Geoeconomic Pressure

Christopher Clayton[*] Antonio Coppola[†] Matteo Maggiori[‡] Jesse Schreger[§]

June 2025

Paper is Updated at High Frequency: Current Version Here

Visualization and Data: GCAP Geoeconomic Monitor

Abstract

Geoeconomic pressure—the use of existing economic relationships by governments to achieve geopolitical or economic goals—is a prominent feature of global power dynamics. This paper introduces a methodology using large language models (LLMs) to systematically identify the application of and response to geoeconomic pressure from large textual corpora. We classify which governments apply pressure to which foreign targets, using which instruments, firms, and products. We demonstrate that firms affected by tariffs respond primarily with price changes whereas firms affected by export controls respond disproportionately by investing in research and development. We document significant heterogeneity in how firms respond to pressure based on whether their home government is applying the pressure, whether their home country is the recipient of the pressure, or whether they are based in an affected third party country. Finally, we quantify the degree of measurement uncertainty generated by the LLM-based analysis by comparing the classifications across multiple open-weight models as well as considering a wide range of variations of our prompts.

Keywords: Geoeconomics, Geopolitics, Artificial Intelligence, LLMs, Economic Coercion, Economic Security, Economic Dependency, Chokepoints, Sanctions, Export Controls, Tariffs.

JEL Codes: C4, F3, F4, G3.

[*] Yale School of Management and NBER; christopher.clayton@yale.edu.

[†] Stanford University Graduate School of Business and CEPR; acoppola@stanford.edu.

[‡] Stanford University Graduate School of Business, NBER, and CEPR; maggiori@stanford.edu.

[§] Columbia Business School, NBER, and CEPR; jesse.schreger@columbia.edu.

We thank Nick Bloom, Matt Gentzkow, Tarek Hassan, Steve Redding, and Christoph Trebesch for useful comments. We thank Chiara Bargellesi, Angelo Cataldo, Max Feng, Serdil Tinda, Yicheng Yang, and Zhongyu Yin for outstanding research assistance. We thank the Smith Richardson Foundation, Stanford Center on China's Economy and Institutions, and Stanford Impact Labs for financial support. Clayton thanks the NSF (2441937) for financial support.

1 Introduction

A defining feature of the current global system is the willingness of the Great Powers to use their economic and financial strength to achieve geopolitical or economic goals. This rise of geoeconomic activity is a major departure from the last twenty years of policymaking and has the potential to dramatically alter the landscape of the global economic and financial system.

Prominent examples of this new wave of geoeconomic pressure are well-known. For instance, in the last days of the Biden administration, the United States announced a new "Framework for Artificial Intelligence Diffusion" to restrict the ability of American firms to sell cutting-edge semiconductors to China and its allies; China is currently engaged in ongoing threats to suspend rare earth exports to the United States in retaliation against the Trump administration's tariffs; and Russia has repeatedly threatened to suspend natural gas shipments to Europe.

As these events drive major shifts in markets, the real economy, and the international policy environment, academics and policymakers are trying to assess and document the use of these policies. The fundamental challenge is that pressure can take many forms. The means used can include sanctions, tariffs, export controls, regulation, boycotts, and other forms of pressure. The threats might or might not be realized on the equilibrium path. The response of the targets might also be multifaceted including supply chain rearrangement, prices, quantities, domestic investment and R&D, inventories, and a range of other adjustments. In many policy applications, the availability of near real-time information also has a huge premium. In this paper, we show that recent advances in large language models (LLMs) and the widespread availability of text that covers these issues, either from the corporate executives or from professional analysts, can help shed light on these crucial topics.

To organize the analysis, we use as a guiding principle the geoeconomics framework developed by Clayton et al. (2023, 2024, 2025a). Hegemons use both domestic and foreign entities (e.g., foreign firms or governments) to apply pressure to targets. These instruments are used in threats that might or might not be realized, and can involve not selling an input to the target (e.g. a semiconductor export control), not buying a product from the target (e.g. an oil sanction on Russian exports), or imposing tariffs or subsidies to attempt to onshore or friendshore a particular line of production. The targets respond to the pressure (and the induced uncertainty) by rearranging their economic activities. Some of the reaction might be the intended purpose of the pressure: e.g., not selling a particular chip to China. Some of the reaction, on the other hand, may be the target's own best response to dilute the pressure, for instance, by developing an in-house product or finding an alternative supplier.

This organization leads to an analysis based on senders, receivers, means, and instruments. For example, when the US government pressured the Netherlands-based company ASML Holdings NV not to sell advanced lithography machinery to China, we categorize the US government as the sender of the pressure, China as the receiver of the pressure, a product in the semiconductor industry by a Dutch firm as the means, and export controls as the policy instrument. Our methodology applies an LLM as a classification tool to documents that might describe this episode. For example, a transcript of an earnings call by ASML's CEO might discuss this episode, or a transcript by a Chinese firm affected by this episode might provide related information. In each case, we also use the LLM to classify the reactions of the firms involved. Is ASML affected positively or negatively? Are the targeted Chinese firms trying to find alternative suppliers, investing in their own ability to create the input they lost access to, or responding on other margins? Ultimately, the LLM is a tool to transform a vast amount of text into a structured dataset, with the structure dictated by our organizing framework.

We combine a range of firm-level text, including English language earnings calls featuring discussions between executives and analysts, earnings calls for the universe of Chinese firms with A-shares, and sell-side analyst research reports at the firm and sector level. We analyze these textual corpora by performing large-scale inference using open-source, open-weights large language models (LLMs) including Meta's Llama 3.3, Google's Gemma 3, and Alibaba's Qwen 2.5. We perform all inference locally on private GPU computing infrastructure with controlled pseudo-randomization to ensure computational reproducibility. Our sampling strategy from the model uses a zero sampling temperature (i.e., a greedy sampling strategy) to abstract from arbitrary sampling variability: we instead, in the last section of the paper, analyze variability induced by model and prompt perturbations. To improve computational efficiency, we first process every document using a basic prompt that flags whether the document discusses a set of policies: tariffs, sanctions, export controls, boycotts, investment screening, and subsidies for on-shoring or friend-shoring. This step is perhaps the closest to what could be accomplished by a traditional dictionary-based NLP method that flags whether particular words are included in the document.

Our analysis focuses on a second step in which documents that have been flagged are then processed by a detailed prompt that is specific to the instrument and type of document being analyzed. For example, if a particular earnings call from ASML has been flagged for export controls, that text is then fed to the LLM again with a detailed prompt designed for export controls. The detailed prompts ask approximately 55 questions requesting structured output for each, and also requesting the LLM to produce a 100-word summary of its entire analysis. The detailed questions come in two types. A first set of questions characterizes the type of pressure: who is applying it (sender), toward whom (receiver), using which firms and their inputs or products (means), and whether the policy tools are currently applied or threatened to be applied in the future. A second set of questions extensively characterizes the firm's reaction: sales, purchases, inventories, customer and suppliers sourcing decisions, labor and hours, investment, R&D, pricing on inputs and products, and profit margins. For each field, we ask the LLM to characterize whether the activity is increasing or decreasing, and also the geography (domestic versus foreign, but also which foreign countries). The LLM produces structured output in the form of either boolean flags (e.g., a zero-one flag for whether domestic investment increased), lists of countries, or less structured text (e.g. the products or inputs that the firm reports being affected by export controls).

The transformer-based LLM architecture embeds self-attention mechanisms which enable a nuanced interpretation of subtle language and complex sentence structure in a document's broader context: these are crucial to performing our classification task. Distinguishing between threats, countries or products receiving them versus those imposing them, or characterizing how the firm is reacting along multiple possible dimensions is a type of task that until recently was not feasible to carry out at scale. The methods can also be implemented in near real time, with the two main constraints being the speed at which text becomes available and the computational power required to run the LLMs on the text. We demonstrate this feature by studying the reported effect in the texts of the ongoing tariff announcements and implementations around and following the US administration's April 2nd "Liberation Day" shock. We aim to update this analysis and the paper frequently.

With our measures of geoeconomic pressure in hand, we then present a range of new stylized facts on the nature of global economic pressure. As a first-pass validation, we document the aggregate time series of firms affected by tariffs, sanctions, and export controls. Reassuringly, these measures spike during well-known periods of geoeconomic pressure. We also demonstrate that the aggregate time series patterns of geoeconomic pressure display remarkable similarity across distinct text samples from corporate earnings calls and sell-side analyst reports, and across different LLM models. We map these aggregate patterns into a bilateral measure of pressure by instrument: who is pressuring whom using which policy instruments. We find that most pressure is applied worldwide by the two hegemons, the United States and China, mainly against each other. Yet, the aggregate patterns mask substantial heterogeneity.

In export controls, the US and China apply most of the pressure toward each other. The means are narrow, with a handful of sectors and firms being involved. The analysis captures the most prominent case of export controls of the US government targeting China largely via semiconductors-related export controls using both US firms (e.g. NVIDIA) and foreign firms (e.g. ASML or TSMC). The Chinese government targets the US largely using rare earth related products (metals, but also magnets). Overall, export controls in the data reflect a "small yard, high fence" approach to pressure: a few sectors and products with near total bans.[1]

Sanctions are mostly applied by the US and other Western countries (European countries, but also Japan and Australia) and have three main targets: Russia, Iran, and China. The sanctions on Russia occur in 2014 and 2022, corresponding to the two invasions of Ukraine. Sanctions are broad-based across sectors (both finance and trade) and include prohibitions to buy and to sell goods and services to Russian entities. Given Russia's composition of exports, natural gas and oil show up as the most prominent sectors. Sanctions on Iran are imposed by the US, especially in the period following the Trump administration's withdrawal from the Iran nuclear deal (JCPOA) in May 2018, and largely target oil and its transport. Sanctions on China are again imposed by the US, but in this case narrowly targeting a few entities, such as Huawei and ZTE (in the period 2019-2021), both on their acquisition and sale of telecommunication products.

[1] The phrase was popularized as the strategy for US export controls by US National Security Advisor Jake Sullivan. See, for instance, Brookings.

Tariffs differ markedly from the above instruments. A first episode is driven by the trade war between China and the US in 2018-19. The tariffs are largely reciprocally imposed and relatively wide in sectoral composition. The 2025 massive spike in tariffs is driven by the US government, broad based in both sectors and countries being targeted. We still find episodes of threats of tariffs toward specific countries and with specific aims (e.g. the February 2025 threat to impose US tariffs on imports from Colombia to induce that government to accept US military deportation planes), but the majority of the activity is broad-based.

We then turn to exploring whether governments are applying pressure via strategic sectors, in the sense of relying on those sectors whose cut-off would cause relatively more economic harm. We use our guiding framework to understand how the government imposing the policy chooses the means (firms-sectors) with which to apply the pressure. An export control, which is a restriction not to sell an input to a foreign entity, is more powerful the more it can reduce the profits of the target. These losses are generally greater when the target is very reliant on that input (high expenditure share) and has access to relatively poor substitutes (either domestic or foreign substitutes). Indeed, in recent years, there has been substantial focus on "chokepoints": parts of the economy (at either the sector or firm level) that are particularly vulnerable to economic pressure. A chokepoint in our framework is a bilateral relationship in which the target has a high expenditure share on the sender of pressure, and few alternatives (low elasticity of substitution). The same can be adapted for policies such as sanctions that often involve restrictions not to buy goods and services provided by the target. Using data on bilateral trade and estimates of elasticities of substitution, we show that the United States is indeed using sectors with these characteristics in both export controls and sanctions.

Having established how the pressure is being applied, we turn to examining its effect on the target, the means used by the sender, and third parties that, while not directly involved, are indirectly affected by the pressure. Starting with the US pressure applied on Chinese tech firms, via both semiconductor export controls and sanctions applied to specific firms, we find that the targets indeed report being negatively affected, which is to an extent the intended outcome of US policy. The Chinese firms, however, also report positive developments. First, they report rearranging their supply chains to source some of the chips from Korea (Samsung) and other East Asian countries (such as Vietnam and Malaysia). Second, they report increasing domestic investment and R&D specifically to produce domestic alternatives to the products they lost access to. At the same time, we document that the firms used as means to apply pressure by the US, e.g., NVIDIA and ASML, report being negatively affected by the export controls with loss of sales concentrated in the Chinese market.

In the case of sanctions on Russia, we find the firms based in Western countries are reacting negatively to loss of business (sales) but this is generally not a widespread concern, highlighting that Russia was not a large business destination for many Western firms, particularly American firms. However, we do find evidence of Western firms being concerned with negative effects of the sanctions in terms of losing access to Russian gas and oil. This was especially true during 2022.

Third parties in India and China often report being positively affected. Two forces drive this effect: firms that experience expanded business opportunities because a previous relationship between a Western and Russian firm was severed and they can step in, or firms that benefit from low oil and energy prices locally (e.g., Indian firms since Russia redirected its oil exports, at a discount, to that market).

Tariffs offer the largest cross-section of the paper since in 2025 over 60 percent of all firms report being affected by the tariffs. Restricting attention to the tariffs imposed and threatened by the U.S. government in 2025, we find evidence of both the consumer-tax and domestic-producer subsidy components of the tariffs, but we do not find evidence of the terms of trade moving in favor of the United States. The vast majority of US firms report being negatively affected by the tariffs. They report facing higher input prices and, as a consequence, increasing their sales prices. Indeed, some firms report not being negatively affected because they can pass through the cost of the tariffs to their US clients. There is a notable number of firms that report being positively affected by the tariffs: those firms with a (more) domestic production line compared to their competitors. We also document that US firms are planning to onshore some production or supplier choice to minimize the effect of the tariffs. Among foreign destinations, supply chains are moving toward US firms sourcing from Mexico and Vietnam and away from China, a rebalancing potentially due to differential tariff rates. Focusing on foreign firms, we do not find systematic evidence that these entities are (planning) to cut prices at which they sell to their products to the US, highlighting that the US does not see a terms of trade improvement.

Classification problems on large-scale text using a sophisticated understanding of the language are a new frontier of research in many fields, including economics. A methodological component of this paper is to work towards building best practices for the use of LLMs in economic analysis. In the last part of the paper, we discuss three types of concerns: data confidentiality and scientific reproducibility, computational costs, and measurement error. In particular, we quantify the degree of LLM measurement uncertainty by comparing the classifications across multiple open-weight models, as well as considering a wide range of variations of our prompts. While there is significant noise, it is generally in the range of well-regarded survey data. Comparing results across LLMs rather than within an LLM across different variations of a prompt generally results in a higher degree of uncertainty. Our takeaway is that these new research tools are here to stay and are expanding the frontier of what is possible in economic research, but have to be handled with care, especially in terms of measurement error.

Related Literature. Our paper relates methodologically to the literature on natural language processing (NLP) in economics, and in particular to the nascent literature using transformer-based large language model architectures. The economic focus of the paper relates to the literature on geoeconomics, building on international economics and political science.

NLP has become an integral tool of analysis in economics at least since the classic work of

Gentzkow and Shapiro (2010) and Baker et al. (2016).[2] In particular, we relate to the pioneering work on firm-level uncertainty using NLP by Hassan et al. (2019).[3] Caldara and Iacoviello (2022) measure discussions of risk of war, military buildups, and terrorism using daily newspaper data, and Caldara et al. (2020) focus on trade policy uncertainty. Juhász et al. (2022) measure industrial policy by undertaking natural language processing from text from Global Trade Alert. Goldberg et al. (2024) explore industrial policy in semiconductors. These papers largely use more traditional dictionary-based approaches in NLP. Our methodological approach is instead to leverage recent advances in artificial intelligence and in particular LLMs to extract more information out of the text. This nascent literature includes contributions by Chen et al. (2022), Ottonello et al. (2024), Bybee (2023), Sarkar (2025), Lagakos et al. (2025), and Fang et al. (2025).[4]

Second, we connect to the literature in economics and political science on sanctions and economic statecraft. One strand of the literature studies the effectiveness of and responses to economic sanctions. In an important contribution, Baldwin (1985) studies the tools of economic statecraft and challenges the idea that these instruments of foreign policy are not effective. The modern empirical literature largely attempts to understand the response to sanctions across various domains. This literature is primarily in political science and is surveyed in Drezner (2024). Drezner (2003) emphasizes that if sanctions are a tool of pressure, then sanctions should frequently be threatened and rarely imposed. Morgan et al. (2009) introduce the Threat and Imposition of Economic Sanctions (TIES) dataset, separately measuring realized sanctions and their threat from newspapers. Recent papers in economics include Ahn and Ludema (2020), Nigmatulina (2022), Keerati (2022), Itskhoki and Mukhin (2022), Fernandez-Villaverde et al. (2025), and Egorov et al. (2025). Felbermayr et al. (2020) assembled a bilateral database of sanctions around the world. A recent literature has focused on the idea of "weaponized interdependence" in Farrell and Newman (2019), Drezner et al. (2021), and Farrell and Newman (2023).

Lastly, we also connect to the fast-growing literature on geoeconomics, which includes Hirschman (1945), Kindleberger (1973), Keohane and Nye (1977), Blackwill and Harris (2016), Dreher et al. (2022), Clayton et al. (2023), Clayton et al. (2024), Thoenig (2023), Kleinman et al. (2024), Alekseev and Lin (2024), Becko and O'Connor (2024), Broner et al. (2024), Liu and Yang (2024), Kooi (2024), Mattoo et al. (2024), Mayer et al. (2025), and Pflueger and Yared (2024).[5]

[2] See also Gentzkow et al. (2019) for a review of the literature.

[3] For additional NLP at the firm-level see Hassan et al. (2024b), Hassan et al. (2024a), and Flynn and Sastry (2022).

[4] See Dell (2024) for a review of deep learning methods for economists.

[5] See also Fernández-Villaverde et al. (2024), Gopinath et al. (2025), Aiyar et al. (2024), Hakobyan et al. (2023), Aiyar et al. (2023), Flynn et al. (2025), Bonadio et al. (2024), Crosignani et al. (2024), Broner et al. (2025), and Clayton et al. (2025b). Mohr and Trebesch (2024) surveys the burgeoning literature.

2 Identifying Pressure Episodes

2.1 Textual Data

In order to measure geoeconomic pressure at the firm level, we need sources of firm-level text. Here, we use two primary sources: firm earnings calls and analyst reports.

Earnings Calls. Our first source of text is the earnings conference calls of publicly traded firms. As discussed in detail in Hassan et al. (2024b), earnings calls are a valuable source of textual information relative to regulatory filings such as 10-K forms because they allow market participants to ask firm leadership questions of their choosing. While the calls begin with a prepared presentation from management, this means that market participants have the opportunity to ask about the issues that they find most pressing. Of course, this in principle also generates the possibility that questions reflect analyst priorities rather than the primary concerns of the firm. However, given the prepared statement, this should not be a major issue.

One part of our dataset is the global earnings call transcripts from Capital IQ via Wharton Research Data Services (WRDS). These are the transcripts most often used in academic and industry research. The current dataset contains over 364 thousand transcripts for the period 2008-2025. However, as Table 1 shows, these transcripts are dominated by US firms. For example, in 2024 they contain transcripts from 3,333 US firms, 161 Chinese firms, and 5,481 firms in the entire Rest of the World. The lack of coverage in China is particularly problematic for our paper since China is likely to be both the sender and receiver of much economic pressure. We, therefore, introduce an additional dataset, the Orbit's China A-Shares Transcripts, that provides transcripts of earnings calls and investor relations meetings for Chinese companies listed in mainland China for the period 2008-2025 (approximately 220 thousand documents). As shown in Table 1, this dataset vastly expands our coverage of Chinese firms' text (e.g., 4,871 firms in 2024).

Earnings calls are a rich source of text that becomes available in almost real time. For our purposes there are three specific concerns. First, whether a firm enters the dataset might be endogenous to geoeconomic events: e.g., in the aftermath of Russia's invasion of Ukraine in 2022, far fewer Russian firms hold earnings calls that are covered in the dataset. Second, even conditional on having a conference call, some relevant event might not be covered. Typical concerns are events that, while big for the target country, may not attract attention on the earnings calls of major multinationals to discuss in its earnings calls. This issue is likely to be particularly important for smaller developing countries. At the other extreme, CFOs and CEOs might be reluctant to candidly discuss pressure being applied by their own governments. To address some of these concerns, we turn to analyst reports at both the firm and sector level as complementary source of text.

