

# Government Influence in Information Production of International Organizations: Evidence from the IPCC\*

Patrick Bayer<sup>†</sup>

Lorenzo Crippa<sup>‡</sup>

January 12, 2024

**Draft submission for PEIO conference. Please do not cite or share**

## Abstract

A core function of international organisations (IOs) is to provide information to address international cooperation problems. Despite convincing evidence for the importance of information and the conditions under which information facilitates cooperative outcomes, we know little about governments' strategic role in shaping the information that IOs provide. Existing research is largely agnostic about the information production process. In this paper, we develop a theoretical argument that models governments' incentives to intervene in an IO's information production process. For cases when such information will be relevant for shaping negotiation outcomes, we specify at what points along the information production and negotiation cycle governments will exercise their influence. Empirically, we draw on text data from the Intergovernmental Panel on Climate Change (IPCC), the United Nation's scientific body on climate change, to find strong support for our argument: for topics in IPCC scientific assessments that governments expect to matter greatly in UN climate negotiations, governments strategically interfere already at the information production stage. These findings have major implications for our understanding of the role of information in IOs and international cooperation more broadly.

**Keywords:** international organizations; information provision; climate science; IPCC; word embeddings.

---

\*We are grateful to Mark Buntaine, Robert Gulloty, Dafni Kalatzi Pantera, Michael Lerner, Krzysztof Pelc, Bernhard Reinsberg, Charles Rodger, Gabi Spilker, and Muzhou Zhang for comments on earlier drafts of the paper. This research was funded by the Economic and Social Research Council (ESRC) as part of the "The Politics of Science in International Climate Cooperation" (ES/W001373/2). For the purpose of Open Access, the author has applied a Creative Commons Attribution (CC BY) to any Author Accepted Manuscript (AAM) version arising from this submission. Rosanna Crawford and Stuart Brown provided excellent research assistance. All errors are our own.

<sup>†</sup>School of Social and Political Sciences. University of Glasgow. Email: [patrick.bayer@glasgow.ac.uk](mailto:patrick.bayer@glasgow.ac.uk)

<sup>‡</sup>School of Social and Political Sciences. University of Glasgow. Email: [lorenzo.crippa@glasgow.ac.uk](mailto:lorenzo.crippa@glasgow.ac.uk)

## Introduction

Information is essential for international cooperation. Often, this information is provided by international organizations (IOs) and their agencies (Keohane, 1984; Milner, 1997; Abbott and Snidal, 1998). The United Nations (UN) Human Rights Office, for example, reports on violations against the UN Convention against Torture, the World Bank informs foreign investors about host countries' business climate, and the World Health Organization tracks the outbreak of diseases, such as Covid-19, cholera, and measles, in its Weekly Epidemiological Reports. In all these cases, so standard scholarship goes, information facilitates international cooperation by constraining governments. Once information about non-compliant or deviant government behavior becomes public, leaders fear they may lose their political reputation (Hafner-Burton, 2008; Lebovic and Voeten, 2006; Tingley and Tomz, 2021), be punished by markets (Simmons, 2000; Morse, 2019), or take a hit at the ballot box (Dai, 2005; Fang, 2008; Chaudoin, 2014)—all with the same effect of increasing governments' willingness to compromise. Information that is provided by IOs, hence, has the power to shape global governance outcomes because it offers (otherwise unobservable) evidence to other governments, market participants, and voters at home, who—at least under some conditions—can constrain offending governments through naming and shaming, market responses, or elections. In autocracies, greater transparency can cause mass protests and lead to political instability (Hollyer, Rosendorff, and Vreeland, 2015, 2018).

However, for information to have this constraining effect on policymakers, much of the existing literature typically assumes that any information that IOs share with international and domestic audiences is free from government influence. Indeed, if national governments could easily meddle with, for instance, IO reports on human rights violations, the contents of infringement proceedings, or the publication of economic performance data, this would considerably limit an IO's capacity to constrain government behavior. The controversies over data irregularities in the World Bank's "Ease of Doing Business" indicator and the cancellation of the associated report in 2021 are one

prominent example. In such cases, the information that IOs release is no longer independent, but becomes endogenous to government influence. Similar to the discussion about whether international treaties constrain or screen governments (von Stein, 2005; Simmons and Hopkins, 2005), information provided by IOs that has actively been shaped and framed by national governments is unlikely to bind and constrain governments in meaningful ways. Hence, an IO's ability to effectively constrain governments *through information* depends on the extent to which IOs are able to firewall their information production processes from undue outside influence. Information production within IOs may rely heavily on principal governments' input, either directly, through data, expertise, or finances (Abbott and Snidal, 1998; Nielson and Tierney, 2003), or indirectly through informal governance arrangements and staff preferences (Chwieroth, 2013; Clark and Dolan, 2021; Clark and Zucker, 2023)—challenging the assumption that there is no government involvement in IOs' information production.

In this paper, we study the *politics of information production* inside IOs. We argue that the way in which IOs produce information is a political process itself; this is to say that governments have strong incentives to interfere with information production in IOs *exactly* because they recognize the constraining effects of information and increased transparency which international institutions can generate (Mitchell, 1994; Dai, 2005; Fang, 2008; Chaudoin, 2014). Information supplied by IOs often structures international negotiations (Morrow, 1994), and so we expect those states to heavily intervene in IOs' information production whose domestic political economy would suffer high cooperation costs when uninterrupted information were to become publicly available. Our theoretical framework explains government interference in information production in IOs as a function of the distributional effects from international cooperation outcomes which differ depending on what type of information an IO provides.

We test this argument about government interference in information production in the context of the Intergovernmental Panel on Climate Change (IPCC), the United Nation's (UN) primary scientific body on climate change (de Pryck and Hulme, 2022). The IPCC regularly publishes

the most extensive assessments of what we know about the physics, the impacts, and mitigation options of climate change. These reports are highly influential in shaping the discourse in annual climate talks under the UN as they set the scientific guardrails for agreeing a political response to the climate crisis. Taking advantage of original data from both governmental review comments on draft text and the line-by-line negotiations of the IPCC’s key “Summary for Policymakers” (SPM) report, we show that governments interfere more with the text production in the IPCC when their national economies are heavily reliant on the continued use of carbon emissions. The empirical evidence draws on multiple approaches that combine descriptive statistics, the statistical modeling of government behavior, and word embeddings around key target words across draft, interim, and final versions of IPCC text.

