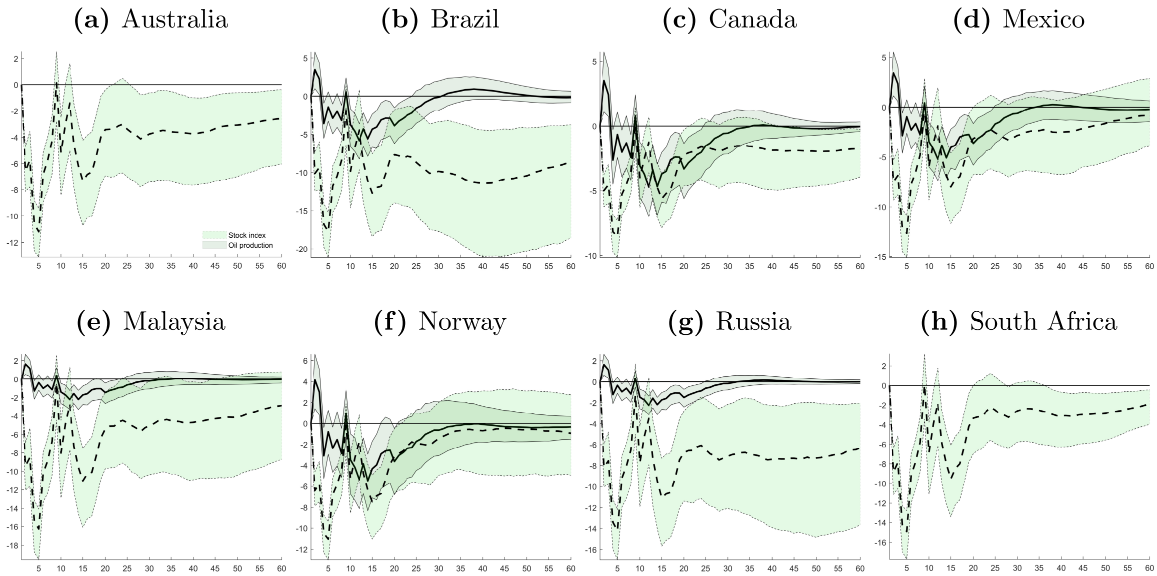

Figure C.13. Partially pooled panel VAR results. Each graph reports the oil production and stock market responses following a one standard deviation exogenous innovation to transition risk. Estimates are obtained assuming a recursive ordering with transition risk ordered last in the VAR system. The color shaded areas are 68% probability bands. All data is standardized prior to estimation. The 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.