Analyst Research Reports. Our second source of firm and sector level textual data comes from the research reports from J.P. Morgan and Fitch Ratings (BMI). Our coverage starts in January 2011 for single-firm reports and 2017 for country-sector reports and includes approximately 200

Table 1: **Summary statistics of text corpora**

(a) Document Level

	Transcripts				Analyst Reports		
	WRDS			Orbit			
Year	USA	CHN	RoW	CHN	USA	CHN	RoW
2011	12630	239	4723	1264	4713	8	5049
2015	11780	302	5789	10482	4846	11	5094
2017	12033	274	7942	12008	5446	90	11223
2019	12219	469	11294	10808	6008	88	11221
2024	12500	491	16800	30089	5573	97	10573
Total	208134	5629	150257	219941	76968	800	123092

(b) Firm Level

	Transcripts				Analyst Reports		
	WRDS			Orbit			
Year	USA	CHN	RoW	CHN	USA	CHN	RoW
2011	3401	76	1598	1088	893	2	1245
2015	3131	92	1803	2200	971	4	1387
2017	3249	93	2894	2925	1025	20	2657
2019	3267	154	3633	3189	1131	18	2665
2024	3333	161	5481	4871	1174	28	2731
Total	58498	1990	52005	48200	15917	214	31726

Notes: The table focuses on counts of documents and unique entities by year, for selected years, across different text corpora. Panel (a) focuses on number of documents split by country and data source. Panel (b) focuses on number of firms, and for analyst reports exclude those at the country-sector level.

thousand individual research reports. While we also have industry-wide, country, or global macro reports, the focus of the present paper is on those reports that are about a single company or a single country-sector, as these can be more easily matched to our geoeconomics framework of sender-means-receiver and analyzed in parallel to the earnings calls. The sample of single-firm analysts reports, like the sample of earnings calls from WRDS, are heavily skewed towards large publicly traded firms from Western markets. The country-sector reports, therefore, help us to supplement our coverage: we rely on reports from BMI, a part of Fitch Ratings. These reports provide an overview of the prospects of a sector in a given country (such as "Argentina Banking" or "Canada Oil & Gas"), discuss the risks facing a given sector, review the strengths and weaknesses of a sector, discuss a forecast of future activity in a sector, and sometimes discuss the situation of the major companies in the sector. While analyzing these reports loses some of the granularity and detail available in earnings calls and single-firm analyst reports, they cover a much wider range of countries and industries. This extended coverage is particularly helpful, for instance, in measuring how Russian firms respond to Western sanctions.

2.2 Methodology

Our methodology performs large-scale inference on the two textual corpora (earnings calls and analyst reports) using frontier pre-trained, open-source large language models (LLMs). All the baseline results shown in the paper use the 70 billion parameter version of the Llama 3.3-Instruct model released by Meta, but we have also repeated the analysis using Google's 27 billion parameter Gemma 3 model, and Alibaba's 72 billion parameter Qwen 2.5-Instruct model. Section 6 discusses the robustness of the results to model choice. We structure the inference task using a two-part prompt design. First, we construct detailed system prompts with explicit instructions on the analysis to be performed. Second, the full text of each document is passed as the user prompt, so that every document is analyzed under the same standardized guidelines.

To ensure computational reproducibility of our analysis, we execute all inference locally using open-weight models (rather than through external APIs and closed-weight models) with controlled pseudo-randomization, and all data is stored on local hardware. The primary local cluster used for inference is based within the Stanford Sherlock high-performance computing system, and it contains eight NVIDIA A100-80GB GPUs and four NVIDIA H100-80GB GPUs. We have also used Stanford's Marlowe, a new GPU-based cluster (Kapfer et al. 2025). To minimize arbitrary stochastic variation in the LLM outputs across repeated runs, we set the sampling temperature to zero (greedy sampling). In the last section of the paper, to quantify the size of model-induced noise, we instead focus on analyzing variation induced by model and prompt perturbation, which has a natural scale and economic interpretation.

Finally, we implement frontier model quantization techniques. Specifically, we deploy 4-bit quantized versions of our models using the AWQ framework. This quantization dramatically reduces the memory footprint, allowing the entire model to fit on a single GPU without sacrificing performance, hence enabling a higher degree of parallelization during the inference procedure.

2.3 Prompt Structure

The computational challenge that we face is to analyze several hundred thousand earnings calls and analyst reports. The advantage of our approach is that it can extract detailed information about a firm's response, such as changes to its investment patterns. However, that comes with substantial computational requirements.[6] To reduce the computational burden, we implement a

[6]These costs are also changing quickly as both models and hardware leap forward. In our experience, for example, inference with the newer NVIDIA H100-80GB GPUs was significantly faster compared to the

two-stage procedure.

In the first step, all documents are processed via the LLM with a basic prompt. The prompt aims to classify whether the document discusses tariffs, export controls, sanctions, boycotts, investment screening, subsidies to onshore or friendshore, and a residual broader category for geoeconomic policy. The full text of our first-stage prompt for earnings calls is reproduced below. The first-stage prompt for analyst reports is analogous—with minimal changes to adapt the language to the different document type—and all prompts used in the paper are shown in their entirety in the appendix. As can be seen below, we prompt the LLM not to simply look for certain keywords or define a dictionary. We allow for the possibility that the words "tariffs," "sanctions," and "export controls" do not even appear in the text to decide whether firms are affected by these forces. Similarly, distinguishing some of these policies can be difficult for a machine either because of their underlying nature or because of nuances in the language. For example, a sanction imposed by the US government on firms that buy Russian oil is different from an export control on oil imposed by the Russian government on its domestic firms. The word tariff is also used by utility providers and telephone companies to mean "fee schedule."

The full text of the first-stage prompt is reported below:

```
You are assisting me in analyzing companies' earnings calls. The transcript of the
    earnings call will be supplied as the user prompt.

Your goal is to determine whether the company discusses impacts on its business due to
    various economic policies.  We define each in turn.

A. Tariffs are defined as taxes imposed on imported foreign goods. In order to be
    considered tariffs, these policies must be imposed by the importing country to be
    classified as tariffs.

B. Sanctions are government-imposed restrictions on trade or financial transactions
    designed to coerce, punish, or deter targeted firms or governments.

C. Export controls are defined as restrictions on which countries or foreign firms a
    company is allowed to sell their goods or services to. In order to be considered
    export controls, these policies must be imposed by the exporting countries, in
    contrast to import tariffs, which are instead taxes imposed by the importing country
    .

D. Boycotts are defined as a form of protest in which individuals or groups refuse to
    buy or use products or services of a firm as a way to express disapproval or to
    force change. It is typically viewed as voluntary and as a social method to
    influence business practices, government policies, or social issues but may be
    officially encouraged.

E. Investment Screening is a regulatory review carried out by government agencies to
    evaluate whether potential foreign investments meet certain criteria before allowing
     the investment to take place.  Authorities review foreign investments (including
    mergers and acquisitions) to ensure they do not pose risks to national security,
```

older A100-80GB GPUs. The newer Gemma 3 27 billion parameter model is similarly faster in inference compared to the Llama 3.3 70 billion parameter model.

critical infrastructure, or public interests. Investment Screening can be on inbound (foreign investment in the domestic economy) or outbound (domestic investment in a foreign economy) investments. In order to be considered investment screening, the policy needs to be carried out by the government. Do not consider private firms analyzing whether a particular investment is a good idea or profitable to be investment screening.

F. Subsidies to onshore or friendshore production are financial incentives provided by governments to encourage companies to manufacture goods domestically (onshoring) or in allied or politically aligned countries (friendshoring) rather than in less trusted or adversarial nations. These subsidies can take various forms, such as: direct grants or tax breaks for building or expanding production facilities, research and development support to boost innovation in key sectors, preferential loans or guarantees to reduce the cost of doing business locally or in partner countries, procurement preferences for products made in certain locations. Please do not include general subsidies unrelated to the goal of moving production to a particular country as part of this definition.

G. Geoeconomic pressure is the use of economic relationships by governments to achieve geopolitical or economic ends. For example, such pressure can take the form of tariffs, sanctions, boycotts, investment screening, or subsidies to onshore or friendshore.

Response Instructions: Part 1 (Structured JSON Output)

The first part of your response should be a structured output in JSON format that recaps your analysis in a structured way.
This part of your response should be enclosed between the tags <JSON> and </JSON>. The JSON output must have the following fields exactly.
Please make sure to enforce the JSON schema specified below strictly: i.e., the column names should correspond exactly to those listed below:

1. A boolean flag called "tariffs_any", which should be 1 if the firm explicitly or implicitly discusses tariffs at any point in the call, and 0 otherwise.

2. A boolean flag called "sanctions_any", which should be 1 if the firm explicitly or implicitly discusses sanctions at any point in the call, and 0 otherwise. Do not classify tariffs (taxes on imports) as sanctions.

3. A boolean flag called "export_controls_any", which should be 1 if the firm explicitly or implicitly discusses export controls at any point in the call, and 0 otherwise. Do not classify tariffs (taxes on imports) as export controls.

4. A boolean flag called "boycotts_any", which should be 1 if the firm explicitly or implicitly discusses boycotts at any point in the call, and 0 otherwise. Do not classify restrictions imposed by governments by law or regulations as boycotts.

5. A boolean flag called "investment_screening_any", which should be 1 if the firm explicitly or implicitly discusses investment screening at any point in the call, and 0 otherwise.

6. A boolean flag called "geo_subsidies_any", which should be 1 if the firm explicitly or implicitly discusses subsidies for onshoring or friendshoring at any point in the call, and 0 otherwise.

```
7. A boolean flag called "geoeconomic_any", which should be 1 if the firm explicitly or
    implicitly discusses any form of geoeconomic pressure (as defined above) during the
    call, and 0 otherwise. The term geoeconomic or geoeconomics are unlikely to be used
    explicitly during the call and any discussion is instead likely to involve the
    underlying policies such as tariffs, sanctions, export controls, boycotts,
    investment screening, and subsidies to onshore or friendshore. You should return a 1
     if the firm discusses impacts on its business that clearly relate to geoeconomic
    pressure.

## Response Instructions: Part 2 ##

Write a single summary of 750 words or fewer that captures only how the firm is (or may
    be) affected by the set of geoeconomic policies discussed in the analysis above,
    including but not limited to as tariffs, export controls, financial and trade
    sanctions, boycotts, investment screening, travel restrictions, subsidies to onshore
     or friendshore. Omit all unrelated content. If the firm is not affected by any of
    these policies, write "Not affected." Enclose this analysis between the tags <
    SUMMARY> and </SUMMARY>.


## Important Notes: ##

- Discussions may be explicit, such as directly mentioning "tariffs," "sanctions," or
    specific policy names.
- Discussions may be indirect, referring to impacts like inability to trade with certain
     countries, compliance with new export regulations, or financial losses due to
    exiting a sanctioned market.
- Do not consider general market volatility, economic downturns, or supply chain and
    travel or shipping issues unrelated to the above policies as indications of
    geoeconomic policies.

## Examples of Relevant Discussions: ##

- Mentioning increased costs due to new tariffs on imported materials.
- Discussing loss of a market due to trade sanctions against a country (e.g., exits from
     Russia due to sanctions following the Russia-Ukraine conflict).
- Referring to compliance challenges with new export controls.
- Mention the risk that an investment might not be approved due to failing the
    investment screening review.
- Receiving a subsidy to relocate production to the domestic economy or away from a
    country that is considered potentially hostile.

## Examples of Non-Relevant Discussions: ##

- Talking about decreased sales due to a general economic recession.
- Discussing delays caused by a natural disaster.
- Discussing purely domestic regulations (e.g., changes in US domestic tax laws or
    Federal Reserve policies).
- Mentioning fluctuations in currency exchange rates.

## Ensure you: ##

- Use the tags exactly as written.
- Choose a boolean flag (0 or 1) that is consistent with your analysis.
- Do not interpret general economic issues as geoeconomic policies unless directly
    linked.
```

The output of this first-stage prompt is a set of 7 boolean flags, each reporting whether a particular policy is being discussed in the specific document.[7] For each of the flagged documents, we run a more detailed second-stage prompt designed for that specific instrument. For this draft we focus exclusively on the analysis of tariffs, sanctions, and export controls. For example, the second-stage prompt for earnings calls flagged for tariffs is reported in its entirety below. Each version of the second-stage prompt is structured in three parts: (1) analysis, (2) producing a structured output in JSON format, and (3) an 100-word summary of the analysis.

```
You are assisting me in analyzing financial analyst reports. Each report is about a
    single company. The analyst report will be supplied as the user prompt.

## Response Instructions: Part 1 (Analysis) ##

The first part of your response should be an analysis of whether the analyst is
    reporting that the firm's decisions are being affected by tariff policies. Tariffs
    are defined as taxes imposed on imported foreign goods. These must be imposed by the
     importing country in order to be classified as tariffs. Tariffs are not export
    restrictions, quotas, embargoes, financial sanctions, boycotts, or non-tariff
    barriers.

This part of your response should be enclosed between the tags <ANALYSIS> and </ANALYSIS
    >. Keep your analysis to 300 words or less. Make sure to cover all of the following
    points in your summary analysis:

- Whether tariffs are discussed explicitly or implicitly (e.g., by not using the word ``
    tariffs'' but referencing impacts on the firm's business that clearly relate to
    tariffs).
- Whether the report discusses current tariffs (i.e., tariffs that have already been
    imposed) or the potential of future tariffs (i.e., tariffs that have not yet been
    imposed).
- Whether the report discusses tariffs on the goods that the firm sells, or on those
    that the firm buys.
- The nature and details of the tariff policies that are discussed.
- The countries that are imposing the tariffs and those that are subject to the tariffs.
- Any impacts on the firm's current profits.
- Any potential impact on the firm's  future profits.
- Any impacts on the firm's behavior (e.g., in terms of investment, pricing, employment,
     inventory, R&D, project delay, supply chains, or other future plans).
- Any details on the geographies affected by the changes in the firm's  behavior (e.g.,
    if the firm reports importing less, report from which destination).

## Response Instructions: Part 2 (Structured JSON Output) ##

The second part of your response should be a structured output in JSON format that
    recaps your analysis in a structured way. This part of your response should be
    enclosed between the tags <JSON> and </JSON>. The JSON output must have the
    following fields exactly. Please make sure to enforce the JSON schema specified
    below strictly: i.e., the column names should correspond exactly to those listed
    below:
```

[7] Our prompt requests also a 750 words short analysis. We use this unstructured output for manual inspection. We are also experimenting on running the second-stage prompts on this summary rather than the original document itself.

```
1. A boolean flag called "effect_any", which should be 1 if the report discusses tariffs
     at any point, and 0 otherwise. Even if the term "tariffs" is not explicitly used
    throughout the report, you should return a 1 if the report discusses impacts on the
    firm's  business that clearly relate to tariffs.
2. A boolean flag called "effect_current", which should be 1 if the report explicitly or
     implicitly discusses tariffs that  are currently in effect and that impact the firm
    's business, and 0 otherwise. Please do not classify tariffs that might imposed in
    the future as a 1 in this flag.
3. A boolean flag called "effect_future", which should be 1 if the report explicitly or
    implicitly discusses tariffs that  might be imposed in the future but are not
    currently in effect and that might impact the firm's business, and 0 otherwise.
    Please do not classify tariffs that are currently in effect as a 1 in this flag.
4. A field called "countries_imposing", listing the countries whose tariffs policy the
    report discusses, if any. For example, if the report discusses concerns about
    tariffs imposed by the US government on goods imported from China, this field should
     say "USA".
5. A field called "countries_receiving", listing the entities (countries) targeted by
    the tariffs discussed by the report, if any. For example, if the report discusses
    concerns about tariffs imposed by the US government on goods imported from China,
    this field should say "China".
6. A field called "product_receiving", listing the goods that the firm sells that are
    targeted by the tariffs, if any. For example, if the report discusses current or
    future tariffs imposed by the US government on steel that the firm sells to its US
    customers, this field should say  "steel".
7. A field called "input_receiving", listing the goods that the firm buys that are
    targeted by tariffs, if any. For example, if the report discusses current or future
    tariffs imposed by the US government on steel that the firm buys from Mexico, this
    field should say "steel".
8. A boolean flag called "investment_up", which should be 1 if the report discusses that
     the firm is increasing or planning to increase its investment as a result of
    current or future tariffs and 0 otherwise. Only set the flag to 1 if the report
    explicitly attributes a change in the firm's investment decisions to tariffs.
9. A boolean flag called "investment_down", which should be 1 if the report discusses
    that the firm is decreasing or planning to decrease its investment as a result of
    current or future tariffs and 0 otherwise. Only set the flag to 1 if the report
    explicitly attributes a change in the firm's investment decisions to tariffs.
10. A field called "countries_investment_up", listing the countries where the report
    discusses that the firm is increasing or planning to increase investment as a result
     of current or future tariffs. If no country is specified or it is unclear, report "
    NaN".
11. A field called "countries_investment_down", listing the countries where the report
    discusses that the firm is decreasing or planning to decrease investment as a result
     of current or future tariffs. If no country is specified or it is unclear, report "
    NaN".
12. A boolean flag called "domestic_investment_up", which should be 1 if the report
    discusses that the firm is increasing or planning to increase its investment in its
    headquarter country as a result of current or future tariffs and 0 otherwise. Only
    set the flag to 1 if the report explicitly attributes a change in the firm's
    investment decisions to tariffs.
13. A boolean flag called "domestic_investment_down", which should be 1 if the report
    discusses that the firm is decreasing or planning to decrease its investment in its
    headquarter country as a result of current or future tariffs and 0 otherwise. Only
    set the flag to 1 if the report explicitly attributes a change in the firm's
    investment decisions to tariffs.
14. A boolean flag called "labor_up", which should be 1 if the report discusses that the
```

```
 firm is increasing or planning to increase its number of workers or their hours
worked as a result of current or future tariffs and 0 otherwise. Only set the flag
to 1 if the report explicitly attributes a change in the firm's employment decisions
 to tariffs.
15. A boolean flag called "labor_down", which should be 1 if the report discusses that
the firm is decreasing or planning to decrease its number of workers or their hours
worked as a result of current or future tariffs and 0 otherwise. Only set the flag
to 1 if the report explicitly attributes a change in the firm's employment decisions
 to tariffs.
16. A field called "countries_labor_up", listing the countries where the report
discusses that the firm is increasing or planning to increase its number of workers
or their hours worked as a result of current or future tariffs. If no country is
specified or it is unclear, report "NaN".
17. A field called "countries_labor_down", listing the countries where the report
discusses that the firm is decreasing or planning to decrease its number of workers
or their hours worked as a result of current or future tariffs. If no country is
specified or it is unclear, report "NaN".
18. A boolean flag called "domestic_labor_up", which should be 1 if the report discusses
 that the firm is increasing or planning to increase its number of workers or their
hours worked in its headquarter country as a result of current or future tariffs and
 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in
 the firm's employment decisions to tariffs.
19. A boolean flag called "domestic_labor_down", which should be 1 if the report
discusses that the firm is decreasing or planning to decrease its number of workers
or their hours worked in its headquarter country as a result of current or future
tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes
a change in the firm's employment decisions to tariffs.
20. A boolean flag called "sales_price_up", which should be 1 if the report discusses
that the firm is increasing or planning to increase the price it charges for the
products it sells as a result of current or future tariffs and 0 otherwise. Only set
 the flag to 1 if the report explicitly attributes a change in the firm's product
pricing decisions to tariffs.
21. A boolean flag called "sales_price_down", which should be 1 if the report discusses
that the firm is decreasing or planning to decrease the price it charges for the
products it sells as a result of current or future tariffs and 0 otherwise. Only set
 the flag to 1 if the report explicitly attributes a change in the firm's product
pricing decisions to tariffs.
22. A field called "countries_sales_price_up", listing the countries where the report
discusses that the firm is increasing or planning to increase the price it charges
for the products it sells as a result of current or future tariffs. If no country is
 specified or it is unclear, report "NaN".
23. A field called "countries_sales_price_down", listing the countries where the report
discusses that the firm is decreasing or planning to decrease the price it charges
for the products it sells as a result of current or future tariffs. If no country is
 specified or it is unclear, report "NaN".
24. A boolean flag called "input_price_up", which should be 1 if the firm says it faces
or expects to face higher prices on the inputs it purchases as a result of current
or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly
attributes a change in the firm's input prices to tariffs.
25. A boolean flag called "input_price_down", which should be 1 if the firm says it
faces or expects to face lower prices on the inputs it purchases as a result of
current or future tariffs and 0 otherwise. Only set the flag to 1 if the report
explicitly attributes a change in the firm's input prices to tariffs.
26. A field called "countries_input_price_up", listing the countries where the report
discusses that the firm is facing or expecting to face higher prices on the inputs
it purchases as a result of current or future tariffs. If no country is specified or
```

```
    it is unclear, report "NaN".
27. A field called "countries_input_price_down", listing the countries where the report
    discusses that the firm is facing or expecting to face lower prices on the inputs it
     purchases as a result of current or future tariffs. If no country is specified or
    it is unclear, report "NaN".
28. A boolean flag called "sales_up", which should be 1 if the report discusses that the
     firm is increasing or planning to increase the quantities of the products that it
    sells as a result of current or future tariffs and 0 otherwise. Only set the flag to
     1 if the report explicitly attributes a change in the firm's product sales to
    tariffs.
29. A boolean flag called "sales_down", which should be 1 if the report discusses that
    the firm is decreasing or planning to decrease the quantities of the products that
    it sells as a result of current or future tariffs and 0 otherwise. Only set the flag
     to 1 if the report explicitly attributes a change in the firm's product sales to
    tariffs.
30. A field called "countries_sales_up", listing the countries where the report
    discusses that the firm is increasing or planning to increase the quantities of the
    products that it sells as a result of current or future tariffs. If no country is
    specified or it is unclear, report "NaN".
31. A field called "countries_sales_down", listing the countries where the report
    discusses that the firm is decreasing or planning to decrease the quantities of the
    products that it sells as a result of current or future tariffs. If no country is
    specified or it is unclear, report "NaN".
32. A boolean flag called "input_up", which should be 1 if the report discusses that the
     firm is  increasing or planning to increase the quantities of the inputs that it is
     purchasing as a result of current or future tariffs and 0 otherwise. Only set the
    flag to 1 if the report explicitly attributes a change in the firm's input
    quantities  to tariffs.
33. A boolean flag called "input_down", which should be 1 if the firm says it  is
    decreasing or planning to decrease the quantities of the inputs that it is
    purchasing as a result of current or future tariffs and 0 otherwise. Only set the
    flag to 1 if the report explicitly attributes a change in the firm's input
    quantities to tariffs.
34. A field called "countries_input_up", listing the countries where the report
    discusses that the firm is increasing or planning to increase the quantities of the
    inputs that it is purchasing  as a result of current or future tariffs. If no
    country is specified or it is unclear, report "NaN".
35. A field called "countries_input_down", listing the countries where the report
    discusses that the firm is decreasing or planning to decrease the quantities of the
    products as a result of current or future tariffs. If no country is specified or it
    is unclear, report "NaN".
36. A boolean flag called "inventory_product_up", which should be 1 if the report
    discusses that the firm is increasing or expecting to increase the amount of
    inventory of the products it sells as a result of current or future tariffs and 0
    otherwise. Only set the flag to 1 if the report explicitly attributes a change in
    the firm's inventory decisions to tariffs.
37. A boolean flag called "inventory_product_down", which should be 1 if the report
    discusses that the firm is decreasing or expecting to decrease the amount of
    inventory of the products it sells as a result of current or future tariffs and 0
    otherwise. Only set the flag to 1 if the report explicitly attributes a change in
    the firm's inventory decisions to tariffs.
38. A boolean flag called "inventory_input_up", which should be 1 if the report
    discusses that the firm is increasing or expecting to increase the amount of
    inventory of inputs it purchases as a result of current or future tariffs and 0
    otherwise. Only set the flag to 1 if the report explicitly attributes a change in
    the firm's inventory decisions to tariffs.
```