Our paper makes two main contributions. First, it demonstrates that governments will interfere in information production in IOs if the otherwise constraining effects of IO information provision were harmful to their domestic political economy. This insight is important because it nuances our understanding of the conditions under which information provided by IOs and international institutions more broadly can credibly constrain governments. Information provision is only likely to effectively limit government behavior if the institutional rules empower IOs to protect information production processes from government influence, highlighting the importance of institutional design when IOs are created ([Abbott and Snidal, 1998](#); [Koremenos, Lipson, and Snidal, 2001](#); [Johnson and Urpelainen, 2014](#)). An immediate implication of this logic for future research is that the constraining power of IOs may be restricted in highly technical policy areas of international cooperation, including cyber security, terrorist financing, banking regulation, global health, or climate change, in which IOs’ information function depends considerably on input by principal governments. Our findings caution against the assumption that information production in IOs is free from politics, but that instead these processes have largely been understudied. Following the growing use of text-as-data approaches in international relations (e.g., [Chaudoin, 2022](#); [Thrall, 2023](#)), we rely on recent methodological advances in modeling the use of words ([Rodriguez and Spirling,](#)

2022; Rodriguez, Spirling, and Stewart, 2023) to empirically study otherwise difficult-to-observe information production processes in IOs.

Second, our research speaks to the existing literature that puts domestic distributional conflict at the core of climate politics (Colgan, Green, and Hale, 2021; Aklin and Mildemberger, 2020; Bayer and Genovese, 2020). We build our argument from the same first principles that concerns about costs from ambitious climate policy will structure opposition to and support for international climate cooperation among firms, sectors, and governments (Bechtel, Genovese, and Scheve, 2019; Genovese, 2019; Kennard, 2020; Cory, Lerner, and Osgood, 2021; Gaikwad, Genovese, and Tingley, 2022; Bayer, 2023). However, we extend this logic one step further and show that the same incentives that are rooted in the domestic political economy operate not only at the level of the international negotiations, but also apply to information production processes that *precede* the actual negotiations. From a normative perspective, this may be disturbing news as it questions the “purity” of information that IOs disseminate and calls for the more systematic study of the politics of information production and, by extension, the politics of science as such.

## **IOs, Cooperation, and Government Influence in Information Provision**

The lack of information typically complicates international cooperation for one of two reasons: either, states are uncertain about a key feature of the cooperation problem itself—commonly referred to as uncertainty about the “true state of the world”—or, states are uncertain about the preferences of others—usually called uncertainty about players’ “types.” While the cost associated with learning about the true state of the world prevents recovering information in the first case, incentives to misrepresent private information about one’s own preferences undermine cooperation in the second case. However, international organizations can help facilitate cooperation in both scenarios.

To begin with, international organizations benefit from economies of scale in information production (Keohane, 1984; Milner, 1997; Koremenos, Lipson, and Snidal, 2001). Pooling resources

reduces the per-unit cost of information gathering and, hence, lowers the risk of free-riding ([Abbott and Snidal, 1998](#); [Mitchell, 1994](#)). International organizations function as an informational clearing house that collects, vets, and disseminates information. In doing so, IOs can effectively change what governments know about (i) the very structure of a cooperation problem itself and about (ii) other governments' cooperation incentives. Greater access to information alone is not a sufficient condition for successful cooperation, of course, as distributional conflict continues to persist ([Morrow, 1994](#); [Goldstein and Martin, 2000](#)), but it certainly reduces the likelihood that cooperation breaks down because of information asymmetries.

International organizations moreover allow governments to make credible commitments. By tying their hands through international law or multilateral agreements, renegeing on previously negotiated cooperation outcomes becomes costly ([Abbott and Snidal, 2000](#); [Koremenos, 2005](#); [Rosendorff, 2005](#)). Scholars have shown that these non-compliance costs can arise from multiple sources. [Simmons \(2000\)](#), [Büthe and Milner \(2008\)](#), and [Gray \(2009\)](#), for example, emphasize a reputational mechanism, whereby legal obligations, often made by signing international agreements or joining IOs, signal valuable information to investors and market participants. Others stress policymakers' fears of losing political (rather than economic) reputation from naming and shaming efforts in cases of treaty violations or attempts to renege on international commitments ([Lebovic and Voeten, 2006](#); [Hafner-Burton, 2008](#); [Tingley and Tomz, 2021](#)). Electoral concerns also shape government incentives for cooperation and compliance through informational constraints created by international institutions ([Dai, 2005](#); [Fang, 2008](#); [Chaudoin, 2014](#)).

No matter what the exact mechanism is that each of these arguments emphasizes, the common thread in the broader literature on credible commitments is that IOs can vouch for the veracity of information. States have incentives to portray their own policies and actions in favorable light. Hence, we generally put little weight on governments' public reassurances that they are, for example, not engaging in torture of political prisoners, that they are not undermining a fragile peace agreement, or that they are not violating IMF mandated loan conditions—unless we have good rea-

sons to believe otherwise. Such reasons are often provided by IOs. As they monitor government behavior and verify information (Mitchell, 1994; Fortna, 2003; Fang and Stone, 2012), IOs benefit from high levels of legitimacy. Since IOs are multilateral, they are perceived as impartial arbiters (Abbott and Snidal, 1998; Dellmuth et al., 2021). This impartiality is at the very core of effective information provision by IOs. Information disseminated by IOs is not tainted by governments' preferences, and can therefore function as a credible (and exogenous) constraint on government behavior.

Notwithstanding the intuitive appeal of this standard argument about IOs as unbiased information providers, its logic, at least on second thought, sits somewhat uncomfortably with the research on government influence in IOs. This literature demonstrates that governments seek control in IOs through both formal and informal means (Stone, 2004; Dreher, Sturm, and Vreeland, 2015; Carter and Stone, 2015; Graham and Serdaru, 2020) as well as through IO staff who may internalize important stakeholder government's interests into their organizational practice (Clark and Dolan, 2021)—giving powerful states disproportional influence in IO decision-making. Building on existing evidence that governments have incentives to shape international cooperation outcomes, we argue that the very same motivations also structure governments' attempts to shape *information production* in IOs.

Specifically, we argue that governments have strong incentives to shape the information that IOs provide because they are aware of the constraining effects of IO information. If this is indeed the case, IO provided information is no longer exogenous, but becomes part of an endogenous political process itself, where governments can shape the information before it becomes publicly released by an IO. The extent to which governments can exercise such control over information production will depend on the institutional design of the specific IO: IOs that largely rely on independent sources, for instance, may be largely shielded from government influence compared to those, such as the Intergovernmental Panel on Climate Change (IPCC), which we study below, for which government inputs and even government approval of the produced information are critical.

In instances where governments have indeed some influence over an IO's information production, we claim that governments seek to protect the distributional losers from the information that IOs release. In our specific case of climate change, and in keeping with evidence from the climate politics literature (Colgan, Green, and Hale, 2021; Aklin and Mildemberger, 2020; Kennard, 2020; Stokes, 2020; Mildemberger, 2020; Green et al., 2022) we expect that governments whose economies are highly reliant on the fossil fuel industry, will attempt to influence the language of the IPCC Summary for the Policymakers texts more aggressively than other countries.

## **Background on the IPCC**

We test our argument in the case of the United Nations' Intergovernmental Panel on Climate Change (IPCC). The IPCC was founded in 1988 by the World Meteorological Organization and United Nations Environment Programme. It is both an IO (with permanent Secretariat in Geneva) and a scientific body which counts 195 member countries. Its goal is to assess and summarize the science, impact, and mitigation options of climate change. Its reports provide crucial policy inputs for governments' international negotiations, such as those happening under the United Nations Framework Convention on Climate Change (UNFCCC). In order to produce such reports, the IPCC does not conduct its own research: instead, it relies on available climate change-related knowledge and draws on peer reviewed, published, and technical literature.