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

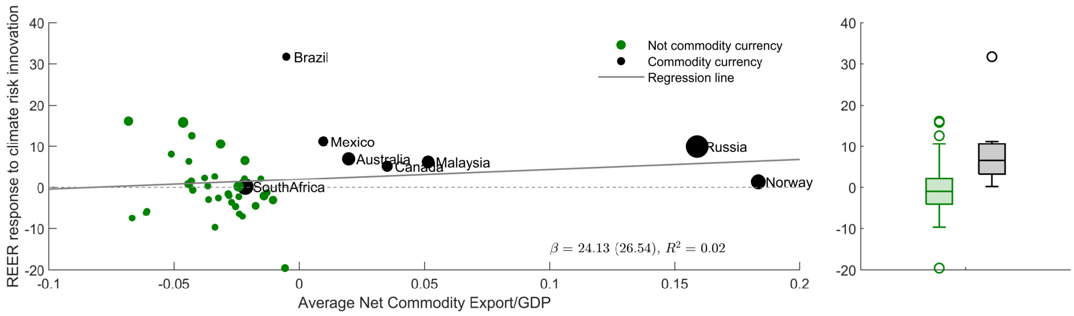

(b) Gavriilidis - 1-year horizon

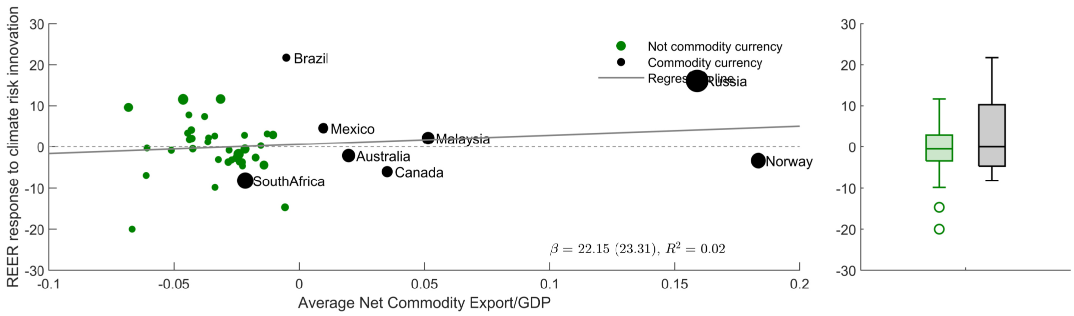

(c) Ardia et. al.: Climate concern - 1-year horizon

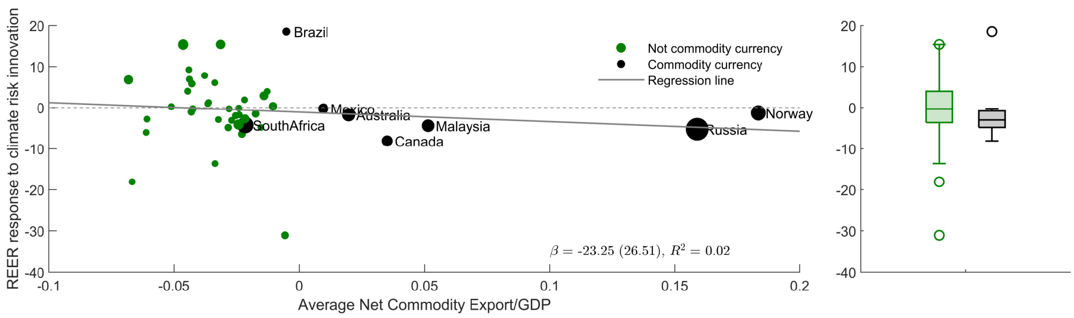

(d) Temperature anomalies - 1-year horizon

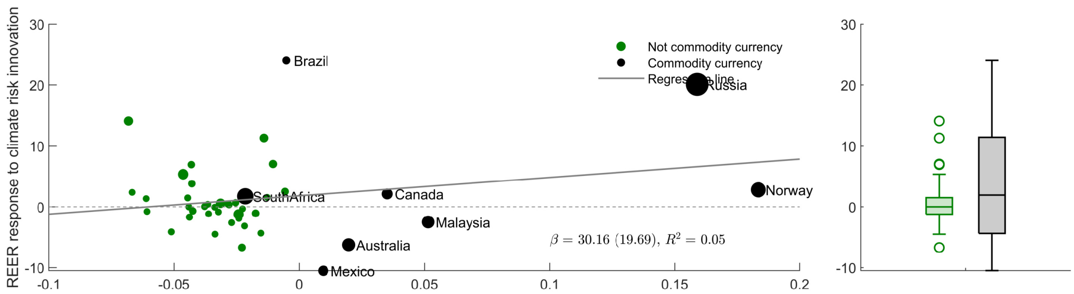

Figure C.14. REER responses for alternative climate risk proxies and commodity export shares. Each graph reports a country's REER response following a one standard deviation exogenous innovation to the Engle et al. (2020), Gavriilidis (2021), and Ardia et al. (2022) climate risk indexes, or temperature anomalies (y-axis: in percentage change). See Figure 6 for additional details.