39. A boolean flag called "inventory_input_down", which should be 1 if the report discusses that the firm is decreasing or expecting to decrease the amount of inventory of inputs it purchases as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's inventory decisions to tariffs.
40. A boolean flag called "rd_up", which should be 1 if the report discusses that the firm is increasing or expecting to increase the amount of research and development it undertakes as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's research and development decisions to tariffs.
41. A boolean flag called "rd_down", which should be 1 if the report discusses that the firm is decreasing or expecting to decrease the amount of research and development it undertakes as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's research and development decisions to tariffs.
42. A boolean flag called "delay", which should be 1 if the report discusses that the firm is delaying a decision as a result of current or future tariffs. Only set the flag to 1 if the report explicitly attributes the delay to tariffs.
43. A boolean flag called "supply_chain_adj", which should be 1 if the report discusses that the firm is adjusting its supply chain as a result of current or future tariffs . Only set the flag to 1 if the report explicitly attributes adjusting the firm's supply chains due to tariffs.
44. A field called "countries_supply_chain_adj_towards", listing the countries where the report discusses that the firm is increasing the exposure of its supply chain as a result of current or future tariffs. If no country is specified or it is unclear, report "NaN".
45. A field called "countries_supply_chain_adj_away", listing the countries where the report discusses that the firm is reducing the exposure of its supply chain as a result of current or future tariffs. If no country is specified or it is unclear, report "NaN".
46. A boolean flag called "lobbying_up", which should be 1 if the report discusses that the firm is increasing or expecting to increase its lobbying activity as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's lobbying policy to tariffs.
47. A boolean flag called "lobbying_down", which should be 1 if the report discusses that the firm is decreasing or expecting to decrease its lobbying activity as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's lobbying policy to tariffs.
48. A boolean flag called "profit_margin_up", which should be 1 if the report discusses that the firm's profit margins have increased or are expected to increase as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's profit margins to tariffs.
49. A boolean flag called "profit_margin_down", which should be 1 if the report discusses that the firm's profit margins have decreased or are expected to decreased as a result of current or future tariffs and 0 otherwise. Only set the flag to 1 if the report explicitly attributes a change in the firm's profit margins to tariffs.
50. A boolean flag called "financing_adj", which should be 1 if the report discusses that the firm is adjusting its financing decisions in response to current or future tariffs. Only set the flag to 1 if the report explicitly attributes the adjustment of the firm's financing decisions to tariffs.
51. A boolean flag called "negative_impact", which should be 1 if the report discusses any negative impact on the firm's business as a result of current or future tariffs and 0 otherwise.
52. A boolean flag called "positive_impact", which should be 1 if the report discusses any positive impact on the firm's business as a result of current or future tariffs and 0 otherwise.

```
## Response Instructions: Part 3 (Summary) ##
Write a single summary of 100 words or fewer that captures only how the firm is (or may
    be) affected by tariffs. Omit all unrelated content. If the firm is not affected by
    any tariffs, write "Not affected." Enclose this analysis between the tags <SUMMARY>
    and </SUMMARY>.

## Response Instructions: Part 4 (Evaluation) ##
The fourth part of your response should be an evaluation of how well the JSON structured
     summary agrees with your initial analysis. Keep the evaluation to 100 words or less
    . This part of your response should be enclosed between the tags <EVAL> and </EVAL>
    and be 100 words or less.

## Important Notes ##
- Do not consider restrictions on exports (imposed by the seller country) or generic
    sanctions as indications of tariffs. Tariffs must be taxes imposed by the importing
    country.
```

The outcome of these detailed second-stage prompts is a structured dataset of boolean flags, code lists, and keywords. In addition, the LLM produces a 100-word summary of the entire analysis that is useful for both inspection and further analysis as we discuss below. The resulting structured datasets are the core of the empirical analysis of geoeconomics in this paper. After this step, the empirical analysis proceeds using methods familiar to economists with structured (potentially measured with error) data. Conceptually, our geoeconomics framework guides the structure of the prompts (what we are asking the LLM to classify, and how) and the selection of appropriate text to classify. The LLM is then used as a classifier that reduces a large amount of unstructured data (text) to a low-dimensional structured dataset.

3 Who Pressures Who Using What

3.1 An Organizing Framework for Economic Pressure

Before turning to a systematic empirical analysis, we begin by introducing a basic organizing framework and the terminology that we then apply throughout the analysis. To make it concrete, we run through the example in the actual data of economic pressure applied by the US government toward China using export controls on semiconductors via both US and European firms. This also helps us to illustrate how the LLM classifies the text.

We start by considering a firm that takes as given a set of aggregate variables Z (e.g., prices, aggregate demand, etc.) and chooses privately how to operate. The value of the firm is $V(x^*, Z, \theta, \tau)$, where $x^* = x(Z, \theta, \tau)$ are the privately optimal choices of the firm given Z and (θ, τ) are the geoeconomic instruments, described below. The firm in question could be NVIDIA or a China-based firm operating in the semiconductor industry, such as Semiconductor Manufacturing International Corporation (SMIC). In the absence of geoeconomic pressure, $(\theta, \tau) = (0, 0)$, the firm's private optimum is $x(Z, 0, 0)$ and its value is $V(x^*, Z, 0, 0)$. We think of pressure as a change in government policy that potentially affects the firm's value and choices. We represent the application of pressure

Figure 1: **Representing pressure on firms**

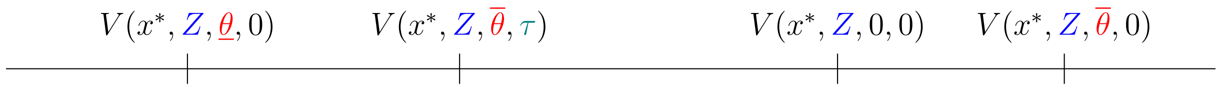

as a vector θ that enters the firm problem. For example, the US government imposing export controls on Chinese firms to limit access to high-end chips changes the value of the targeted firm to $V(x^*, Z, \underline{\theta}, 0) \leq V(x^*, Z, 0, 0)$, where the optimal choices $x^*(Z, \underline{\theta}, 0)$ for a subset of economic activity are influenced by the government imposing the pressure, in this case the unavailability of certain inputs. However, pressure does not have to be negative: a government might for example decide to subsidize a firm under certain conditions. We represent positive inducements by $\bar{\theta}$, so that we have $V(x^*, Z, \bar{\theta}, 0) \geq V(x^*, Z, 0, 0)$.

Firms might be uncertain about the policies θ that they might face in the future, or the policies could be conditional on the firms taking specific actions that the imposing government desires. For example, a government might pressure a foreign firm by threatening a low $\underline{\theta}$ unless the firm agrees to modify its behavior, that is to change some or all of the values of the vector x that the firm sets. We represent these changes in behavior by a vector of costly actions τ that the government imposes on the firm such that the firm's action is now $x^*(Z, \overline{\theta}, \tau)$ and its value is $V(x^*, Z, \overline{\theta}, \tau)$. The vector τ could contain quantity restrictions, such as sanctions and export controls, but also ad-valorem taxes and subsidies, such as tariffs (taxes on imports) or production subsidies. Figure 1 tracks these instruments and the notation. As in Clayton et al. (2023) and Clayton et al. (2024), our framework in this paper focuses on a participation constraint and off-path and on-path pressure applied by governments to domestic and foreign firms. Ultimately, firms comply with government demands when they would otherwise be worse off, that is $V(x^*, Z, \overline{\theta}, \tau) \geq V(x^*, Z, \underline{\theta}, 0)$.

Let us consider as a concrete example the recent pressure applied by the US government on the semiconductor industry. In this case, the US government is the sender of the pressure and the target are firms in the semiconductor industry in China. The instrument of the pressure is export controls and the means used are both US domestic firms and some firms in Europe and in other American allied countries. Starting from the domestic side, the US government imposed strict export controls toward China on advanced chips produced by NVIDIA. Since NVIDIA is an American firm, the US government had the legal means to directly dictate part of the firm behavior in the case of national security concerns, in this case via the Commerce Department's Bureau of Industry and Security (BIS). We think of $\tau_{USA,NVIDIA}$ as the set of export controls that on-path the US dictated to NVIDIA, and $\underline{\theta}_{USA,NVIDIA}$ as the threat of legal consequences for NVIDIA had it not complied with US export controls. Governments have extensive threats and legal power over their domestic firms, and NVIDIA complied with the government orders. We think of the vector $\tau_{USA,NVIDIA}$ as including an element for NVIDIA's export of advanced chips to China for which the tax rate is set to infinity (a full ban). However, NVIDIA actively worked to develop new products to satisfy Chinese demand while remaining in legal compliance, also demonstrating the limits of what a government can ask of its own firms, as captured in the participation constraint. In our framework, NVIDIA's optimal actions, x^*, were both to comply with the τ and to minimize the impact on its business. We show below both excerpts of the raw corporate text by NVIDIA and the LLM classifications.

In its earnings call on November 16, 2022, NVIDIA focuses extensively on the new export controls announced by the United States on cutting edge semiconductor sales to China. NVIDIA's Chief Financial Officer Colette Kress begins by summarizing the situation:

> "During the quarter, the U.S. government announced new restrictions impacting exports of our A100 and H-100 based products to China, and any product destined for certain systems or entities in China. These restrictions impacted third quarter revenue, largely offset by sales of alternative products into China. That said, demand in China more broadly remains soft, and we expect that to continue in the current quarter. We started shipping our flagship H-100 data center GPU based on the new Hopper Architecture in Q3. H-100-based systems are available starting this month from leading server makers including Dell, Hewlett Packard Enterprise, Lenovo and Supermicro."

To begin, our first-stage prompt correctly flags this earnings call as discussing export controls. Our second-stage prompt summarizes its analysis in 100 words or less as:

> "The firm is affected by U.S. export controls on their A100 and H-100 products to China, which impacted their third-quarter revenue. However, they were able to offset some of this impact by selling alternative products."

The LLM successfully classifies the firm as being affected by current export controls, "USA" as the country imposing the export controls, "China" as the country receiving export controls, and the product or input being affected as "A100 and H-100 based products". The LLM, therefore, has successfully classified the sender-means-receiver of the pressure in this particular document. Turning to how NVIDIA responded to the export controls, it reports that NVIDIA responded by adjusting supply chains away from China, and that the overall impact on the firm was negative. The LLM also picks up that the sales were not affected as the negative impact of the export control is offset by the sale of other products to China. In the context of our theory τ is imposing export controls on NVIDIA-specific products to China, as a result $V(x, Z, , ., \tau)$ for NVIDIA is lower than $V(x, Z, ., 0)$ since the export controls are imposing costs on NVIDIA.

We turn next to the US government pressuring ASML Holdings NV, a Dutch firm, into not selling or servicing advanced lithography machines to Chinese firms. Since ASML is a foreign firm, the US government had more limited ability to apply pressure compared to NVIDIA to induce ASML to comply. In such cases, governments generally opt for two approaches: either they apply pressure directly to the firm itself, or apply pressure to the firm's government that can in turn use its domestic regulatory and taxation powers on the firm. In the case of ASML the US government seems to have pursued both avenues. At the firm level, the threat $\underline{\theta}_{USA,ASML}$ was allegedly to impose the Foreign Direct Product Rule (FDPR). This rule would have allowed the US government to directly regulate ASML as long as it used technology and inputs produced by US firms as part of its business.[8] It is harder to assess the threats and inducements that the US government applied on the Dutch government since those are often part of confidential diplomatic negotiations. Ultimately, ASML faced export controls imposed by either the US or the Netherlands towards Chinese mainland producers of semiconductors such as SMIC and Yangtze Memory Technologies.

In our data, ASML is repeatedly flagged for export controls by our first-stage prompt in 25 earnings calls between January 2021 and January 2025. The LLM correctly captures the countries imposing the controls as being either the United States or the Netherlands, and the target of the pressure as being China. The products involved are classified correctly by the LLM as "EUV systems", "lithography tools", and "EUV and DUV systems". The main reactions of the firm are classified as lower sales, the country of lower sales being China, and ASML as being overall negatively affected.

The ultimate targets of the US pressure were Chinese firms that either aimed to purchase the finished chips or aimed to develop their own ability to produce them. For example, Semiconductor Manufacturing International Corporation (SMIC) is a Chinese semiconductor foundry company that was targeted by the American (and Dutch) controls. In this case the threat $\underline{\theta}_{USA,SMIC}$ was realized since the ask would have amounted to this Chinese firm either not developing its own production capabilities or sufficiently severing its relationships with the Chinese government and other state-owned enterprises to convince the US government that the firm posed no national security threat to the US. We focus therefore on how SMIC reacted to the threat being realized, the term $V(x, Z, \underline{\theta}, 0)$.

In our data, SMIC reports being negatively affected by the US-imposed export controls. For example, in an earnings call in November 2022, SMIC mentions that:

> "Due to the weak demand in mobile phone and consumer, overlapping with the impact from that some customers need time to interpret the newly released U.S. export control rules revenue is expected to decline by 13% to 15% sequentially with gross margin in the range of 30% to 32%."

and in an earlier call in February 2022, it mentions that:

> "In the past year, the company has made persistent efforts to meet customer demand and alleviate shortages in the supply chain in the face of the huge challenges posed by the entity list. [...] Internally, we reach consensus on development with the Board of Directors, quickly launched the construction of new fabs and capacity, streamlined management processes and advanced marketing, selling, planning, procurement and operations side-by-side."

The LLM correctly captures SMIC as being negatively affected by US-imposed export controls on China and reports domestic sales going down as well as the firm being negatively affected. It also reports, in the case of the February 2022 call that SMIC is increasing domestic investment and R&D in response to the export controls.

[8] See for example press reports about the US threatening FDPR by Bloomberg.

Figure 2: **Geoeconomic pressure: aggregate trends**

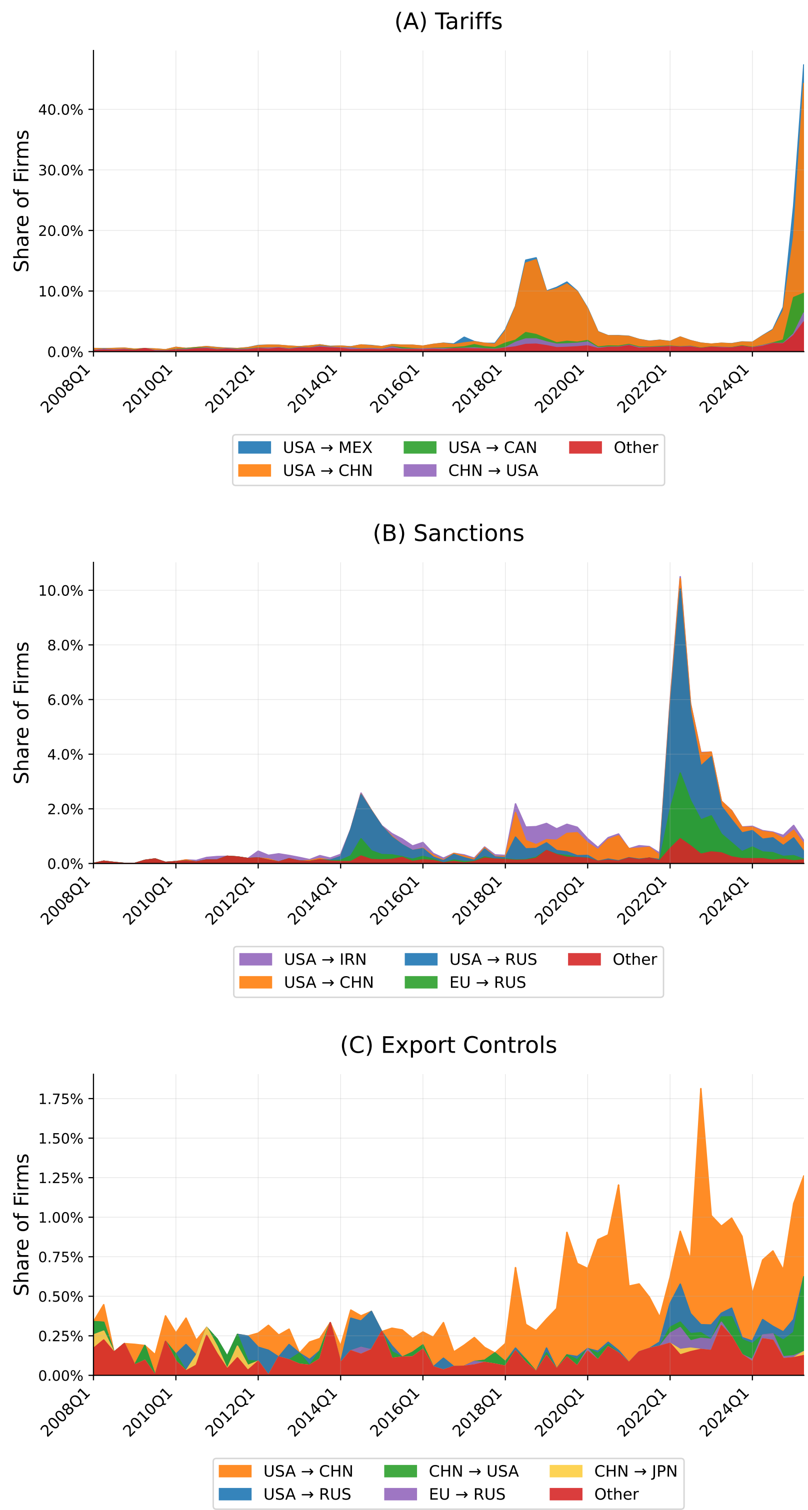

Notes: Figures plot the share of firms in earnings calls reporting that they are affected by each of the policy instruments (tariffs in Panel A, sanctions in Panel B, export controls in Panel C) in a given quarter the sender and receiver of each form of pressure. The notation $X \rightarrow Y$ denotes that country X is the sender of pressure and country Y is the receiver of the pressure. Observations with missing sender or receiver are imputed from the distribution of the observed data.