The IPCC is organized in three working groups (WGs) dedicated, respectively, to summarizing available knowledge on: the physical basis of climate change (WGI); impact, adaption, and vulnerability of climate change (WGII); and options to reduce CO2 emissions and mitigate climate change (WGIII). The IPCC operates in "assessment cycles", i.e., rounds of five to seven years which end up in the production of an assessment report (AR), such as the latest AR6 report in 2023.<sup>1</sup> In each cycle, every WG produces at least three relevant documents: a Longer Report (in IPCC jargon this is often referred to as the "underlying report"), a Technical Summary, and a Sum-

---

<sup>1</sup> See: <https://www.ipcc.ch/assessment-report/ar6/>.



mary for Policymakers (SPM). These three documents contribute to the AR that gets produced in a given cycle. IPCC summary reports are usually organized in “headline statements” that synthesize a given set of sub-paragraphs (which we refer to as “sub-headline statements”).

We study the production of the SPM in AR6 of WGIII. We focus on WGIII as it deals with mitigation options and greenhouse gas emission reduction, both particularly relevant policy issues in climate negotiations. Among the various text documents produced by the IPCC, we focus on SPMs because of the significant public attention that they tend to receive. Their content is often reported on newspapers and media sources. Moreover, SPMs are reviewed, discussed, and approved line-by-line by government delegates. The process of revision and government approval produces three different versions of a single SPM. First, WG authors produce an initial SPM draft. We refer to this as the “draft” version of the SPM. In the case of the SPM produced by WGIII for AR6, the document was dated November 28, 2021. The draft is then sent to governments, who submit comments in a process of review on this initial version. WG authors receive comments and incorporate them in a second version, which we refer to as the “interim” SPM. In our case, this version is dated March 16, 2022. The interim version is then discussed over multiple days in a plenary session, where government delegates discuss the SPM line-by-line and must reach consensus on raised issues. For WGIII, in assessment cycle 6, plenary sessions started on March 21. Although initially scheduled to finish on Friday, April 1, a final approval of the SPM was only achieved on Sunday night, April 3, 2022. The outcome of this stage is a third version of the SPM, which we refer to as the “final” version.

This two-stage process of revision can make the final version of an SPM significantly different from the initial one. Figure 1 shows in red insertions and deletions for the single headline statement B.2 of the SPM of WGIII in Assessment Cycle 6, when moving across versions. It exemplifies how the text changed significantly at the end of government approval. The review stage (1a) significantly toned down the extent to which improvements in energy efficiency failed to compensate growing CO<sub>2</sub> emissions from fossil fuels and industry. Moreover, it deleted references to changes

B.2 Net GHG emissions have increased since 2010 across all major sectors globally. For CO<sub>2</sub> An increasing share of emissions is accounted for by activities in urban areas. Emissions reductions in CO<sub>2</sub> from fossil fuels and industry, due to improvements in energy efficiency-intensity of GDP and carbon intensity of energy, have not been sufficient to compensate for growing more than offset by increasing global activity levels in industry, energy supply, transport, buildings, agriculture and land-use change, as well as urbanisation buildings. (high confidence) {2.2, 2.4, 6.3, 7.2, 8.3, 9.3, 10.1, 11.2}

(A) Draft vs interim version

B.2 Net anthropogenic GHG emissions have increased since 2010 across all major sectors globally. An increasing share of emissions is accounted for by activities in urban areas. Emissions reductions in CO<sub>2</sub> from fossil fuels and industrial processes, due to improvements in energy intensity of GDP and carbon intensity of energy, have been less than more than offset by increasing emission increases from rising global activity levels in industry, energy supply, transport, agriculture and buildings. (high confidence) {2.2, 2.4, 6.3, 7.2, 8.3, 9.3, 10.1, 11.2}

(B) Interim vs final version


**FIGURE 1:** Tracked changes across versions of an IPCC SPM. Headline Statement B.2, Working Group III, assessment cycle 6

in land usage. Changes after the plenary session (1b), then, led to a headline statement which further softened the degree to which growing CO<sub>2</sub> emissions have been appropriately reduced.

Textual changes are not limited to this statement. In Figure 2, we describe changes in the whole document across the three versions of WGIII's SPM in AR6 (draft, interim, and final). As unit of analysis, we consider sub-headline statements—individual SPM paragraphs—that were present across all text versions. The figure plots the distributions of two similarity scores across the three different versions of the same sub-headline statement. First, we calculate Levenshtein similarity<sup>2</sup> across versions of the same sub-headline statement (left-hand panel). The text changed significantly between the draft and final version of the SPM: the average sub-headline statement in the draft version has a similarity score of just 0.39 with its final variant. However, we note that a significant change is already detected after the review stage: the average sub-headline statement has already a low Levenshtein similarity score (0.51) between the draft and the interim version.

<sup>2</sup> The measure is based on the Levenshtein distance  $D(a, b)$  between string  $a$  and string  $b$ , defined as the minimum number of changes (insertions or deletions) of characters required to turn  $a$  into  $b$  (or vice-versa). The Levenshtein similarity  $L(a, b)$  is:  $L = 1 - \frac{D(a, b)}{\max(\text{length}(a), \text{length}(b))}$ , where  $\text{length}(\cdot)$  returns the number of characters of a string. As such,  $L$  is bounded between 0 (all characters of  $a$  should be changed in order to produce  $b$  or vice versa) and 1 ( $a$  and  $b$  are identical).

A similar picture is observed when computing cosine similarity between term frequency–inverse document frequency (TF-IDF), which quantifies the extent to which different versions use a similar vocabulary across versions of the same sub-headline statement (right-hand panel).



**FIGURE 2:** Similarity between versions of SPM text sections. Left panel shows the distribution of the Levenshtein similarity. Right panel shows the distribution of the cosine similarity when considering TF-IDF


## Empirics

We can leverage various sources in order to test our argument. At the current, preliminary stage, we concentrate on explaining decisions by governments to intervene on a specific part of the SPM draft. IPCC procedures and data allow us to document otherwise difficult-to-observe attempts by governments to interfere with the information produced by IOs. To this aim, we begin by obtaining data on the three SPM text versions from the IPCC data repository. Excluding the introductory section, the draft SPM produced by WGIII for AR6 contained a total of 135 sub-headline statements, i.e., paragraphs numbered in a progressive manner (e.g., B.1, B.1.1, B.1.2, ...). Each sub-headline statement is a self-contained summary of available scientific evidence related to different aspects

of mitigation or CO2 emission reduction options.