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

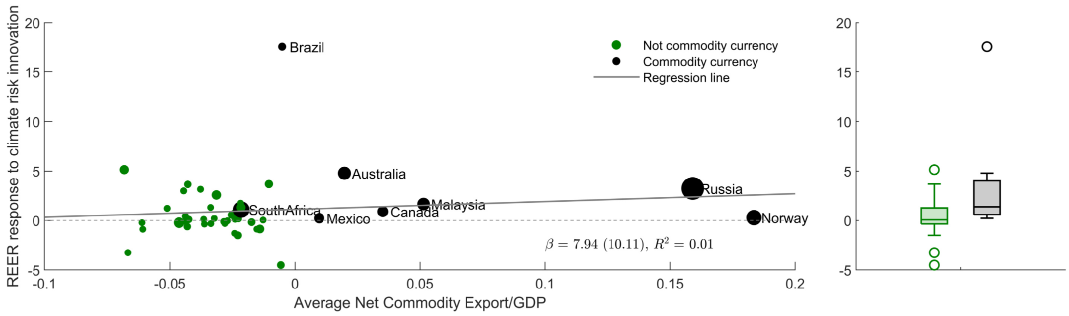

(b) Gavriilidis: Climate policy uncertainty - 5-year horizon

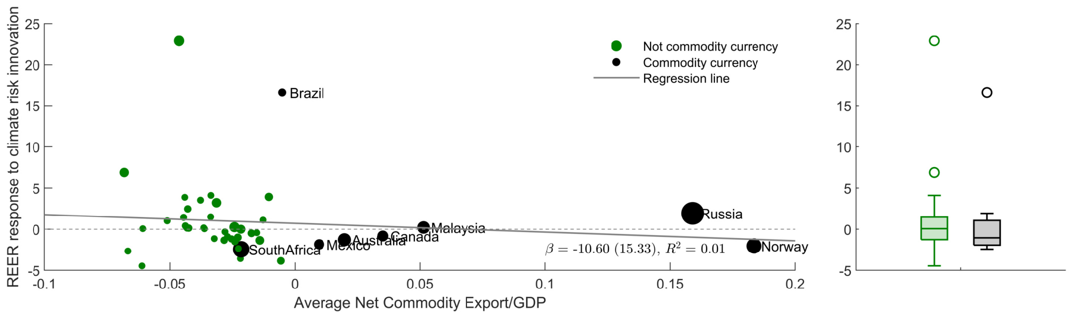

(c) Ardia et. al.: Climate concern - 5-year horizon

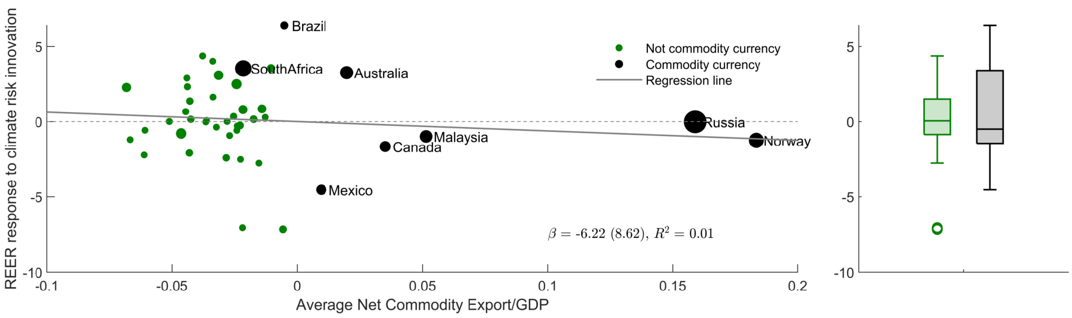

(d) GISS: Temperature anomalies - 5-year horizon

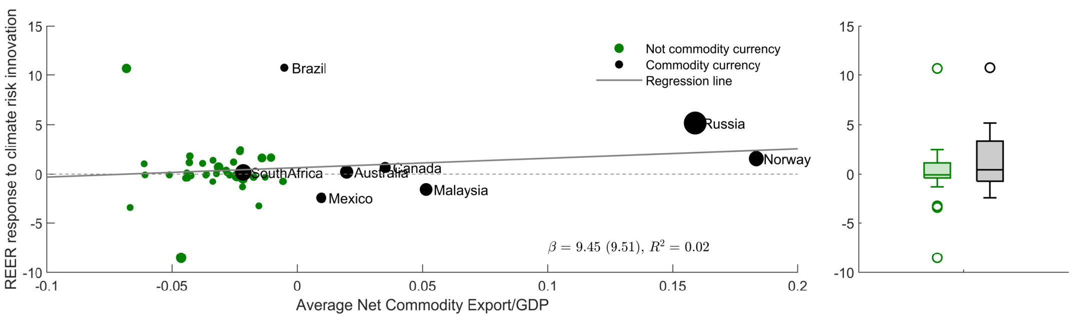

Figure C.15. REER responses for alternative climate risk proxies and commodity export shares. Each graph reports a country's REER response following a one standard deviation exogenous innovation to the Engle et al. (2020), Gavriilidis (2021), and Ardia et al. (2022) climate risk indexes, or temperature anomalies (y-axis: in percentage change). See Figure 6 for additional details.

Appendix D 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.

D.1 Partially pooled panel VAR model

First, rewriting (3) 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 $\boldsymbol{b}_c \sim N(0, \boldsymbol{\Sigma}_b)$. 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:

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

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.

D.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 $\lambda_1^0 = 0.01$.

D.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):

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

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