3.2 Aggregate Trends in Geoeconomic Pressure: From Whom to Whom

We begin our empirical analysis by exploring the aggregate time series and showing that they conform to well-known events. In particular, we begin by exploring quantitative patterns in the boolean fields indicating whether the firms are affected at all by export controls, financial sanctions, or tariffs. These boolean fields are generated by the first question in the second-stage prompt.[9] We also show a decomposition of who is imposing the pressure on whom for each instrument. For example, we decompose the time series for the tariffs by which country is imposing the tariff on which other countries. To do so, we use the output of questions 4 and 5 of the second-stage prompt.[10]

In Figure 2, we report the share of firms discussing being affected by these three instruments of geoeconomic pressure in their earnings calls. There are a number of key findings that emerge from this analysis. First, in the aggregate, we observe spikes for each of the three tools around well-known episodes when geoeconomic pressure was applied or anticipated.

The spike in tariff discussion in the first and second quarters of 2025 is the largest event in our sample. We note that the second quarter is still ongoing at the time of writing and that we will continue adding data as it becomes available. Prior to 2025, the largest spike in tariffs is during the US-China Trade War of 2018-2019. Looking at the combinations of countries imposing andreceiving the tariffs, the first event in 2018-2019 is quantitatively driven by US tariffs imposed on China, while the second in 2025 is driven by the US (threatening to) imposing tariffs on China and a broad set of other countries. Although we do see an increase in firms affected by China's tariffs on the United States in 2018-2019 and in the ongoing trade war, the increase is small compared to the spike in firms concerned about US tariffs on China.

For sanctions, two notable spikes occur in 2014 and 2022. Both are related to sanctions imposed on Russia by Western countries, most prominently the US. These sanctions occurred in response to Russia's invasion of Crimea in 2014 and broadening of the invasion of Ukraine in 2022. Three other bilateral pairs stand out. First, US sanctions targeting China in the period 2019-2021: these are mostly sanctions applied on specific companies in China such as Huawei and ZTE. Second, we observe a rise in firms affected by US sanctions targeting Iran and in particular its oil sector starting in the second quarter of 2018 in the wake of the Trump Administration's withdrawal from the "Iran nuclear deal" (the Joint Comprehensive Plan of Action, JCPOA) in May 2018.

For export controls, the patterns are more diffuse over time but with a notable time trend in

[9] For example, for tariffs the question in the prompt is: "1. A boolean flag called "effect_any", which should be 1 if the report discusses tariffs at any point, and 0 otherwise. Even if the term "tariffs" is not explicitly used throughout the report, you should return a 1 if the report discusses impacts on the firm's business that clearly relate to tariffs."

[10] For example, for tariffs the questions in the prompt are: "4. A field called "countries_imposing", listing the countries whose tariffs policy the report discusses, if any. For example, if the report discusses concerns about tariffs imposed by the US government on goods imported from China, this field should say "USA"."; "5. A field called "countries_receiving", listing the entities (countries) targeted by the tariffs discussed by the report, if any. For example, if the report discusses concerns about tariffs imposed by the US government on goods imported from China, this field should say "China".

the fraction of firms affected. Most export controls are imposed by US and China targeting each other. The Euro Area countries mostly impose export controls toward Russia, probably as part of sanction packages.

One concern with earnings calls and politically sensitive topics such as geoeconomic pressure is that CEOs, CFOs, and their investor relations teams are reluctant to be candid about the pressure in fear of upsetting relations with their own government or with that of their major customers or suppliers. This could particularly affect the composition of who applies the pressure to whom. In Section 6, we show that the results in Figure 2 are overall consistent with changing the text to be single-firm analyst reports. Professional analysts, being external to the firm and having career incentives to accurately analyze the available information, reduce concerns of misreporting.

3.3 The Means of Geoeconomic Pressure: Sectors and Countries

Having established who is applying the pressure on whom using each instrument, we turn next to identifying the means used to apply the pressure. We focus on sectors and products involved in each episode.

We map each firm that is associated with a document in our dataset into a primary sector (at the SIC4 code level) using the Factset Symbology package.[11] Geoeconomic pressure, however, is often applied using very specific products that might be hard for the target to buy or sell elsewhere. For example, a firm might be a mining conglomerate, but it is only the rare earths that are being used as means of pressure.

To narrow down these finer means of pressure, we use the LLM unstructured answers to questions 6 and 7 of the second-stage prompts that specifically identify the items affected by the policy instrument either on the product (sales) or input (purchases) side.[12] The output of the LLM for these fields is unstructured text, the answers for example might be words such as: semiconductors, H100 GPUs, rare earths, steel, lasers, etc. In order to map these unstructured fields into sectoral codes, we build an algorithm to map these words into SIC4 sector codes.[13]

[11]For WRDS earnings calls transcripts, we start from the Capital IQ company id, a unique identifier of each firm associated with each transcript provided in the raw data. For both ORBIT earnings calls transcripts and for single-firm analyst reports, we start from equity Tickers and market of listing information associated with each document and map them to companyid using Capital IQ table wrds_ticker. We then map each of these identifiers into Factset entity ids, a unique identifier of each entity across all Factset datasets that we use. We then use Factset to obtain information such as the country of registration of the firm, the SIC code of its primary industry, etc.

[12]For example, for tariffs the question in the prompt is: "6. A field called "product_receiving", listing the goods that the firm sells that are targeted by the tariffs, if any. For example, if the report discusses current or future tariffs imposed by the US government on steel that the firm sells to its US customers, this field should say "steel"."; "7. A field called "input_receiving", listing the goods that the firm buys that are targeted by tariffs, if any. For example, if the report discusses current or future tariffs imposed by the US government on steel that the firm buys from Mexico, this field should say "steel"."

[13]We use a pretrained sentence-embedding model 'bge-large-en-v1.5' created by the Beijing Academy of Artificial Intelligence (BAAI) to compare the textual similarity between each entry of "input_receiving' and "product_receiving" with official SIC4 description. We accept assignments with sufficiently high similarity, and manually review those with similarity below the threshold.

Figure 3 is a Sankey chart of who applies pressure to whom using which means grouped by the most frequently used sectors.[14] For export controls, very specific products-inputs are being used, but which products are used differ substantially based on who is the sender and receiver of the pressure. When the US pressures China using export controls, they concentrate in the semiconductor industry and its supply chain. When China pressures the US some export controls are also in semiconductors, but the bulk focus on rare earth and products downstream of these materials (such as magnets and batteries).[15]

For sanctions, the means being used are wider than for export controls but still involve fairly specific sectors. For example, in the case of the US sanctioning Iran or specific entities in Russia, the means of applying the pressure are concentrated in the oil industry (and its transport) and in telecommunications equipment, respectively. Sanctions applied by the US and the EU on Russia use broader means. Oil and natural gas (and their transportation) all play an important role, but financial sanctions on payments and other financial services (like insurance and brokerage) also play a role. Financial sanctions are also often indirectly reported as the background reason why acquiring a specific input is impaired.

For tariffs, the patterns are substantially different from the other two instruments. The US government, during both Trump administrations, has imposed or threatened to impose tariffs on a very broad set of products and countries.

The picture that emerges from this analysis of geoeconomic pressure points to a few stylized facts. First and foremost, the pressure (at least the one firms talk about) is largely applied by hegemons such as the US and China. Most of the pressure is targeted toward each other. Export controls and sanctions are applied using very specific sectors and target countries, a "small yard, high fence" approach to geoeconomics. Tariffs, especially in the recent case of the US, are instead applied broadly both in terms of the countries receiving the pressure and the sectors being used, or more of a "whole yard, massive fence."

3.4 Chokepoints: Which Firms and Products Are Used to Apply Targeted Pressure?

In the case of targeted pressure, the previous subsection documented a distinct pattern of which sectors and products are used to apply pressure in each bilateral relationship. Here we build on the geoeconomics framework by Clayton et al. (2023, 2024, 2025a) to systematically analyze whether countries are using the products that those theories predict would generate the most pressure in a given bilateral relationship. In particular, Clayton et al. (2024) show that the pressure that country

[14] The sectors are assigned starting from the product and input receiving mapped to SIC4. For visualization purposes, we focus on SIC2 level code. In Panel (b) and (c) we further aggregate some sectors to more easily visualize the underlying effects. In particular we group sectors: 32, 35, 36, and 38 into a miscellaneous "Semiconductors Supply Chain"; 10, 14, 33, and 34 into "Metals and Minerals"; 13 and 29 in "Oil and Gas"; 11 and 24 in "Agriculture and Forestry".

[15] See Alfaro et al. (2025) for a study of directed technical change in rare earths in response to shortages and export controls.

Figure 3: **Characterizing pressure: from who to whom using what**

(a) Export Controls

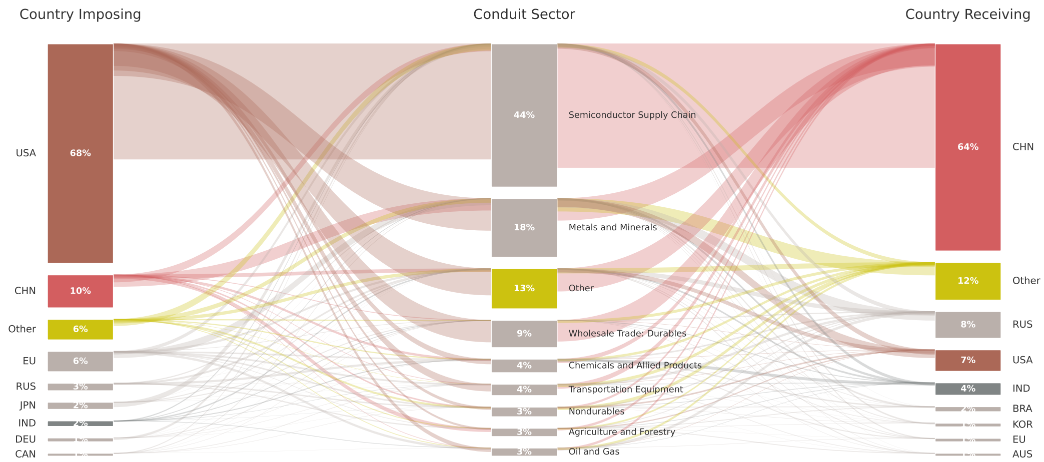

(b) Sanctions

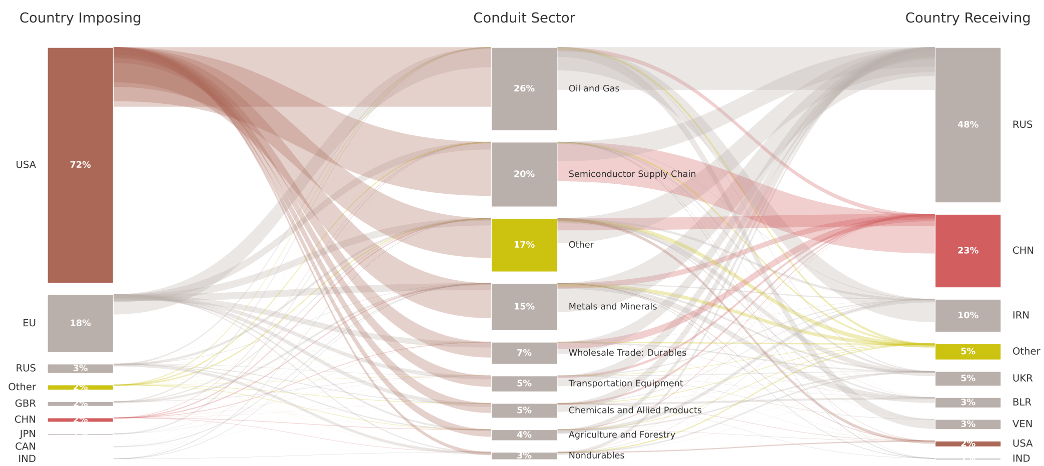

(c) Tariffs

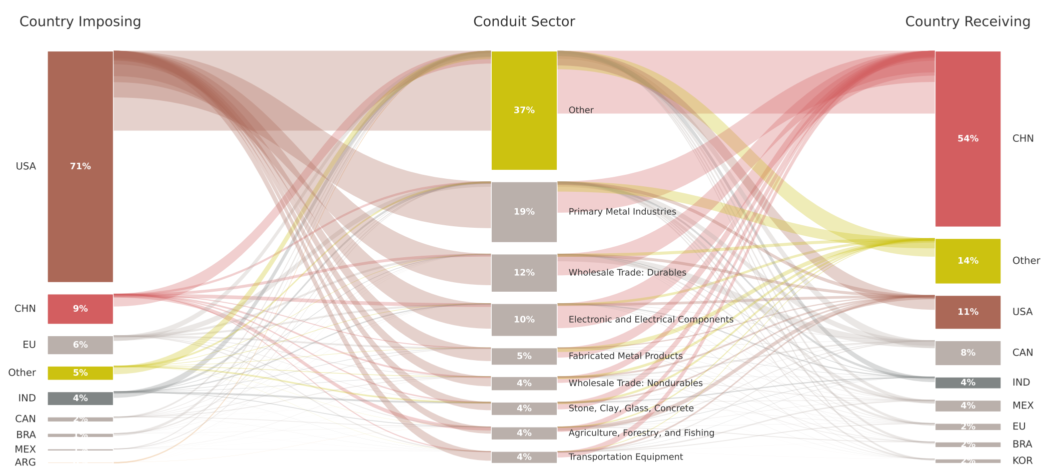

Notes: Figure displays the share of tariffs, sanctions, and export controls imposed by each sender country on receiver countries through inputs in a given SIC2 code (or aggregate category).

m can apply to country n via restrictions not to sell a set of inputs it controls takes the form:

$$\text{Power}_{mn} = \frac{\beta}{1-\beta}\frac{1}{1-\varrho}\log\left(\sum_{G\in\mathcal{G}}\Omega_{nG}\left(\sum_{J\in\mathcal{J}_G}\Omega_{nGJ}\left(1-\Omega_{nJR}+\Omega_{nJR}\left(1-\omega_{nJR_m}\right)^{\frac{\varsigma_J-1}{\sigma_J-1}}\right)^{\frac{\rho_G-1}{\varsigma_J-1}}\right)^{\frac{\varrho-1}{\rho_G-1}}\right) \tag{1}$$

where Ω_{nG} is the expenditure share by country n on sector G, Ω_{nGJ} is the share of sector G spending on sub-sector J, Ω_{nJR} is the share of sub-sector J spending on foreign inputs, and ω_{nJR_m} is the share of foreign input spending in sub-sector J controlled by the hegemon. β and ρ are parameters of the aggregate production function invariant by sector.

Returning to our conceptual framework of Section 3.1, our measure of power assesses the difference in values (V) between the inside and outside option that the hegemon can induce with its threat. That is, a threat $\underline{\theta}$ is more powerful the more it lowers the V of the target.

In this paper, we simplify the formulation by taking the limit of some elasticities going to 1 (Cobb-Douglas), in particular $\varrho = \rho_G \to 1$, focusing on only the manufacturing sector G, and focusing only on internationally traded goods (assume $\Omega_{nJR} = 1$). We then consider how much power the hegemon country m has over country n if it were to cut off its supply of its output from a single industry J rather than all of its industries simultaneously. Under these conditions, the power formula for a single industry simplifies to:

$$\text{Power}_{mnJ} = -\frac{\beta}{1-\beta}\frac{1}{\sigma_J-1}\Omega_{nJ}\log(1-\omega_{nJR_m})$$

Intuitively, cutting an input from country m to country n inflicts more losses on country n whenever the input is in a sector J in which country n imports more from country m (a higher import share ω_{nJR_m}), the country spends more on foreign inputs in sector J (a higher Ω_{nJ}, where $\Omega_{nJ} = \Omega_{nG}\ \Omega_{nGJ}$), and a low elasticity of substitution among the foreign varieties of the input (a low σ_J).

In this paper we assess whether in a particular country $m-n$ relationship, and conditional on an input being restricted (by either export controls or sanctions), the sectors J being used are those in which the sender of pressure m would have the most power over the receiver n. Given the findings of Figure 2 Panel (c) that the overwhelming share of export controls in our sample captures export controls by the United States on China from 2018 to the present, we focus on analyzing the case where m is the United States and n is China. We define an indicator variable for whether a given American firm i producing goods in industry J faces export controls on selling its good to China at time t, $EC_{i(J),t}$.

We express our analysis in logit regression form as:

$$P(EC_{i(J),t} = 1) = \Lambda(\alpha_t + \beta \cdot \text{Power}_{mnJ}), \tag{2}$$

where $\Lambda(x) = \frac{e^x}{1+e^x}$ is the logistic inverse link function. We then turn our estimates into our coefficient of interest, the average marginal semi-elasticity effect δ, as $\delta = \mathbb{E}\left[\frac{\partial \log P(EC_{i(J),t}=1)}{\partial \text{Power}_{mnJ}}\right]$. We construct three different versions of the indicator variable for our dependent variable, $EC_{i(J),t}$ using the output of our LLM-based classification. First, we classify the indicator variable as equal to 1 if a given American firm reports that it is affected by export controls, regardless of whether the LLM classifies the receiving country as China. This approach maximizes the number of times the indicator flags positive, since the country receiving the controls might be missing. It is a useful benchmark assuming that most instances of export controls by the US are toward China. The results following this approach are in the first two columns, labeled “Any EC”, in Table 2. A second approach is to impose the stricter requirement that the LLM also classifies China as the receiving country of the pressure. The results are in columns 3 and 4, labeled “ EC on China.” The third approach is to consider both export controls and sanctions on China, since some export controls might be imposed in the form of sanctions on specific firms in China. The corresponding results are in columns 5 and 6.

We take the estimates of $\text{Power}_{mnJ,t}$ directly from Clayton et al. (2024) who estimate the input shares using bilateral trade data and input-output domestic data, and elasticities of substitutions σ_j from Fontagné et al. (2022).[16] We calibrate the returns to scale parameter $\beta = 0.8$ as in Clayton et al. (2024). To avoid the impact of previously imposed export controls affecting the power measures, we measure $\text{Power}_{mnJ,t}$=$\text{Power}_{mnJ,2018}$ using 2018 trade data, since Figure 2 shows export controls were largely imposed after this date.

We estimate this as a logit panel regression, with even-numbered columns in Table 2 including time fixed effects. We cluster our standard errors at the SIC4 level given that this is the level of variation of our power measure. For every specification, we find that American firms are significantly more likely to report being affected by export controls if they produce goods whose restriction would cause a more negative effect on the Chinese economy. The estimates imply that a 0.00134 increase in Power_{mnJ} (corresponding to moving from the median industry to the 99th percentile in the distribution of Power_{mnJ}) increases the probability that an American firm will be involved in export controls by 1.0228 log points (using the estimates from column 5).