Next, from the IPCC data repository we obtained the exhaustive list of comments submitted by governments when reviewing the draft SPM. This amounts to a total of 3,092 comments, submitted by 43 governments in all. The table includes information on: the government submitting comments, the exact pages and lines of the draft SPM that each comment refers to, the type of comment (classified as either “figures/tables”, “substantive”, or “editorial”), and the text of the submitted comment. Cleaning this data source reduced the number of comments that can be used in our analysis. First, we discard comments submitted by the European Union, as we are unable to clearly allocate them to any specific member state. This brings comment sources down to 42 countries. Furthermore, at the current preliminary stage, we only focus on substantive and editorial comments in our analysis. Finally, a number of comments appeared to refer to non-existing pages and lines of the draft SPM, for which we manually retrieved the correct page and line information of the draft SPM. We discarded comments that we were unable to attribute to any correct page or line. These selections bring the number of comments down to 2,826 individual text annotations, submitted by 42 countries in all.

Figure 3 shows the distribution of the number of comments submitted by each government, distinguishing between substantive and editorial comments. Consistently with our expectations, the most active countries in submitting attempts to meddle with IO-produced information tend to be countries with significantly high stakes in the fossil fuel production or consumption (e.g., the US, France, the UK, Saudi Arabia, Germany, Norway). According to our explanation, these countries are those that would be most vulnerable to ambitious climate actions. However, the figure also shows that smaller countries, for instance small island states threatened by the most dire consequences of climate change, tend to attempt to intervene frequently in the review process (e.g., St. Kitts & Nevis, Jamaica, St. Lucia. . .). Similar patterns do not emerge when looking at editorial comments, which suggests that the mechanisms explaining decisions to submit comments related to the content or the form of the text are distinct.



**FIGURE 3:** Number of review comments submitted by governments for the IPCC SPM produced by WGIII for AR6, draft version

## Variables and Model Specification

In order to explain such attempts by governments to interfere with the information eventually produced by the IPCC, we pair each paragraph from the draft SPM (135 in total) with every country intervening in the review process (42 countries). Our unit of analysis is thus a paragraph-country pair. For each pair, we measure whether the country submitted at least one substantive comment pertaining to the paragraph. We thus obtain a binary variable describing substantive attempts by specific governments to change the content of a given paragraph. We then obtain the same variable related to editorial comments, which we intend as a placebo test. Consistently with evidence shown in Figure 3, the probability that a paragraph-country pair will see a government intervening with at least one substantive comment is rather high, at 0.23. Instead, editorial comments are much rarer (with a probability of 0.05).

To explain our dependent variable, we collect country-level covariates that pertain to our argument. We draw on the World Bank World Development Indicators to measure fossil fuel rents per each country. In particular, we measure oil, coal, and natural gas rents as percentage of gross domestic product (GDP). These variables are our explanatory variables of interest, meant to indicate the extent to which the economy of a country is vulnerable to ambitious climate action that would significantly cut down on fossil fuel usage.

We also gather data on variables that would likely confound our relationship of interest. First, we measure total CO<sub>2</sub> emissions (in Kg of CO<sub>2</sub> per constant GDP), to ensure that results do not reflect the role of more polluting economies. Second, we control for the total natural resources rent (as GDP percentage) to rule out that results do not represent a spurious relationship driven by dependency on natural resources. We also control for constant-price GDP and for percentage GDP growth, because richer economies could have better capacity to submit comments and they could also be more dependent on fossil fuels to sustain their economy (or its growth). Finally, we control

for whether the country is classified as a Small Island Developing State<sup>3</sup>. All our covariates are computed as country-specific averages over the entire assessment cycle 6 and until the year of the review (that is, from 2015 to 2021).

We explain our dependent variables in linear models estimated using ordinary least squares (OLS). Because our covariates are all defined at the country-level, we completely remove paragraph-level heterogeneity with the inclusion of a paragraph-fixed effect across all our specifications. This is meant to account for all factors that contribute to create variation across paragraphs such as: topic, length, technicality, or imprecisions in the reported content. With the inclusion of this fixed effect, our models explain within-paragraph decisions by countries to intervene as a function of their fossil fuel dependence (and covariates), then averages estimated effects across paragraphs. The key identifying assumption that needs to hold, in order for the sign of our estimates to be reliable, is thus a version of the conditional independence assumption relative to country-level features only. No omitted country-level variable should exist that simultaneously increases (decreases) fossil fuel rents and the likelihood to intervene on an IPCC paragraph. Across our specifications, all standard errors are clustered using two-way clustering over paragraph and country, to account for likely correlation in the review activity at these two levels.

## Results

Table 1 shows our initial results when estimating the fixed-effect models described above. We introduce control variables step-wise, to avoid suppression effects (Lenz and Sahn, 2021). As such, in our simplest specification we only include variables related to fossil fuel rents and paragraph fixed effect. Consistently with our argument, we find that the likelihood to submit a substantive comment on a paragraph increases by about 0.01 when oil rents increase by 1 percentage point of GDP. This effect remains significant, and increases in magnitude, when including control variables. For instance, Model 3 estimates an increase in probability by about 0.08, for each percentage point

---

<sup>3</sup> See: <https://www.un.org/ohrrls/content/list-sids>.

increase in oil rents over GDP. Across specifications, effects are always statistically significant at the 0.05 conventional level.

**TABLE 1:** Probability that a country submits a substantive comment on a paragraph of the IPCC WGIII SPM

	Model 1	Model 2	Model 3	Model 4	Model 5
Oil rents (% of GDP)	0.014* (0.006)	0.026*** (0.005)	0.077*** (0.016)	0.062*** (0.014)	0.061*** (0.014)
Coal rents (% of GDP)	0.046 (0.097)	0.273* (0.125)	0.343** (0.121)	0.299*** (0.082)	0.311*** (0.080)
Natural gas rents (% of GDP)	-0.093* (0.044)	-0.047 (0.036)	0.020 (0.033)	0.020 (0.032)	0.023 (0.031)
CO2 emissions (kg per 2015 US\$ of GDP)		-0.386* (0.155)	-0.358** (0.131)	-0.425*** (0.116)	-0.441*** (0.112)
Total natural resources rents (% of GDP)			-0.054** (0.015)	-0.035* (0.013)	-0.034* (0.013)
GDP (constant 2015 US\$, trillions)				0.036*** (0.005)	0.036*** (0.005)
GDP growth (annual %)				0.006 (0.011)	0.008 (0.012)
Small Island Developing State					0.050 (0.097)
Sub-headline statement FE	Yes	Yes	Yes	Yes	Yes
Num.Obs.	5535	5535	5535	5535	5535
R2	0.087	0.112	0.135	0.219	0.220
R2 Adj.	0.064	0.090	0.113	0.199	0.200

+ p < 0.1, \* p < 0.05, \*\* p < 0.01, \*\*\* p < 0.001

An effect consistent with our argument is also detected when looking at other fossil fuel rents. For example, larger coal rents are associated with an increase in the likelihood to intervene on a paragraph. Models 2 to 5 predict that the likelihood to intervene on a paragraph increases by about 0.27 to 0.31 for countries that derive coal rents that are 1 percentage point of GDP larger. The effect is large and significant across these specifications, but in our sparsest model (Model 1). An effect that contradicts our argument is, instead, observed in Model 1, for what concerns natural gas rents. Here, the estimate is negative and statistically significant at a 0.05 conventional level. However, this effect is not robust to the inclusion of control variables.