This analysis suffers from several shortcomings. For example, we are not incorporating indirect linkages between the US and China (e.g. Taiwan in semiconductors), and bilateral trade data for disaggregated sectors is notoriously noisy (and estimates of elasticities even more so). Yet, the data shows a clear pattern that, in imposing export controls, the US government is looking for particular products that are "chokepoints." In ongoing work, we are working to expand our classification of vulnerability to the firm rather than sector level using disaggregated trade data as well as expand the set of countries and instruments considered.

[16]For each firm we select the sector J based on their primary SIC4 code from Factset. We then map the BACI trade data from HS6 to SIC4 using the WITS Concordance to measure the firms and trade data at the same level of aggregation. The elasticity σ is the average estimated elasticity at the HS6 for all products within the SIC4 group.

Table 2: **American power over China and imposed export restrictions**

	Any EC		EC on China		EC or Sanctions on China	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
$\hat{\delta}$	722.7**	724.3**	703.3*	711.3*	761.1**	777.6*
	(304.7)	(318.4)	(406.2)	(425.2)	(387.1)	(408.0)
Quarter FE	No	Yes	No	Yes	No	Yes
Observations	37,396	37,396	37,396	37,396	37,396	37,396
Pseudo-R^2	0.02	0.05	0.01	0.06	0.02	0.06

Notes: This table shows the estimates from the logit regression specification in equation (2). We show marginal semi-elasticity effects. Standard errors are clustered at the SIC4 level and obtained via the delta method.

3.5 A Global Map of Geoeconomic Pressure

Before turning to how firms around the world respond to the pressure that they face, we briefly present a global snapshot of exposure to pressure. To expand the coverage of our analysis, instead of relying on the earnings calls of publicly traded firms, we use analysts reports at the country-sector level. From 2017 to 2025, we have 44,495 of these country-sector reports. For larger countries like China, this leads to hundreds of reports, whereas for smaller countries the coverage is much more sparse. For instance, while we include Syria in the global map, we note that we only have 5 reports on the country. The reports cover a range of sectors, such as Oil & Gas, Banking, Information Technology, Defense, and Pharmaceuticals.

For each report, we run our first stage prompt and classify whether the document mentions that firms in the sector are affected by tariffs, sanctions, export controls, and the other forms of geoeconomic pressure. Then, for each country, we average this dummy variable for all sectoral reports for that country and calculate the share of reports that mention firms in the sector being affected by a given form of geoeconomic pressure. The results are shown in Figure 4. Consistent with our analysis of earnings calls and single-firm analysts reports, tariffs are reported to be the most commonly used form of pressure, followed by sanctions, and export controls. While some of the patterns in Figure 2 can also be seen here, such as the exposure of American firms to tariffs and Chinese firms to export controls, the expanded country coverage in Figure 4 captures pressure felt by firms in smaller countries whose firms are less likely to have earnings calls or to be the focus of a single-firm analyst report from a Western investment bank. For instance, here, we see clearly the exposure of Sudan, Venezuela, Iran, and Myanmar to sanctions. In ongoing work, we are adapting our second-stage prompts for a more detailed analysis of the global response to pressure to complement the existing analysis. While there is limited coverage of firms in the aforementioned sanctioned countries in the earnings calls or single-name analyst reports, we are working to explore their responses using this related sectoral text.

4 Firm-Level Responses to Geoeconomic Pressure

In this section, we analyze how firms respond to geoeconomic pressure by examining qualitative details extracted from the earnings calls and analyst reports. We focus on the various strategies firms adopt in reaction to current or potential future sanctions, tariffs, and export controls.

In Figure 5, we summarize how firms respond to the different forms of geoeconomic pressure throughout our sample period. In the top panel, we measure responses using the earnings call sample and in the bottom panel we use the analyst reports. For every earnings call in which the firm reports being affected by the particular form of geoeconomic pressure (either present or future), we classify whether the firm responded using any particular action. For this figure, we plot the average net response within a category, conditional on being affected. For instance, consider the response of firm R&D: if a firm reports raising its research and development activity in response to sanctions but not lowering it, it receives the value of 1 for "R&D". If it reports lowering its investment but not raising it, it receives the value of -1. If it reports both raising and lowering R&D in response to tariffs (say increasing research in one product line and cutting it in another), it receives a value of zero. The reported value in Figure 5 is then the average net response of firms reporting being affected by tariffs, sanctions, and export controls.

A number of findings emerge. First, we find that firms reporting being affected by tariffs are more likely to also report facing higher input prices ("Input Price") and are more likely to report raising the price they charge for their product ("Sales Price"). In addition, firms reported as being affected by tariffs also report having the largest reduction in their net reported profit margin ("Profit Margin"). By contrast, we find that firms reporting being affected by export controls on net raise their research and development ("R&D"), investment ("Investment"), and in particular their investment in their home country ("Domestic Investment"). Finally, firms reporting being affected by sanctions report the largest net drop in input purchases ("Input"), sales ("Sales"), although both effects are small in magnitude, and report on average being the most net negatively affected. For all three instruments, the firms report on net being negatively affected by the application of pressure.

Of course, there are many ways that a firm can be "affected" by tariffs, sanctions, and export controls. In Figure 6, we investigate separately the responses of firms based in the country imposing the pressure ("sender country") relative to those based in countries receiving the pressure ("receiver country"). For instance, if we consider an American firm that we classify as reporting being affected by American tariffs on China, its responses would be analyzed in the "sender" graph as its home country, the United States, is the source of the pressure. If, instead, it is the text of a Chinese firm that reports being affected by the same American tariffs imposed on US imports from China, this firm would be analyzed in the "receiver" figure as its home country, China, is the target of the pressure.

Figure 4: **A global map of geoeconomic pressure**

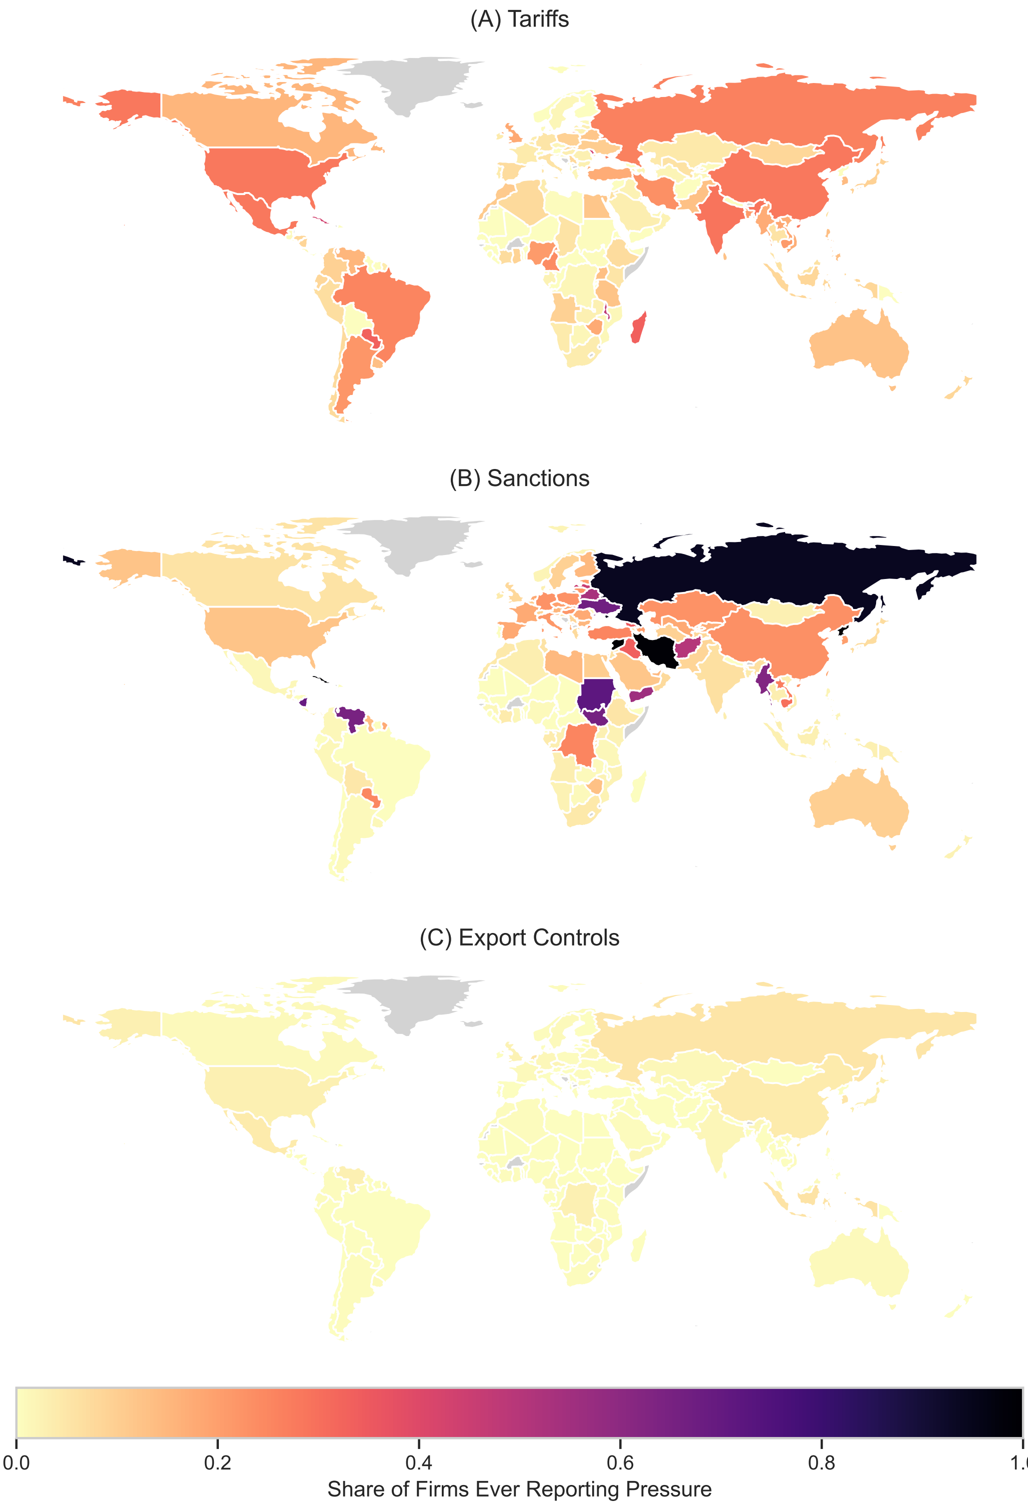

Notes: This map plots the share of all country-sector reports for a given country that report firms in that sector being affected by tariffs (Panel A), sanctions (Panel B), and export controls (Panel C) for the period 2017-2025.

Figure 5: **Firms' responses to geoeconomic pressure**

(a) Earnings Calls

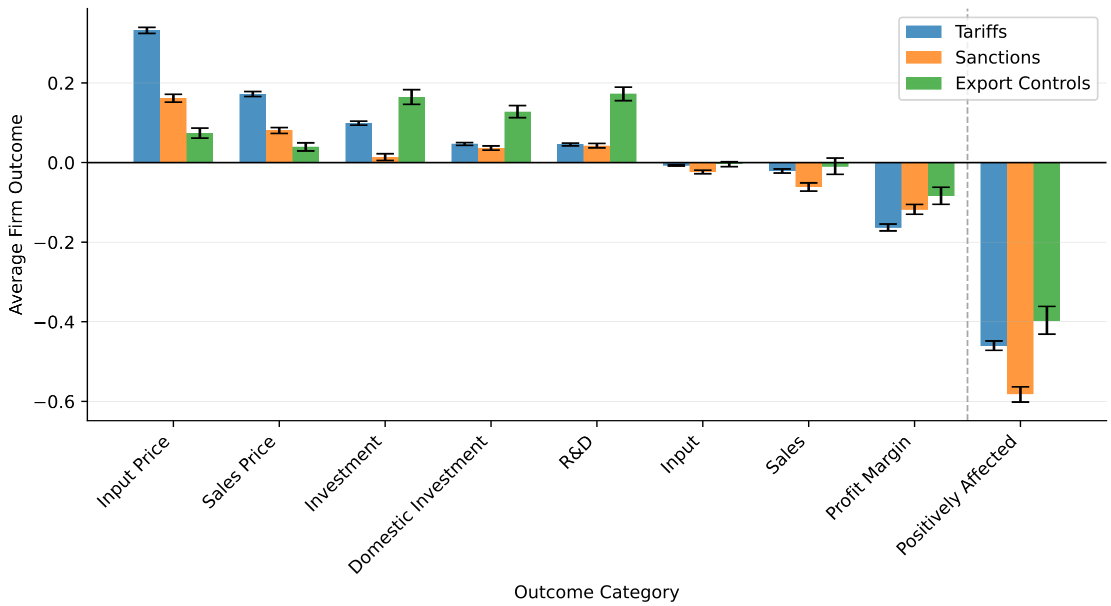

(b) Analyst Reports

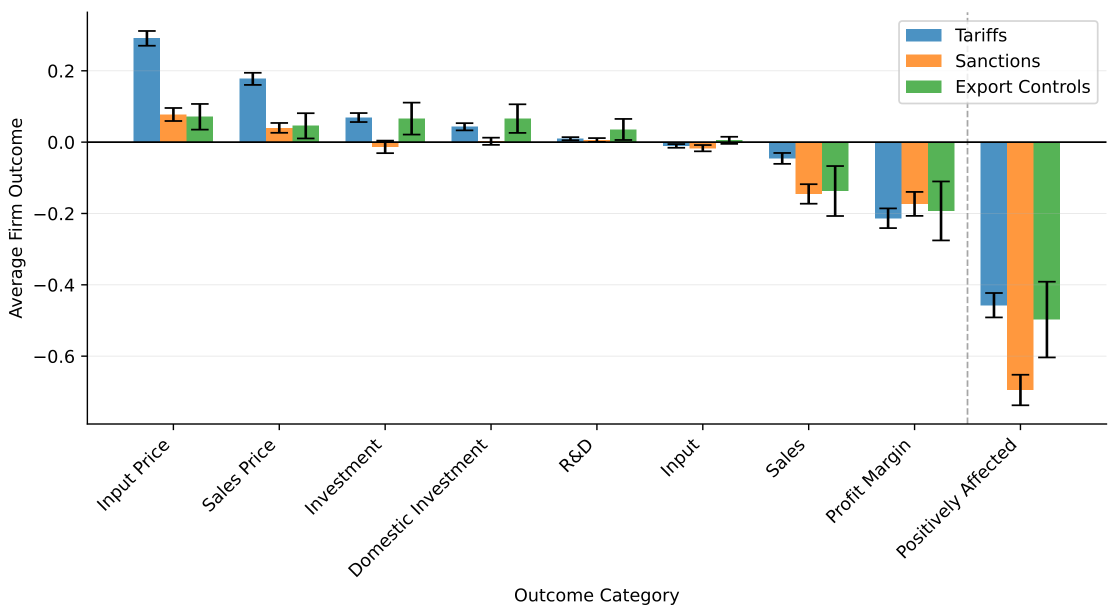

Notes: This figure plots the net response of firms affected by each form of geoeconomic pressure. The outcomes are coded as 1 for an increase, -1 for a decrease, and 0 for no directional change reported; average outcomes conditionally on the firm being affected are reported. Two standard error bands are added to each bar. Panel (a) is based on firms' earning calls, while Panel (b) is based on single-firm analysts reports.

Figure 6: **Responding to pressure, by country role**

(a) Firms in Sender Country

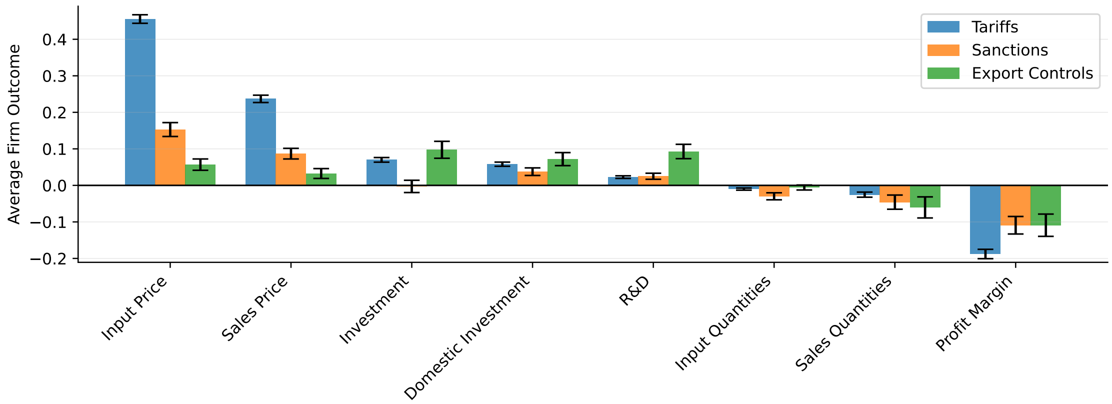

(b) Firms in Receiver Country

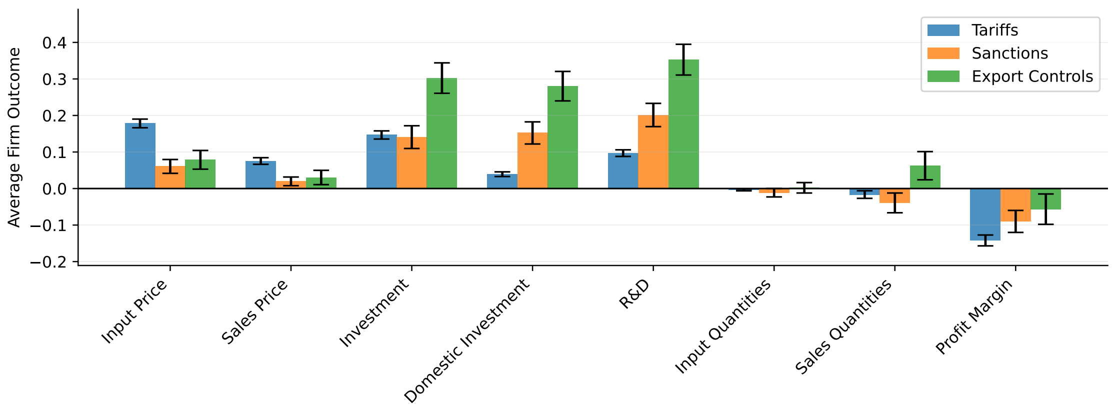

(c) Firms in Impacted Third-Party Country

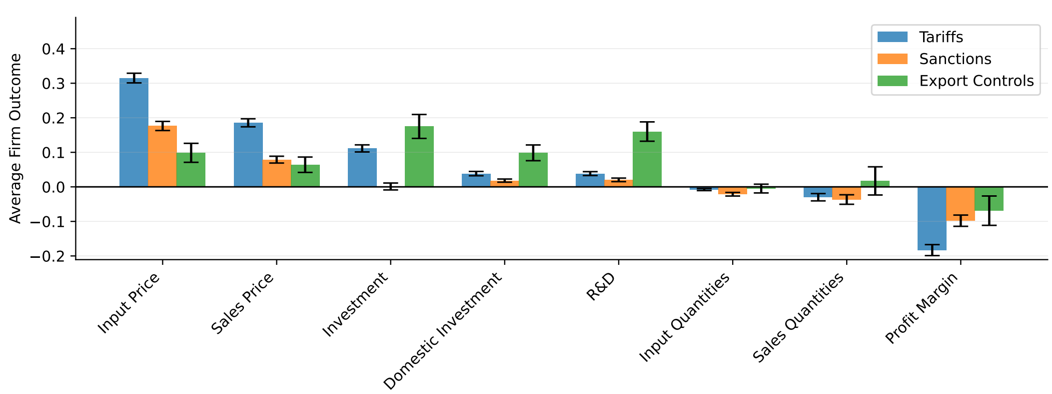

Notes: This figure plots the net response of firms affected by each form of geoeconomic pressure. Two standard error bands are added to each bar. Panel A plots the average net response of firms located in the country the government of which is the sender of the pressure. Panel B plots the average net response of firms located in the country that is the receiver of the pressure. Panel C plots the average net response of firms located in countries that are neither the sender nor the receiver of the pressure. Each panel displays the response separately by instrument of pressure: tariffs, export controls, and sanctions.