Next, we offer a placebo test to bolster our results from the previous table. We replicate all models from 1 replacing our binary dependent variable with the analogous version capturing whether a government submitted an editorial comment on a paragraph. If our logic is correct, and countries attempt to intervene on the information produced by an IO in order to shield vulnerable domestic constituencies, our variables of interest should not generate strong effects on this dependent variable. Table 2 reports our findings. We find small to non-significant effects for oil rents, as expected. Instead, we do detect a positive and statistically significant effect for coal and natural gas rents (models 2 to 5). However, these effects appear smaller than previously estimated.

**TABLE 2:** Probability that a country submits an editorial comment on a paragraph of the IPCC WGIII SPM

	Model 1	Model 2	Model 3	Model 4	Model 5
Oil rents (% of GDP)	-0.003* (0.001)	0.002 (0.002)	0.010 (0.007)	0.007 (0.006)	0.007 (0.006)
Coal rents (% of GDP)	0.019 (0.032)	0.120* (0.048)	0.130* (0.051)	0.121** (0.045)	0.110* (0.047)
Natural gas rents (% of GDP)	0.014 (0.013)	0.035+ (0.018)	0.045* (0.019)	0.044* (0.019)	0.041* (0.019)
CO2 emissions (kg per 2015 US\$ of GDP)		-0.171* (0.066)	-0.167* (0.070)	-0.176* (0.070)	-0.161* (0.070)
Total natural resources rents (% of GDP)			-0.008 (0.007)	-0.005 (0.006)	-0.006 (0.007)
GDP (constant 2015 US\$, trillions)				0.006* (0.002)	0.005* (0.002)
GDP growth (annual %)				-0.0004 (0.005)	-0.002 (0.006)
Small Island Developing State					-0.046+ (0.025)
Sub-headline statement FE	Yes	Yes	Yes	Yes	Yes
Num.Obs.	5535	5535	5535	5535	5535
R2	0.045	0.063	0.064	0.072	0.075
R2 Adj.	0.021	0.039	0.040	0.048	0.051

+  $p < 0.1$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

## Content of Review Comments: Word Embeddings

So far, we have only studied the likelihood that governments will intervene on a specific paragraph as a function of the importance of fossil fuel for their countries' domestic economies. We have completely abstracted the content of the submitted comments from our analysis. We propose here one test that bolsters our claim and complements the analysis proposed. We delve into the content of governments' submitted comments by studying the semantics of key words used in such review process. Our aim is to show that governments' understanding of salient mitigation policy options changes as a function of domestic fossil fuel rents.

Our approach to the study of governments' understanding of mitigation policy options relies on word embeddings (Rodriguez and Spirling, 2022). This methodology is gaining popularity in political science (Rheault and Cochrane, 2020; Rodman, 2020). Embeddings are representations of “features” of the text—i.e., words or multi-word expressions—as multi-dimensional vectors. Under this representation, each instance of a feature corresponds to a  $1 \times D$  vector, defined in terms of its proximity to  $D$  other features. Underlying this approach is the “distributional hypothesis”, i.e., the idea that the particular meaning of a feature depends on its proximity to other features used in the text. This hypothesis allows to account for polysemy: the feature “bank” will be closer to financial terms if it means a financial institution, but it will be associated to navigation terms if it refers to the land alongside a river. More interestingly for us, representing words in vector spaces allows to estimate semantic differences among different usages of the same features, as a function of subgroup of the corpus—e.g., defined based on characteristics of the actors producing the text.

We adopt the integration of word embeddings in a regression framework proposed by Rodriguez, Spirling, and Stewart (2023), which allows us to estimate the embeddings of target features that occur rarely. This is convenient in our highly technical case, where governments discuss salient mitigation policy options only in the context of SPM paragraphs that directly pertain to them.<sup>4</sup> The

---

<sup>4</sup> For instance, one of the target features we focus on below appears in just 115 instances.

methodology is designed to study semantic differences in the usage of a certain “target feature”, i.e., a word or multi-word expression of interest. The framework relies on three main ingredients. First, a set of pre-trained word embeddings covering a sufficiently large vocabulary, usually estimated based on large corpora. We follow the authors and use the Global Vectors (GloVe) word vectors estimated by Stanford University based on the 2014 corpus of Wikipedia pages (Pennington, Socher, and Manning, 2014).<sup>5</sup> Second, a transformation matrix used to discount the weight of very common words like “stop-words”. We use the matrix estimated by Khodak et al. (2018). Finally, a local corpus containing instances of the target feature that are of interest, associated with document-level variables that can be used as explanatory covariates. In our case, the corpus is represented by the 2,457 substantive comments submitted by governments on the SPM draft. We treat each comment as a separate document. Covariates are the same country-level variables defined in the previous analysis.

The methodology computes the word embedding of a specific instance of the target feature “a la carte” (ALC), i.e., by drawing on the pre-trained GloVe embeddings of words appearing around the target. We define a symmetrical ten-feature-long “context”, that is, we consider ten features to the left and to the right of each instance of the target in order to compute the ALC embedding for each instance. The pre-trained GloVe embeddings of words in the context are discounted based on the transformation matrix to down-weight very common words. This produces one vector per each instance of the target word, with dimensions  $1 \times D$ , where  $D$  is the total length of the vector space (300, in our case). Next, the  $N$  computed ALC vectors—one per instance of the target word—are stacked in an  $N \times D$  vector. This represents our dependent variable of interest, that is explained in a multivariate regression model as a function of the document-level covariates. Each coefficient is, in turn, a  $1 \times D$  vector. Such vectors are then normalized, by taking the Euclidean norm, to produce a single scalar. With the same logic, standard errors are estimated via bootstrap, and

---

<sup>5</sup> Among the various versions of GloVe, we rely on the Wikipedia 2014 + Gigaword 5 (6B tokens, 400K vocabulary, uncased, 300-dimensional vector).

empirical p-values for tests of hypotheses can be computed by running permutation tests. Finally, interpretation is done by studying which (pre-trained) embeddings of features the target tends to be more often associated to, according to specific values of the covariates.

We focus on two particularly salient mitigation policy options to identify target features in our corpus of government comments. The first is carbon capture and storage (CCS), consisting in the process of capturing CO<sub>2</sub> emitted from industrial activities and storing it underground. The technology is controversial, due to leakage risks, low effectiveness, and high costs. Furthermore, the utilisation of captured CO<sub>2</sub>, for instance in the production of blue hydrogen, chemicals, and soft drinks, is criticized as an effort to delay the renewable energy transition. Controversy is also reflected in our data. In the draft version of the IPCC SPM, paragraphs referring to the cost and adoption of low-emission technologies (B.4 and B.4.1) are among the single paragraphs that receive the most comments in the review process—a total of 61 and 53 comments, making them respectively first and third paragraph by number of comments received. Paragraphs mentioning CCS’ implementation for achieving net zero GHG emissions from industry (C.5, C.5.2, C.5.3, C.5.4) receive a total of 136 comments. Table 3 reports a sample of ten comments mentioning CCS, their context, and the government reporting the comment.

**TABLE 3:** Random sample of comments including the target feature ‘CCS’

Country	Text
United States	what is the role of fossil_fuels with [carbon_capture_and_storage] ccs given government investment and planning in these technologies it
France	with co2 storage capacity but it is better to say [ccs] capacity
St. Kitts & Nevis	increasing evidence that blue hydrogen produced with natural gas [ccs] has high ghg_emissions and high fugitive methane emissions in particular
China	original text in section of the underlying report [ccs] also addresses a key challenge of the technical aspect of
Norway	page line where it is stated that [ccs] is included in pathways and on page line
France	technological issues and trade_off of [ccs] should be added to this list
Saudi Arabia	which allows all types of technologies to be employed additionally [ccs] has been presented in the most of the reviewed literature
Germany	be avoided in the chemical industry please add along with [ccs] for process related co2_emissions
Australia	unclear from the current text how much [ccs] has been assumed in these reductions are these reductions in
United States	estimating the potential role of ccs can the confidence that [ccs] will be employed really be high based on the literature

The second mitigation policy option we consider relates to land usage. Agriculture, Forestry, and Land Usage (AFOLU) is an acronym used by the IPCC and UNFCCC to refer to CO<sub>2</sub> emissions from agriculture and forestry. A related acronym is Land Use, Land-Use Change and Forestry (LULUCF). Because CO<sub>2</sub> in the atmosphere can be abated in the form of carbon in vegetation,

AFOLU, and in particular deforestation, is a main contributor of global carbon dioxide emissions. In this context, LULUCF is considered a valuable mitigation option with short-term effects of reducing carbon dioxide in the atmosphere, through activities like afforestation or reforestation. However, such practices are reversible, vulnerable to human activity or natural disasters, and thus potentially fail at presenting a long-term solution to mitigate CO<sub>2</sub> emissions.

We analyze the semantics used by government delegates to refer to these discussed policy options by studying variation in the embeddings around two different sets of mitigation policy target words: CCS<sup>6</sup> and LULUCF.<sup>7</sup> Our aim is to study that semantics of these policy options change significantly as a function of domestic fossil fuel rents. We also test differences in semantics around the usage of generic target features related to emissions.<sup>8</sup> Finally, we run a placebo test aimed at showing that semantics do not change significantly as a function of fossil fuel rents when our target features are common words related to the review process.<sup>9</sup> We fit a multivariate regression model of the embedding spaces for each of these target words drawing on the full specification of our previous analysis, excluding paragraph fixed-effects.

Table 4 reports our results. Semantics around our target words of interest change significantly as a function of domestic fossil fuel rents (oil, coal, and natural gas), even when holding constant sources of variation captured by important confounders.<sup>10</sup> Instead, significant differences are not detected for the placebo expressions, which suggests that the differences observed are specific to the salience of mitigation policy options.

We interpret the results from model 1, which explains changes in semantics around the target feature CCS. In order to interpret our findings substantively, we compute the cosine similarity between the predicted target word's embedding, at given levels of oil and coal rents, and the em-

---

<sup>6</sup> Including: "ccs", "carbon\_capture\_and\_storage".

<sup>7</sup> Including: "lulucf", "afolu", and misspelling of the acronym "aflo" and "afolu".

<sup>8</sup> Including: "co2", "co2\_emission", "co2\_emissions", "emission", "emissions", "ghg\_emission", "ghg\_emissions".

<sup>9</sup> We study placebo target words considering exhortation to rephrase certain parts of the text ("rephrase", "rephrased", "rephrasing"), consider changes ("consider", "considerd", "considere", "considered", "considers", "reconsider"), and make notes ("note", "noted", "notes").

<sup>10</sup> In model 2, we are unable to include the binary variable for SIDS due to singularity issues.

**TABLE 4:** Multivariate regression models of target words' embeddings in governments' submitted comments

	Test			Placebo		
	CCS <sup>1</sup>	LULUCF <sup>2</sup>	Emissions <sup>3</sup>	Rephrase <sup>4</sup>	Consider <sup>5</sup>	Note <sup>6</sup>
Oil rents (% of GDP)	0.353*	0.347*	0.088***	1.206	0.610	0.705
	(0.061)	(0.063)	(0.008)	(0.871)	(0.184)	(0.467)
Coal rents (% of GDP)	2.139***	1.543	0.420***	4.626	1.675	3.899
	(0.282)	(0.371)	(0.030)	(2.351)	(0.373)	(4.705)
Natural gas rents (% of GDP)	0.484*	0.553**	0.158***	2.032+	0.744	3.534
	(0.199)	(0.325)	(0.013)	(2.425)	(0.258)	(6.736)
CO2 emissions (kg per 2015 US\$ of GDP)	1.380**	1.390+	0.430***	5.120	1.380	3.156
	(0.296)	(0.465)	(0.023)	(2.204)	(0.377)	(1.280)
Total natural resources rents (% of GDP)	0.333*	0.298*	0.087***	1.145	0.582	0.712
	(0.050)	(0.071)	(0.008)	(0.864)	(0.174)	(1.879)
GDP growth (annual %)	0.102	0.115***	0.027***	0.637	0.063+	0.128
	(0.046)	(0.038)	(0.002)	(0.326)	(0.009)	(0.087)
Small Island Developing State	0.521+		0.260***	1.350	0.624	0.697
	(0.083)		(0.017)	(0.354)	(0.119)	(0.132)
Target word occurrences	115	155	1382	48	228	49
Unique features in context	606	713	2591	361	1136	476
Vector space length	300	300	300	300	300	300
Skip-gram window size	10	10	10	10	10	10

<sup>1</sup> Instances of target feature: carbon\_capture\_and\_storage, ccs

<sup>2</sup> Instances of target feature: aflolu, aflou, afolu, lulucf

<sup>3</sup> Instances of target feature: co2, co2\_emission, co2\_emissions, emission, emissions, ghg\_emission, ghg\_emissions