Figure 6 sheds further light on the patterns documented in Figure 5. In particular, we find that on net nearly 50% of firms in the country imposing tariffs ("sender") report facing higher input prices whereas less than 20% of those in the target ("receiver") country do. The asymmetry is also striking for the price firms charge for their goods, with nearly 25% of firms in the source country reporting raising the price they charge, whereas less than 10% in the destination country do so. Tariffs are sometimes imposed with the aim of lowering prices of imports as the foreign exporters absorb part of the tariff. Our analysis does not find substantial evidence of this mechanism. Instead, we find evidence that the firms in the importing country face higher input prices, including the tariff, and pass through some of that into higher sales prices. This is the consumer-tax component of a tariff.[17]

Turning to export controls, we find that the increase in R&D, investment, and domestic investment is essentially entirely concentrated in the countries targeted by the export controls. For instance, in September 2023, Rigol Technologies, a Chinese electronics manufacturer, responded to a question in an earnings call by stating:

> Due to foreign technological blockade,The impact of a series of issues such as export controls,There are significant bottlenecks in the upward development of the domestic electronic measuring instrument industry,There is a need to strengthen research and development, Expand production, Investment such as mergers and acquisitions integration. *[Punctuation as in original data]*

On the basis of these statements and other parts of the call, the LLM classifies the firm as increasing its domestic investment and raising its R&D in response to the export controls impeding its access to foreign technology. More generally, we find that much of this pattern is driven by Chinese firms targeted by American export controls reporting that they are increasing their R&D and investment inside China in order to create substitutes for the blocked American products. This is consistent with the findings in Flynn et al. (2025) that sectors more exposed to foreign political risk conduct more innovative activity. The response to sanctions largely follows similar qualitative patterns as export controls, with a smaller response of R&D and investment.

Figure 6 Panel (c) focuses on firms in third-party countries that are neither the sender nor receiver of the pressure. For example, the earnings call of a Dutch firm might be discussing export controls imposed by the US toward China on semiconductors, or an Indian firm might discuss how it is affected by the sanctions imposed by the US and Europe on Russia in oil and natural gas. This third-party country variation not only contributes to identifying overall episodes of pressure, but is also of independent interest since it tells us how pressure from a given country A on country B affects firms in country C.

The third-party country effects have two main cases behind them. In one set of cases, the imposing government is using a foreign entity (with or without that entity's government involvement) to apply the pressure. This is the case, for example, of the US applying or threatening to apply

[17]The pass-through of tariffs to prices see Cavallo et al. (2021), Fajgelbaum et al. (2020), Amiti et al. (2019), especially for the tariffs imposed in the 2018-19 period by the US government.

secondary sanctions on foreign financial institutions that facilitate transactions with Russia or Iran irrespective of whether the transactions involve a US entity. Another example is the US exerting pressure on Dutch company ASML to prevent the sale of its advanced lithography machines to certain entities in China. In the other set of cases, the third-party firms are affected because of a change in their business environment as a byproduct of the pressure. For example, an Indian firm might discuss being affected by the US sanctions on Russia because as a result it is now obtaining oil cheaply from Russia or because it has now acquired new customers for its products in the US since US firms lost their relationship with a competitor supplier based in Russia.

4.1 The Response of Global Supply Chains to American Pressure

Global responses to US pressure. We next turn to exploring how firms adjust their supply chains when facing the various forms of geoeconomic pressure. To do so, we use the results of two fields: the list of countries the firm reports adjusting their supply chains towards as a result of pressure (field 44 in our second-stage prompt, "countries_supply_chain_adj_towards") and the list of countries the firm reports adjusting their supply chains away from as a result of pressure (field 45, "countries_supply_chain_adj_away"). Here, we focus on tariffs, sanctions, and export controls where the United States is the imposing country. For each policy, we consider the response of three different groups of firms: American firms, Chinese firms, and firms from the rest of the world. In each graph, we display the number of firms reporting to be reallocating their supply chains towards a given country by moving upwards on the y-axis and we report the number of firms moving away from each country as negative numbers on the y-axis.

Figure 7 Panel (a) displays the results for tariffs. In response to American tariffs, American firms disproportionately shift their supply chains towards the United States (panel a.1). In addition, we find reallocation towards Vietnam, Malaysia, and India. As a result of these US-imposed tariffs, we find a strong move of American supply chains away from China. This result is consistent with the stated objective of the US government of onshoring supply chains. Turning next to the response of Chinese firms in Panel (a.2), we find striking differences compared to the American firms' responses. Chinese firms report that as a result of American tariffs, they move activity away from both the United States and China, although in both cases a number of Chinese firms also report moving their activity towards those countries. For Chinese firms, the strongest pattern in response to the American tariffs is the expansion of their supply chains throughout Asia, moving their supply chains to Vietnam, Thailand, Malaysia, India, and others. In addition, we find that Chinese firms report expansion into Mexico. These patterns are consistent with some firms re-routing China-US trade via East Asia and Mexico to minimize the impact of the tariffs.[18] Finally, Panel (a.3) explores the response of firms in the rest of the world to American tariffs. We find a striking reallocation of supply chains away from China and towards the United States, as well as into other Asian countries and Mexico.

[18]See Grossman et al. (2024), Alfaro and Chor (2023), Freund et al. (2024), Utar et al. (2023), and Gopinath et al. (2025) for evidence on the reshaping of global supply chains.

Figure 7: **Response of global supply chains to US pressure**

(a) Tariffs

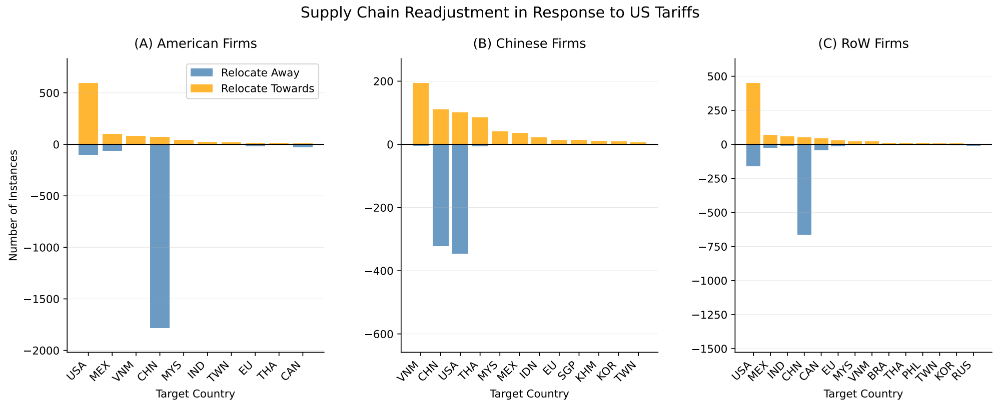

(b) Sanctions

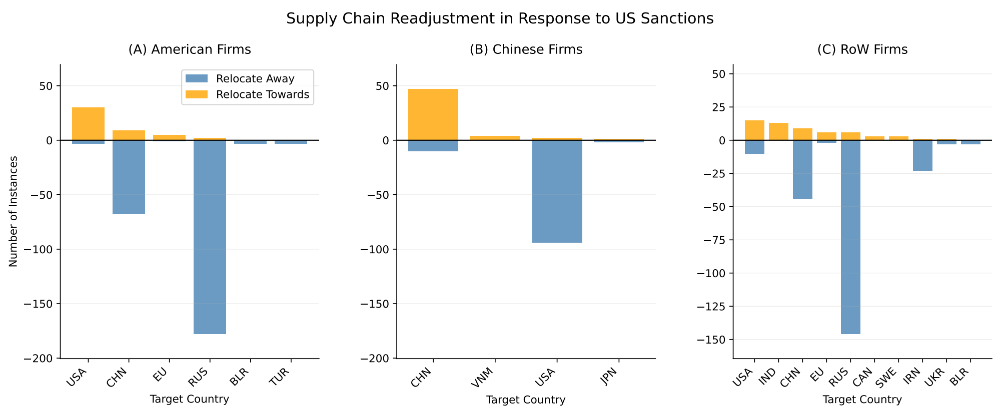

(c) Export Controls

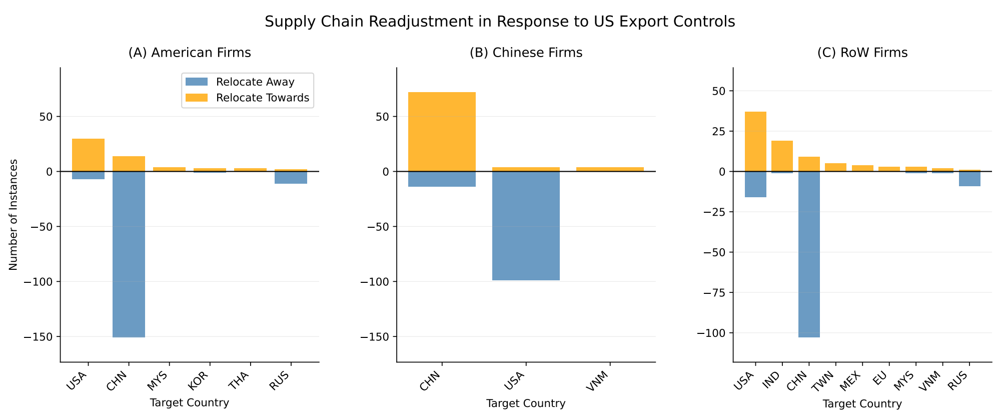

Notes: This figure plots the number of firms in each country or group of countries reporting that they are adjusting their supply chains towards (positive numbers, yellow) or away from (negative numbers, blue) each country in response to US government pressure. The top row considers the response to tariffs, the middle row the response to sanctions, and the bottom row to export controls. In all three cases the policies are imposed by the US government. The left column reports the response of American firms, the middle column reports the response of Chinese firms, and the right column reports the response of firms from the rest of the world (RoW).

In the bottom two Panels of Figure 7, we perform analogous analyses for sanctions and export controls. In response to American sanctions, American firms reallocate towards the United States and away from China and Russia, which—as shown in Figure 2—were the primary recipients of American sanctions along with Iran. By contrast, Chinese firms respond to American sanctions by shifting away from the United States and back home towards China. The rest of the world is in between with offsetting movements into and out of the United States, large shift out of the targeted countries of China, Russia and Iran, and movements towards India. One challenge in this analysis, which we will return to shortly, is that the coverage of Russian firms' earnings calls is generally low and shrinks even further once sanctions are imposed.

For export controls, the dominant movement of American firms is towards the United States and away from China and to a lesser extent Russia. In the case of China, we find a major reallocation of supply chains towards China and away from the United States. For the rest of the world, we find a net movement away from China and towards the United States, although we also observe firms moving away from the United States as well.

The response to sanctions of Russian firms. We next turn to exploring how Russian firms respond to sanctions in the wake of Russia's invasion of Ukraine in 2022. As mentioned above, earnings calls as a source of text have a serious shortcoming due to the limited availability of Russian corporate text, especially after the sanctions. We therefore turn to country-sector analysts reports. While these do not directly cover individual firms, they do cover 16 sectors of the Russian economy at a quarterly frequency. Figure 8 plots the response of Russian firms using these sectoral reports. We find that nearly 40% of Russian sector-quarter reports mention Russian firms reallocating their supply chains towards China, with around 15% moving to India, and smaller moves towards Kazakhstan, Iran, and Turkey. By contrast, we find that nearly 40% mention Russian sectors moving away from countries in the European Union, and over 20% moving away from the United States, with a smaller movement away from the United Kingdom and Japan.[19]

5 An Ongoing Look at the Trade War of 2025

In this section, we zoom in on the effects of the ongoing trade war that started in the first quarter of 2025. Given the rapid developments in global trade policy at the time of writing, we anticipate that this section will change significantly as ongoing events unfold. Here we analyze text from the first quarter of 2025 and any text that had become available to us by mid May 2025, thus capturing the lead up to the April 2nd announcement of tariffs from the US administration and the initial response after the shock.[20]

Tariff announcements by the US administration have come at high frequency in 2025 and often

[19] See Egorov et al. (2025) for a detailed study of the responses of Russian firms to trade sanctions.

[20] For papers tracking and assessing the ongoing trade war see: Auclert et al. (2025); Rodríguez-Clare et al. (2025); Kalemli-Özcan et al. (2025); Ignatenko et al. (2025).

Figure 8: **The supply chain response of Russian firms to sanctions, 2022-2025**

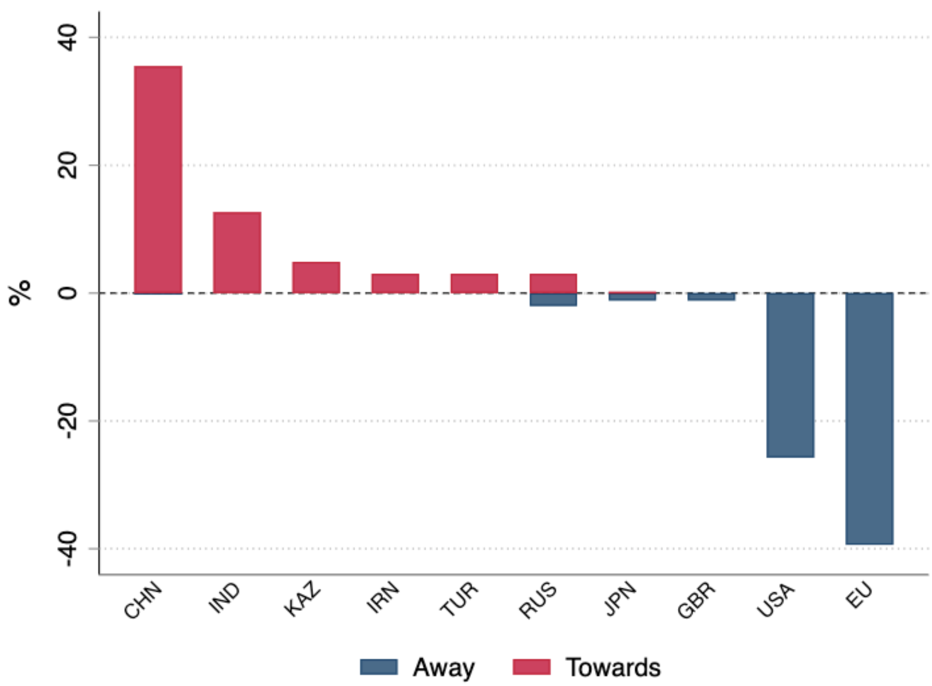

Notes: This figure reports the percentage of country-sector analyst reports focused on the Russian economy in which firms are classified as moving their supply chains towards or away from certain countries as a result of sanctions. The data covers the period 2022-2025.

changed the announced tariffs before they were ever implemented. Figure 9 quantifies the relative importance of firms expecting future tariffs versus being affected by existing ones. The vertical dotted lines are the fourth quarter of 2016 and 2024: in both instances, President Trump won a presidential election in November running on a campaign that made tariffs a major feature of his presidential policy. In both instances, firms reported being affected by the prospect of future tariffs. In the 2016 episode, the firms eventually transitioned by 2019 from expecting future tariffs to being affected by tariffs actually imposed by President Trump and eventually largely maintained by the Biden presidency. In 2025, thus far, we observe firms being affected by both current and future tariffs. This is consistent with the Trump administration both immediately imposing tariffs (for example on steel) and announcing further tariffs to be imposed at an uncertain future date.

Figure 9: **The effect of current and future tariffs**

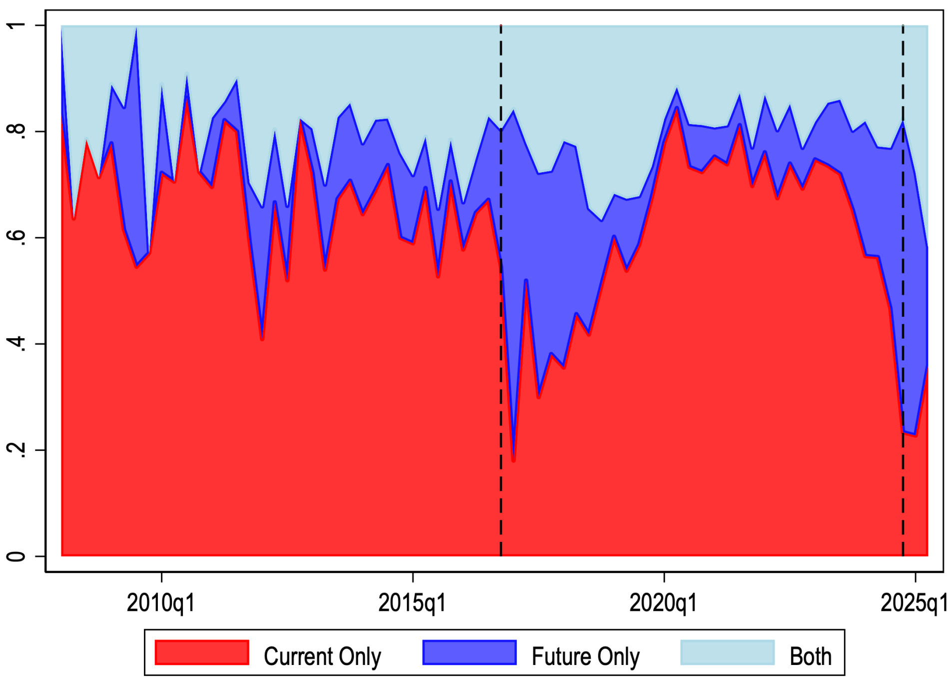

Notes: This figures displays the share of firms reporting being affected by current tariffs only, future tariffs only, or both current and future tariffs. The vertical dotted lines are the fourth quarters of 2016 and 2024.

Figure 10 displays the share of both American and non-American firms that report being negatively or positively affected by tariffs. The time series for American firms is plotted in the top panel of Figure 10 and shows that the share of American firms reporting being negatively affected by tariffs is at a high for our sample at the beginning of the second quarter of 2025.

Interestingly, a number of firms report being positively affected by the tariffs. We inspected both the LLM-generated summaries for these firms and their full transcripts. Indeed, some American firms report being able to increase their market share because their competitors, either other

Figure 10: **The self-reported effect of tariffs on American and non-American firms**

(a) American Firms

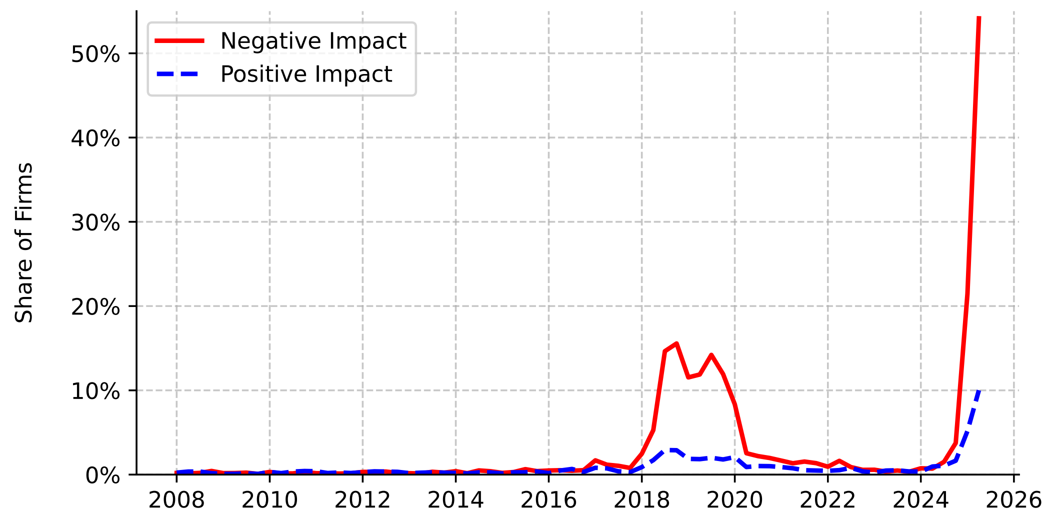

(b) Non-American Firms

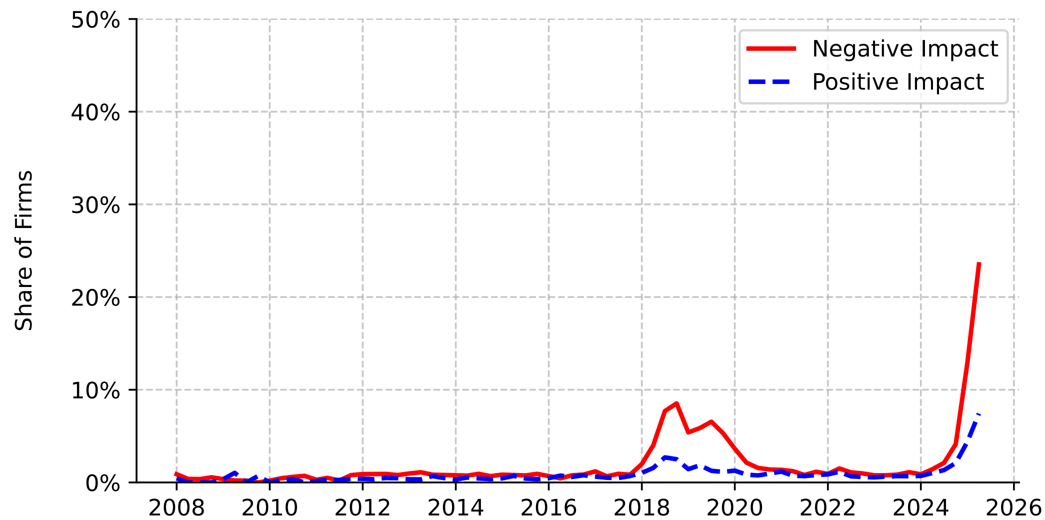

Notes: This figure displays the share of firms reporting being positively or negatively affected by tariffs. Panel (a) focuses on American firms, and Panel (b) on firms in the rest of the world.

American firms or foreign firms that sell in the U.S., are negatively affected by the tariffs. Some firms also mention facing better profit margins and being able to increase product prices more than increases in input costs. Indeed, in trade models the imposition of tariffs can cause domestic producers that do not face any increases in input costs to gain market share and/or raise their prices to take advantage of the fact that their competitors, who face higher input costs due to a tariff, are passing through some of the tariff into their price. As an example, the CEO and CFO of Cleveland-Cliffs Inc, an Ohio-based steel manufacturer, mention in the company earnings call on February 25th 2025 that:

> "With the Trump administration in office, action is being taken and we are starting to see positive signs ahead of us. We at Cleveland-Cliffs appreciate the recently announced 25% tariffs on steel importers from all countries. These tariffs are critical to addressing the problem, and we thank the Trump administration to have the courage to implement these tariffs. [...] Cleveland-Cliffs is not depending on imported inputs and we do not rely on foreign supply chains that can be disrupted overnight. The tariffs will penalize the foreign competitors who have been playing by a different set of rules while strengthening the domestic producers who actually invest in American workers, American manufacturing and American supply chains. [...] Of course, the short term, there will be kind of a rearrangement of the supply chains. But look, nobody can say that it's efficient to like the CEO of one steel company said some parts move across the border between the United States and Mexico 7 times. And then they call that efficiency? My goodness, that's stupid, with all due respect. So let's produce everything here in the United States and get things back where they belong. And don't forget, for a country like Mexico, tariffs will be stacked. So it's 25% plus 25% makes up 50%. So for the ones that relied in Mexico, time to get another thing to do."

As it turns out, in March 2025 Cleveland-Cliffs laid off 1,200 workers in their Michigan and Minnesota operations, but the management maintained an upbeat assessment for the long-run impact of the tariffs.

We next turn to how firms are responding to the ongoing trade war. In the top panel of Figure 11, we compare the response of American and non-American firms to tariffs in 2025. The clearest difference between the response of American and non-American firms is that American firms are far more likely to report being affected by higher input prices. Given that the United States was the imposing country for the key tariffs, this increase in costs corresponds to foreign exporters not (fully) absorbing the tariff by cutting their export prices to the US. Next, we find that American firms are more likely to report raising the price they sell their own goods for. American companies are more likely to report a drop in their profit margin and to be more negatively affected.

In the bottom panel of Figure 11, we explore how the responses of American firms differ conditional on whether they report being positively or negatively affected. First, confirming the consistency of our measure, we find that firms that report being positively affected are far more likely to

Figure 11: **Responses to tariffs of American and non-American firms, 2025**

(a) American firms vs. non-American firms

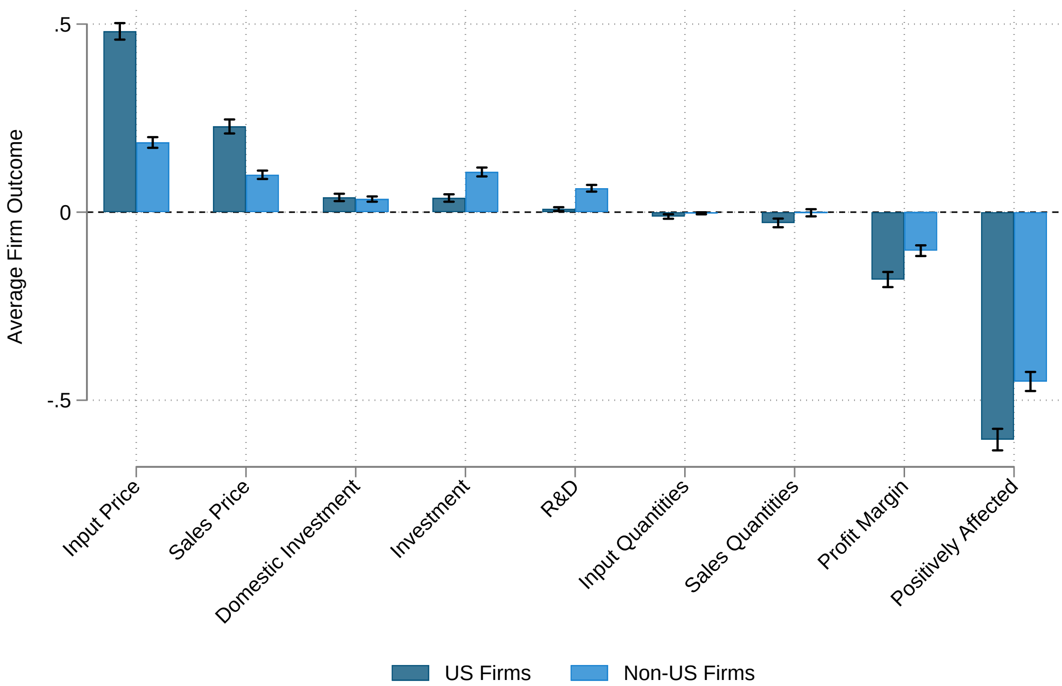

(b) American firms reporting positive vs. negative impact

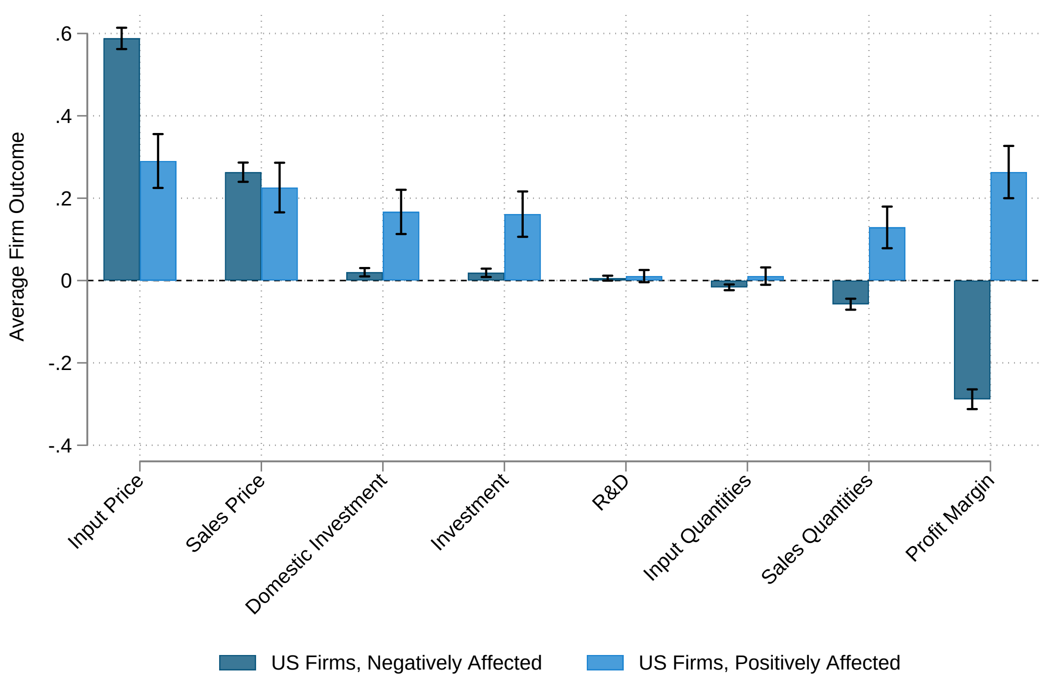

Notes: This figure shows how firms describe adjusting their business in response to tariffs in 2025Q1 and 2025Q2. Panel (a) groups the responses by American firms and firms in the rest of the world. Panel (b) focuses only on American firms and groups the responses by whether the firms report being overall positively or negatively affected.

report increases in their profit margins than are firms that report being negatively effected. These positively affected firms reported increases in their sales and reported responding to the tariffs by raising their investment. One key difference between positively and negatively affected firms is that negatively affected firms are far more likely to be facing higher input prices. Interestingly, positively and negatively affected firms report raising their own sales price at similar frequencies, suggesting that this higher price represents a rise in profit margins for the positively affected firms that don't rely on imported inputs affected by the tariffs, whereas their negatively affected peers are raising their prices to partially offset the higher costs of inputs. Indeed, this is a central mechanism in trade theory of how tariffs lead to higher prices even from those domestic producers who are not directly affected by the tariffs on their inputs.

6 Measurement Uncertainty and Reproducibility

In our empirical analysis, the key outcomes are generated by an LLM. Since LLM outputs are algorithmic proxies, an important question is how much of the variation in the measurement reflects signal versus noise that is inherent to the measurement methods. In this section, we use a model-perturbation and prompt-perturbation methodology to develop a quantitative understanding of the magnitude of the measurement error induced by the LLM-based methodology. We discuss which aspects of the properties of measurement error this approach sheds light on, as well as its limitations, and we connect to the recent literature on econometric approaches to inference with algorithmically generated variables. We also review broader sources of potential measurement error (i.e., ways in which we may imperfectly capture our desired notion of geoeconomic pressure and related firm responses) that stem from the nature of the data and not just from the LLM-based methods. Lastly, we discuss how we designed our computational practices to optimize for data confidentiality and scientific reproducibility.

Measurement reliability: framework. We start with a simple conceptual framework that allows us to distinguish formally between signal and LLM-induced noise. For each document i, we let Y_i be the vector of measured outcomes. This outcomes vector crucially depends on the particular prompt specification p and on the chosen pre-trained LLM m, so that $Y_i = Y_i(p, m)$, which we write as a deterministic mapping since all our analyses use a deterministic (i.e., zero-temperature) sampling strategy: this lets us abstract from stochasticity that simply occurs because of random sampling from the LLM's generated distribution over output tokens.[21]

The particular prompt p used for the baseline analysis is, to an extent, arbitrary, since there are many different prompts that would be equivalent, in the sense of having the same semantic meaning (i.e., the same substance) and specifying the same output schema, but with different

[21]We do not focus on an analysis of the size of the measurement error induced by the sampling since it does not have inherent economic content, as the scale of the sampling variability is arbitrarily determined by the chosen sampling temperature. We instead analyze how the outcomes Y_i change as we span the space of admissible prompts p and use different models m, which has a natural (and economically meaningful) scale.

wording and/or prompting structure. We let $\mathcal{P}$ be the set of all such admissible equivalent prompts. Naturally, the set $\mathcal{P}$ is in principle infinitely large, and therefore we need to proxy for it in a way that is computationally tractable. In practice, we approximate $\mathcal{P}$ by constructing a finite, carefully designed set of prompt perturbations that are meant to span this space, from prompts that are very similar to the baseline specification p to others that are drastically different in both wording and structure. We refer to our set of prompt perturbations as $\tilde{\mathcal{P}}$, with elements $\tilde{p}_k$ whose construction we describe below.

This gives rise to a series of measurement outcomes $Y_i(\tilde{p}_k, m)$, which are the values obtained for document i under prompt perturbation k and using model m. Each field d corresponds to one of the elements of $Y_i(\tilde{p}_k, m)$, which we write as $Y_{id}(\tilde{p}_k, m)$. We model these outcomes as:

$$Y_{id}(\tilde{p}_k, m) = \bar{\eta}_d + \eta_{id} + \varepsilon_{idkm}, \qquad \text{with} \qquad \sigma^2_{\eta,d} = \mathrm{Var}\left(\eta_{i,d}\right), \; \sigma^2_{\varepsilon,d} = \mathrm{Var}\left(\varepsilon_{idkm}\right),$$

where η_{id} is the signal (the across-documents variation of interest) and ε_{idkm} is the noise introduced by the LLM as we span the set of admissible prompts and vary the model choice (which is assumed independent of η_{id}, a point that we return to in our discussion of non-classical measurement error below). The amount of signal relative to noise for a field d is then quantified by the *reliability ratio* R_d, defined as the variance of the signal relative to the total variance of $Y_{id}(\tilde{p}_k, m)$:

$$R_d = \frac{\sigma^2_{\eta,d}}{\sigma^2_{\eta,d} + \sigma^2_{\varepsilon,d}}.$$

In a setting where measurement error is classical and ε_{idkm} is the only source of error, the reliability ratio has a natural interpretation as the exact attenuation factor one would face when using the given variable as a regressor in an OLS setting.[22]

Prompt and model perturbations. For each of the baseline prompts used in the analysis (one first-stage prompt and multiple second-stage prompts, as described in Section 2.3), we construct $K = 20$ prompt perturbations. Each of the 20 variations may vary the prompt wording, the language used to describe the economic policies and fields of interest, the section ordering and formatting cues, or the cognitive instructions provided to the model—all while preserving task semantics. The set includes 10 minor variations, each altering only a few very specific sentences; 5 medium variations, which retain the structure of the baseline prompts but modify the wording and language; and 5 major variations, which are major departures from the baseline prompts meant to give as little instruction and structure to the LLM as possible while preserving the schema of the data (an attempt to deliberately "break" the model).

More specifically, the minor variations substitute groups of words (up to two sentences) with semantically equivalent variations and reorder the components of the prompt. Medium variations

[22]When specializing the reliability statistics to perturbations within a single model, naturally the variances and the reliability ratio would all acquire an additional subscript m.

use substantively different language to describe all the policies and fields of interest, while leaving the structure unchanged. Major variations significantly alter both the language and the structure of the baseline prompts, greatly reducing their length and complexity to the minimum possible while retaining satisfactory output schemas: these are designed to provide intentionally brief and bare-bones instructions, with the aim of stress-testing the model with minimal guidance. Importantly, there is no guarantee that the mapping $Y_i(p, m)$ has any particular regularity properties: given the highly non-linear nature of the deep learning architectures underlying the LLM models, arbitrarily small changes to the prompt p (such as changing a single word) can in principle lead to significant changes in the output, possibly as large as those induced by much larger alterations to the prompt—hence our approach using a range of variations, from minor to major, to span the set $\tilde{\mathcal{P}}$.

With respect to model perturbations, we use $M = 2$ pre-trained, open-weight LLMs. The first one is the 70 billion parameter version of Llama 3.3, released by Meta (Llama 3.3-70B, denoted as $m = m_L$): this is the model that is used as the baseline for all the results shown in the other sections of the paper. The second model, which we use for this perturbation analysis and to assess robustness of the other results, is Google's Gemma 3-27B ($m = m_G$). This generates a total of 40 combinations $(\tilde{p}_k, m)$. In ongoing work, we are also integrating perturbations using Alibaba's Qwen 2.5-72B model.

Measurement reliability: empirical implementation. We use our prompt and model perturbations to measure the variances. To estimate the noise variance $\sigma^2_{\varepsilon,d}$, we use within-document, across-perturbation sample variances:

$$\hat{\sigma}^2_{\varepsilon,d} = \frac{1}{N}\sum_{i=1}^{N}\widehat{\mathrm{Var}}(Y_{id}(\tilde{p}_k, m)),$$

where $\widehat{\mathrm{Var}}(\cdot)$ is the empirical sample variance operator (which here acts on the indices k and i), and N is the number of documents. Analogously, the signal variance σ^2_η is estimated using the across-document variation fixing the baseline prompt specification:

$$\hat{\sigma}^2_{\eta,d} = \widehat{\mathrm{Var}}(Y_{id}(p, m)),$$

where the empirical sample variance operator now acts across documents (over the i index). Given the large compute resources required to run all the perturbations, we estimate these variances using a stratified sample of 3,000 documents. The stratified sample contains 750 samples (25%) flagged as negative from the first-stage prompt in our baseline and 2,250 samples (75%) flagged as positive.

We show our empirical estimates in Figure 12 for a sample range of fields drawn from both the first-stage and second-stage prompts. Figure 12a shows the reliability ratios using all perturbations except for the major ones, while Figure 12b also includes the major ones. The specification which excludes major perturbations is our preferred one, as the latter one is constructed to deliberately stress the estimation by providing very little guidance to the model—nonetheless, they show some-

Figure 12: **Reliability ratios from prompt and model perturbations**

(a) Excluding major prompt perturbations

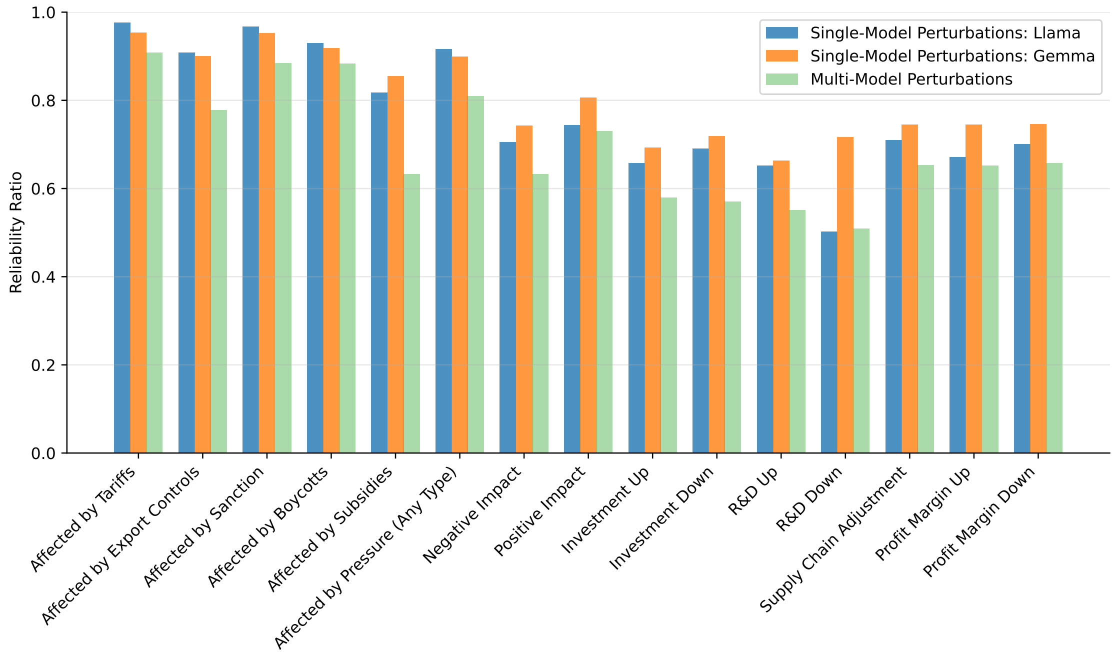

(b) Including major prompt perturbations

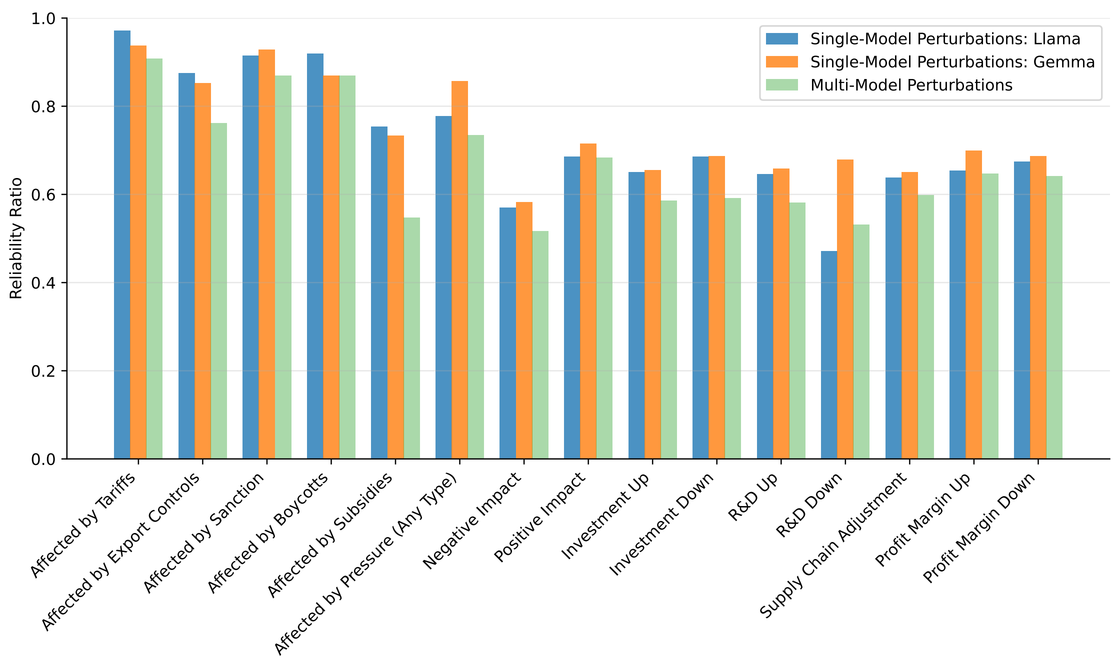

Notes: We show estimated reliability ratios for various fields through prompt-perturbation and model-perturbation analyses.

thing akin to a worst-case scenario. The bars labeled "single-model perturbations" show the ratios estimated by fixing the model (either Llama or Gemma), while the bars labeled "multi-model perturbations" include both model outputs while computing variances and hence incorporate across-model variation.