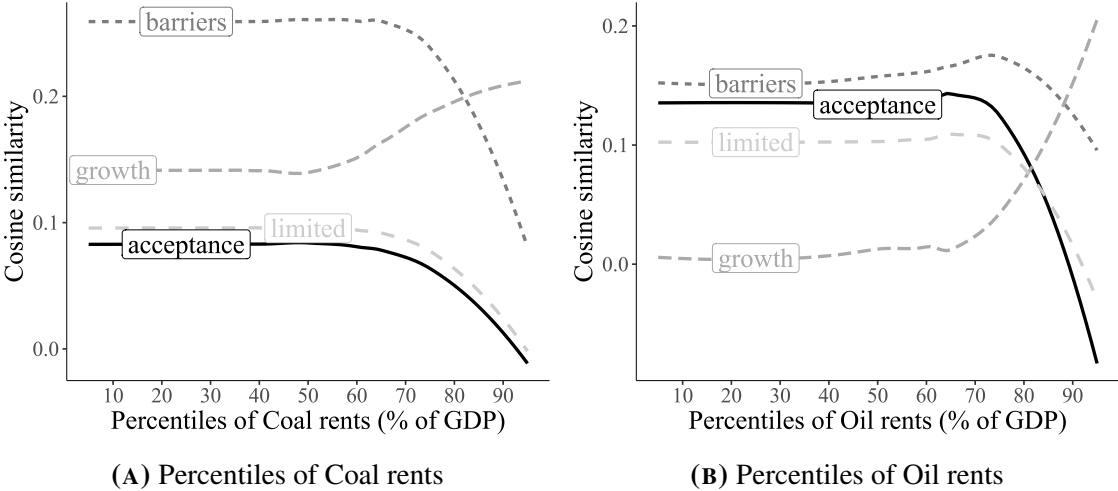
<sup>4</sup> Instances of target feature: rephrase, rephrased, rephrasing

<sup>5</sup> Instances of target feature: consider, considerd, considere, considered, considers, reconsider

<sup>6</sup> Instances of target feature: note, noted, notes

+ p < 0.1, \* p < 0.05, \*\* p < 0.01, \*\*\* p < 0.001

beddings of other features in the context. We choose four features from the context: “barriers”, referring to the limits of CCS; “limited” and “acceptance”, referring to the public skepticism towards the technology; and “growth”, referring to the drop in costs of CCS as a technology for emission mitigation. Figure 4 shows that model 1 predicts that countries with lower fossil fuel rents tend to discuss the target feature “CCS” by referring to its “barriers” or its “limited” “acceptance”, more than do countries with higher domestic fossil fuel rents. The opposite is true when looking at “growth”. The embeddings for CCS tend to be more similar to those for the words “barriers”, “limited”, and “acceptance” when looking at lower percentiles of the distributions of Coal (4a) and Oil rents (4b). Instead, the reverse relationship is found when looking at the similarity of the target word embedding with the embedding of “growth” and fossil fuel rents.



**FIGURE 4:** Cosine similarity between word embeddings of target feature (CCS) and a set of context features, at percentiles of fossil fuel rents

## Conclusion

International organizations hold power because of the information they can provide. While existing research treats information typically as exogenous to cooperation problems, we argue in this paper that information production itself is likely to be part of a political process: governments

have incentives to shape the information that IOs provide to other governments, non-governmental observers, and their domestic public. When governments are given the opportunity to influence the production of IO information, such as in the case of the IPCC, where government participation is mandated in the rules of this UN body, we expect to see government attempt to shape the wording of the text in the highly influence IPCC reports.

Drawing on a combination of descriptive evidence, regression models, and word embeddings, we provide strong evidence that government influence in the IPCC's Summary for Policymakers text is a function of a country's structural dependence on fossil fuels. These findings have important implications. For scholars studying the role of information in international organizations, we caution that information may not always be free from government interests which offers new insight into the conditions under which information can effectively catalyze international cooperation. Future research on the international political economy of climate politics can build on our results which extend the logic of distributional effects of ambitious decarbonization beyond international negotiations themselves to the prior stage of information production in IOs.



## References

- Abbott, Kenneth W., and Duncan Snidal. 1998. "Why States Act through Formal International Organizations." *The Journal of Conflict Resolution* 42 (1): 3–32.
- Abbott, Kenneth W., and Duncan Snidal. 2000. "Hard and Soft Law in International Governance." *International Organization* 54 (3): 421–456.
- Aklin, Michaël, and Matto Mildemberger. 2020. "Prisoners of the Wrong Dilemma: Why Distributive Conflict, Not Collective Action, Characterizes the Politics of Climate Change." *Global Environmental Politics* 20 (4): 4–27.
- Bayer, Patrick. 2023. "Foreignness as an Asset: European Carbon Regulation and the Relocation Threat among Multinational Firms." *Journal of Politics* FirstView.
- Bayer, Patrick, and Federica Genovese. 2020. "Beliefs About Consequences from Climate Action Under Weak Climate Institutions: Sectors, Home Bias, and International Embeddedness." *Global Environmental Politics* 20 (4): 28–50.
- Bechtel, Michael M., Federica Genovese, and Kenneth F. Scheve. 2019. "Interests, Norms and Support for the Provision of Global Public Goods: The Case of Climate Co-operation." *British Journal of Political Science* 49 (4): 1333–1355.
- Büthe, Tim, and Helen V. Milner. 2008. "The Politics of Foreign Direct Investment into Developing Countries: Increasing FDI through International Trade Agreements?" *American Journal of Political Science* 52 (4): 741–762.
- Carter, David B., and Randall W. Stone. 2015. "Democracy and Multilateralism: The Case of Vote Buying in the UN General Assembly." *International Organization* 69 (1): 1–33.
- Chaudoin, Stephen. 2014. "Audience Features and the Strategic Timing of Trade Disputes." *International Organization* 68 (4): 877–911.
- Chaudoin, Stephen. 2022. "How International Organizations Change National Media Coverage of Human Rights." *International Organization* .
- Chwieroth, Jeffrey M. 2013. "'The Silent Revolution': How the Staff Exercise Informal Governance over IMF Lending." *Review of International Organizations* 8 (2): 265–290.
- Clark, Richard, and Lindsay R. Dolan. 2021. "Pleasing the Principal: U.S. Influence in World Bank Policymaking." *American Journal of Political Science* 65 (1): 36–51.
- Clark, Richard, and Noah Zucker. 2023. "Climate Cascades: IOs and the Prioritization of Climate Action." *American Journal of Political Science* .
- Colgan, Jeff D., Jessica F. Green, and Thomas N. Hale. 2021. "Asset Revaluation and the Existential Politics of Climate Change." *International Organization* 75 (2): 586–610.

- Cory, Jared, Michael Lerner, and Iain Osgood. 2021. "Supply Chain Linkages and the Extended Carbon Coalition." *American Journal of Political Science* 65 (1): 69–87.
- Dai, Xinyuan. 2005. "Why Comply? The Domestic Constituency Mechanism." *International Organization* 59 (2): 363–398.
- de Pryck, Kari, and Michael Hulme, eds. 2022. *A Critical Assessment of the Intergovernmental Panel on Climate Change*. Cambridge, UK: Cambridge University Press.
- Dellmuth, Lisa, Jan Aart Scholte, Jonas Tallberg, and Soetkin Verhaegen. 2021. "The Elite–Citizen Gap in International Organization Legitimacy." *American Political Science Review* 116 (1): 283–300.
- Dreher, Axel, Jan-Egbert Sturm, and James R. Vreeland. 2015. "Politics and IMF Conditionality." *Journal of Conflict Resolution* 59 (1): 120–148.
- Fang, Songying. 2008. "The Informational Role of International Institutions and Domestic Politics." *American Journal of Political Science* 52 (2): 304–321.
- Fang, Songying, and Randall W. Stone. 2012. "International Organizations as Policy Advisors." *International Organization* 66 (4): 537–569.
- Fortna, Virginia Page. 2003. "Scraps of Paper? Agreements and the Durability of Peace." *International Organization* 57 (2): 337–372.
- Gaikwad, Nikhar, Gederica Genovese, and Dustin Tingley. 2022. "Creating Climate Coalitions: Mass Preferences for Compensating Vulnerability in the World's Two Largest Democracies." *American Political Science Review* 116 (4): 1165–1183.
- Genovese, Federica. 2019. "Sectors, Pollution, and Trade: How Industrial Interests Shape Domestic Positions on Global Climate Agreements." *International Studies Quarterly* 63: 819–836.
- Goldstein, Judith, and Lisa L. Martin. 2000. "Legalization, Trade Liberalization, and Domestic Politics: A Cautionary Note." *International Organization* 54 (3): 603–632.
- Graham, Erin R., and Alexandria Serdaru. 2020. "Power, Control, and the Logic of Substitution in Institutional Design: The Case of International Climate Finance." *International Organization* 74 (4).
- Gray, Julia. 2009. "International Organization as a Seal of Approval: European Union Accession and Investor Risk." *American Journal of Political Science* 53 (4): 931–949.
- Green, Jessica, Jennifer Hadden, Thomas Hale, and Paasha Mahdavi. 2022. "Transition, Hedge, or Resist? Understanding Political and Economic Behavior toward Decarbonization in the Oil and Gas Industry." *Review of International Political Economy* 29 (6): 2036–2063.
- Hafner-Burton, Emilie. 2008. "Sticks and Stones: Naming and Shaming the Human Rights Enforcement Problem." *International Organization* 62 (4): 689–716.

- Hollyer, James R., B. Peter Rosendorff, and James Raymond Vreeland. 2015. "Transparency, Protest, and Autocratic Instability." *American Political Science Review* 109 (4): 764–784.
- Hollyer, James R., B. Peter Rosendorff, and James Raymond Vreeland. 2018. *Information, Democracy, and Autocracy: Economic Transparency and Political (In)Stability*. Cambridge, MA: Cambridge University Press.
- Johnson, Tana, and Johannes Urpelainen. 2014. "International Bureaucrats and the Formation of Intergovernmental Organizations: Institutional Design Discretion Sweetens the Pot." *International Organization* 68 (1): 177–209.
- Kennard, Amanda. 2020. "The Enemy of My Enemy: When Firms Support Climate Change Regulation." *International Organization* 74 (2): 187–221.
- Keohane, Robert O. 1984. *After Hegemony: Cooperation and Discord in the World Political Economy*. Princeton, NJ: Princeton University Press.
- Khodak, Mikhail, Nikunj Saunshi, Yingyu Liang, Tengyu Ma, Brandon Stewart, and Sanjeev Arora. 2018. "A la carte embedding: Cheap but effective induction of semantic feature vectors." *arXiv preprint arXiv:1805.05388* .
- Koremenos, Barbara. 2005. "Contracting around International Uncertainty." *The American Political Science Review* 99 (4): 549–565.
- Koremenos, Barbara, Charles Lipson, and Duncan Snidal. 2001. "The Rational Design of International Institutions." *International Organization* 55 (4): 761–799.
- Lebovic, James H., and Erik Voeten. 2006. "The Politics of Shame: The Condemnation of Country Human Rights Practices in the UNCHR." *International Studies Quarterly* 50 (4): 861–888.
- Lenz, Gabriel S, and Alexander Sahn. 2021. "Achieving statistical significance with control variables and without transparency." *Political Analysis* 29 (3): 356–369.
- Mildenberger, Matto. 2020. *Carbon Captured: How Business and Labor Control Climate Politics*. Cambridge, MA: MIT Press.
- Milner, Helen V. 1997. *Interests, Institutions, and Information: Domestic Politics and International Relations*. Princeton, NJ: Princeton University Press.
- Mitchell, Ronald B. 1994. "Regime Design Matters: Intentional Oil Pollution and Treaty Compliance." *International Organization* 48 (3): 425–458.
- Morrow, James D. 1994. "Modeling the Forms of International Cooperation: Distribution Versus Information." *International Organization* 48 (3): 387–423.
- Morse, Julia C. 2019. "Blacklists, Market Enforcement, and the Global Regime to Combat Terrorist Financing." *International Organization* 73 (3): 511–545.

- Nielson, Daniel L., and Michael J. Tierney. 2003. "Delegation to International Organizations: Agency Theory and World Bank Environmental Reform." *International Organization* 57 (2): 241–276.
- Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. pp. 1532–1543.  
**URL:** <http://www.aclweb.org/anthology/D14-1162>
- Rheault, Ludovic, and Christopher Cochrane. 2020. "Word embeddings for the analysis of ideological placement in parliamentary corpora." *Political Analysis* 28 (1): 112–133.
- Rodman, Emma. 2020. "A timely intervention: Tracking the changing meanings of political concepts with word vectors." *Political Analysis* 28 (1): 87–111.
- Rodriguez, Pedro L., and Arthur Spirling. 2022. "Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research." *Journal of Politics* 84 (1): 101–115.
- Rodriguez, Pedro L., Arthur Spirling, and Brandon M. Stewart. 2023. "Embedding Regression: Models for Context-Specific Description and Inference." *American Political Science Review* FirstView.
- Rosendorff, B. Peter. 2005. "Stability and Rigidity: Politics and Design of the WTO's Dispute Settlement Procedure." *American Political Science Review* 99 (3): 389–400.
- Simmons, Beth A. 2000. "International Law and State Behavior: Commitment and Compliance in International Monetary Affairs." *American Political Science Review* 94 (4): 819–835.
- Simmons, Beth, and Daniel J. Hopkins. 2005. "The Constraining Power of International Treaties: Theory and Methods." *American Political Science Review* 99 (4): 623–631.
- Stokes, Leah C. 2020. *Short Circuiting Policy: Interest Groups and the Battle over Clean Energy and Climate Policy in the American States*. Oxford: Oxford University Press.
- Stone, Randall W. 2004. "The Political Economy of IMF Lending in Africa." *The American Political Science Review* 98 (4): 577–591.
- Thrall, Calvin. 2023. "Informational Lobbying and Commercial Diplomacy." Working Paper.
- Tingley, Dustin, and Michael Tomz. 2021. "The Effects of Naming and Shaming on Public Support for Compliance with International Agreements: An Experimental Analysis of the Paris Agreement." *International Organization* 76 (2): 445–468.
- von Stein, Jana. 2005. "Do Treaties Constrain or Screen? Selection Bias and Treaty Compliance." *American Political Science Review* 99 (4): 611–622.

# **Government Interference in Information Production of International Organizations**

—SUPPLEMENTARY MATERIALS—

## **Contents**

**1 Section**

**App-2**

# 1 Section