The estimated reliability ratios show that while the LLM-generated fields as expected have a substantial amount of noise, quantitatively the magnitude of the noise relative to the signal is overall in line with other data sources that are commonly used in empirical work in economics. Across fields, the reliability ratios have an average of 78.4% for single-model perturbations and 69.6% for multi-model perturbations. Including the major prompt perturbations, the average ratios are reduced only modestly, to 73.7% and 67.3% respectively, suggesting that even our base set of perturbations does well in spanning the set $\mathcal{P}$. The interquartile range is 70.6% to 90.6% for single-model and 62.8% to 79.4% for multi-model perturbations. More straightforward fields, such as the top-level booleans from the first-stage prompt, exhibit higher reliability (often in excess of 90%). Also, including cross-model variation tends to result in lower estimated reliability. There are some outlier fields that have especially low reliability, such as R&D down, but none go below about 50%. These magnitudes are comparable to those in high-quality administrative survey data such as the Panel Study of Income Dynamics, the Current Population Survey, or even NIPA components, which are commonly estimated to have reliability with similar typical ranges (e.g., Bound and Krueger 1991; Pischke 1995; Schmillen et al. 2024). Naturally, since the amount of noise is substantial, any analysis that uses LLM-generated variables as independent regressors would be best served by adopting IV strategies to correct for the measurement error: in this paper, we do not use IV corrections since all of the current analysis uses them as dependent variables.[23]

Importantly, the perturbation analysis presented here has some key limitations. While the analysis is helpful in quantifying the magnitude of the error, it is silent on whether the error is classical or not. In other words, while we can speak to variances, our analysis cannot make a conclusion as to whether the LLM-generated outcomes exhibit bias or correlation with other analysis variables. Determining these other properties of the error typically requires at least a subsample that is measured without error (a "ground truth" sample). Error correction approaches studied in the econometrics literature include Battaglia et al. (2024) and Ludwig et al. (2025). Further, the error induced by the LLM methodology is unlikely to be the only noise source, which we turn to next.

Broader sources of error. While the reliability analysis above quantifies the magnitude of a particular form of measurement error (induced by reasonable prompt and model variations), we stress that other sources of error are also possible in this context, having to do not just with the LLM-based methodology, but with the textual data itself. Indeed this set of challenges, which we remark on here, is common in the broader literature on text analysis and computational linguistics applied to economics.

[23] In ongoing work, we are testing whether model fine-tuning is helpful to increase reliability.

First, selection into the observable corpus is likely to be non-random. For earnings calls, for example, firms voluntarily choose whether to hold a call, they may have some degree of control over which questions to accept from analysts, and they choose how to answer them, all of which may be strategically motivated. Downward-biased coverage in earnings calls is particularly acute for jurisdictions under severe sanctions (e.g., Russian or Iranian firms in recent years). Even within a call, managers may have incentives to accentuate or downplay the salience of geoeconomic frictions depending on anticipated market reactions. These strategic reporting choices inject a source of measurement error that is not addressed by our prompt- or model-level perturbations. Diversifying the range of text sources allows us to obtain different perspectives from authors that are not subject to the same set of reporting incentives, mitigating these concerns. The present version of the paper includes analyst reports as well, and in ongoing work we are exploring further text sources. As an example, Figure 13 below reports the results of Figure 2 but restricting the text to be purely coming from single-firm analysts reports. We are reassured that, at least at this broad level, different types of text, fixing model and prompt, generate broadly similar patterns.

Second, both earnings calls and analyst reports may be subject to a novelty bias. The impact of a particular instance of geoeconomic pressure is most likely to be reported in the corpus when it first occurs, but it is possible that both the firm managers and the analysts may stop discussing a particular policy if it is not newsworthy anymore, even though it may continue to affect the firm across multiple quarters. As a consequence, our time series results should be interpreted keeping in mind this property of the data.

Third, regarding firm responses, we take no stance on whether the firms or the analysts are correctly assessing how policy changes will affect them nor whether they are truthfully reporting how they are considering to react. We are interested in whether our LLM approach correctly identifies what the firms and analysts are reporting. However, we note that a voluminous previous literature has shown that information extracted from this text is actually related to firm actions.

Computational practices, reproducibility, and data confidentiality. We conclude this section with a set of additional remarks outlining the principles we followed in our computational approach. Our computational design is built to optimize for scientific reproducibility and data confidentiality. We rely on open-weights models precisely because they balance the needs of reproducibility and confidentiality: by executing all inference locally, we ensure that we can both hold the exact environment fixed (including randomization features such as random seeds and scheduling) and that the data always remains on our secure servers. This differs substantially from uploading documents to an online API of a closed-weight model. As mentioned earlier in the paper, we guarantee computational reproducibility by holding the LLM weights and the compute hardware fixed. Our sampling strategy is deterministic: we use zero-temperature sampling (with a fixed random seed to break any ties in the mode of the next-token distribution). At the same time, by quantizing the models and distributing the compute load across multiple research-grade GPUs, we are able to keep compute times manageable.

Figure 13: **Geoeconomic pressure: aggregate trends with analysts reports**

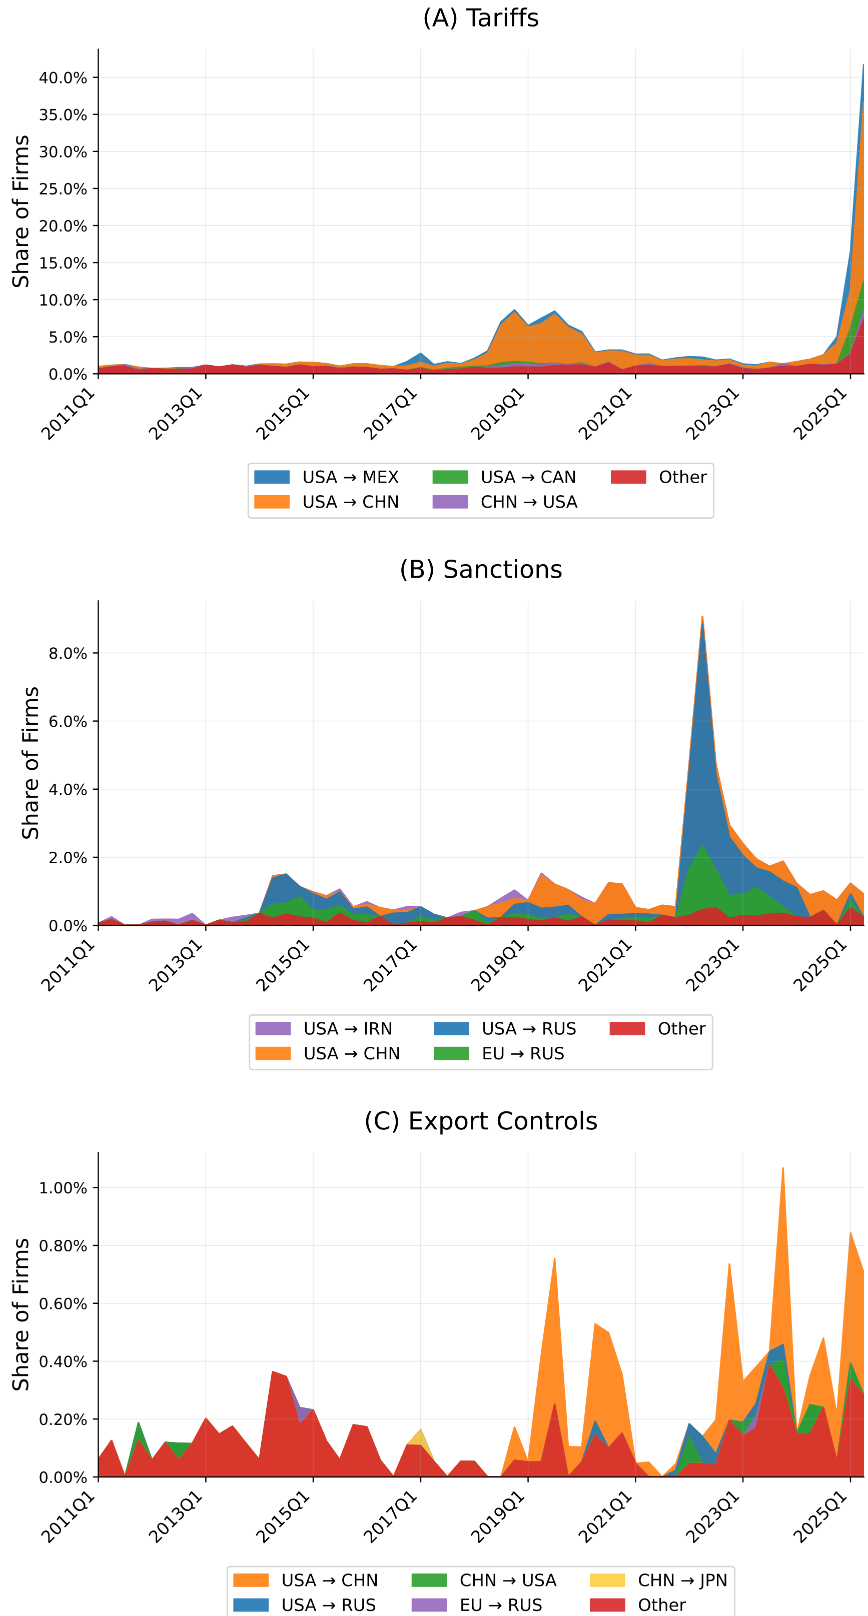

Notes: This figure plots the share of firms discussing each of the tools in a given quarter inferred by single-firm analysts reports. This provides a counterpart to Figure 2 but with a fully different dataset.

7 Conclusion

This paper develops a systematic methodology leveraging large language models to measure and analyze the growing role of geoeconomic pressure in global economic relations. By leveraging advanced NLP techniques and large language models, we construct a comprehensive database of coercive economic actions and link them to firm-level economic responses. We document a substantial rise in the frequency and variety of coercive economic actions, including tariffs, sanctions, and export controls. Using firm-level textual data from earnings calls and analyst reports, we find that firms significantly adjust their strategic decisions—such as pricing, investment, and supply chain structures—in response to both implemented and threatened economic pressure. Our results reveal important heterogeneity across industries and countries, with certain sectors and firms in targeted nations particularly vulnerable. An important advantage of our approach is the ability to assess the effects in near real-time and we intend to update our analysis regularly as the global trade war continues to unfold.

References

Ahn, Daniel P and Rodney D Ludema, "The sword and the shield: The economics of targeted sanctions," *European Economic Review*, 2020, *130*, 103587.

Aiyar, Shekhar, Andrea Presbitero, and Michele Ruta, *Geoeconomic Fragmentation the Economic Risks from a Fractured World Economy*, Centre for Economic Policy Research, 2023.

Aiyar, Shekhar, Davide Malacrino, and Andrea F Presbitero, "Investing in friends: The role of geopolitical alignment in fdi flows," *European Journal of Political Economy*, 2024, p. 102508.

Alekseev, Maxim and Xinyue Lin, "Trade policy in the shadow of conflict: the case of dual-use goods."," 2024.

Alfaro, Laura and Davin Chor, "Global supply chains: The looming "great reallocation"," Technical Report, National Bureau of Economic Research 2023.

Alfaro, Laura, Harald Fadinger, Jan S Schymik, and Gede Virananda, "Trade and Industrial Policy in Supply Chains: Directed Technological Change in Rare Earths," Technical Report, National Bureau of Economic Research 2025.

Amiti, Mary, Stephen J Redding, and David E Weinstein, "The impact of the 2018 tariffs on prices and welfare," *Journal of Economic Perspectives*, 2019, *33* (4), 187–210.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "The Macroeconomics of Tariff Shocks," Technical Report, National Bureau of Economic Research 2025.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring economic policy uncertainty," *The quarterly journal of economics*, 2016, *131* (4), 1593–1636.

Baldwin, David A, *Economic Statecraft*, Princeton University Press, 1985.

Battaglia, Laura, Timothy Christensen, Stephen Hansen, and Szymon Sacher, "Inference for Regression with Variables Generated by AI or Machine Learning," 2024.

Becko, John Sturm and Dangiel G. O'Connor, "Strategic (Dis)Integration," 2024.

Blackwill, Robert D and Jennifer M Harris, *War by other means: Geoeconomics and statecraft*, Harvard University Press, 2016.

Bonadio, Barthélémy, Zhen Huo, Elliot Kang, Andrei A. Levchenko, Nitya Pandalai-Nayar, Hiroshi Toma, and Petia Topalova, "Playing with blocs: Quantifying decoupling," 2024.

Bound, John and Alan B Krueger, "The extent of measurement error in longitudinal earnings data: Do two wrongs make a right?," *Journal of labor economics*, 1991, *9* (1), 1–24.

Broner, Fernando, Alberto Martin, Josefin Meyer, and Christoph Trebesch, "Hegemonic Globalization," 2024.

Broner, Fernando, Alberto Martin, Josefin Meyer, Christoph Trebesch, and Jiaxian Zhou Wu, "Hegemony and International Alignment," 2025.

Bybee, J Leland, "The ghost in the machine: Generating beliefs with large language models," *arXiv preprint arXiv:2305.02823*, 2023.

Caldara, Dario and Matteo Iacoviello, "Measuring geopolitical risk," *American Economic Review*, 2022, *112* (4), 1194–1225.

Caldara, Dario, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo, "The economic effects of trade policy uncertainty," *Journal of Monetary Economics*, 2020, *109*, 38–59.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, "Tariff pass-through at the border and at the store: Evidence from us trade policy," *American Economic Review: Insights*, 2021, *3* (1), 19–34.

Chen, Yifei, Bryan T Kelly, and Dacheng Xiu, "Expected returns and large language models," *Available at SSRN 4416687*, 2022.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "A Framework for Geoeconomics," Working Paper 31852, National Bureau of Economic Research November 2023.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "A Theory of Economic Coercion and Fragmentation," *Available at SSRN 4767131*, 2024.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "Putting Economics Back Into Geoeconomics," 2025.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "The Political Economy of Geoeconomic Power," 2025.

Crosignani, Matteo, Lina Han, Marco Macchiavelli, and André F Silva, "Geopolitical Risk and Decoupling: Evidence from US Export Controls," *FRB of New York Staff Report*, 2024, (1096).

Dell, Melissa, "Deep learning for economists," Technical Report, National Bureau of Economic Research 2024.

Dreher, Axel, Andreas Fuchs, Bradley Parks, Austin Strange, and Michael J Tierney, *Banking on Beijing: The Aims and Impacts of China's Overseas Development Program*, Cambridge University Press, 2022.

Drezner, Daniel W, "The hidden hand of economic coercion," *International organization*, 2003, *57* (3), 643–659.

Drezner, Daniel W, "Global economic sanctions," *Annual Review of Political Science*, 2024, *27.*

Drezner, Daniel W, Henry Farrell, and Abraham L Newman, *The uses and abuses of weaponized interdependence*, Brookings Institution Press, 2021.

Egorov, Konstantin, Vasily Korovkin, Alexey Makarin, and Dzhamilya Nigmatulina, "Trade Sanctions against Russia: Stylized Facts," in "AEA Papers and Proceedings," Vol. 115 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2025, pp. 578–582.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The return to protectionism," *The Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Fang, Hanming, Ming Li, and Guangli Lu, "Decoding China's Industrial Policies," Technical Report, National Bureau of Economic Research 2025.

Farrell, Henry and Abraham L Newman, "Weaponized interdependence: How global economic networks shape state coercion," *International security*, 2019, *44* (1), 42–79.

Farrell, Henry and Abraham Newman, *Underground Empire: How America Weaponized the World Economy*, Henry Holt and Co., 2023.

Felbermayr, Gabriel, Aleksandra Kirilakha, Constantinos Syropoulos, Erdal Yalcin, and Yoto V Yotov, "The global sanctions data base," *European Economic Review*, 2020, *129*, 103561.

Fernández-Villaverde, Jesús, Tomohide Mineyama, and Dongho Song, "Are We Fragmented Yet? Measuring Geopolitical Fragmentation and its Causal Effects," 2024.

Fernandez-Villaverde, Jesus, Yiliang Li, Le Xu, and Francesco Zanetti, "Charting the Uncharted: The (Un) Intended Consequences of Oil Sanctions and Dark Shipping," Technical Report, CESifo Working Paper 2025.

Flynn, Joel P and Karthik Sastry, "The macroeconomics of narratives," *Available at SSRN 4140751*, 2022.

Flynn, Joel P., Antoine Levy, Jacob Moscona, and Mai Wo, "Foreign Political Risk and Technological Change," 2025.

Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice, "Tariff-based product-level trade elasticities," *Journal of International Economics*, 2022, *137*, 103593.

Freund, Caroline, Aaditya Mattoo, Alen Mulabdic, and Michele Ruta, "Is US trade policy reshaping global supply chains?," *Journal of International Economics*, 2024, *152*, 104011.

Gentzkow, Matthew and Jesse M Shapiro, "What drives media slant? Evidence from US daily newspapers," *Econometrica*, 2010, *78* (1), 35–71.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, "Text as data," *Journal of Economic Literature*, 2019, *57* (3), 535–574.

Goldberg, Pinelopi K, Réka Juhász, Nathan J Lane, Giulia Lo Forte, and Jeff Thurk, "Industrial policy in the global semiconductor sector," Technical Report, National Bureau of Economic Research 2024.

Gopinath, Gita, Pierre-Olivier Gourinchas, Andrea F Presbitero, and Petia Topalova, "Changing global linkages: A new Cold War?," *Journal of International Economics*, 2025, *153*, 104042.

Grossman, Gene M, Elhanan Helpman, and Stephen J Redding, "When tariffs disrupt global supply chains," *American Economic Review*, 2024, *114* (4), 988–1029.

Hakobyan, Shushanik, Sergii Meleshchuk, and Robert Zymek, "Divided We Fall: Differential Exposure to Geopolitical Fragmentation in Trade," Technical Report, International Monetary Fund 2023.

Hassan, Tarek A, Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun, "Sources and transmission of country risk," *Review of Economic Studies*, 2024, *91* (4), 2307–2346.

Hassan, Tarek A, Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun, "Firm-level political risk: Measurement and effects," *The Quarterly Journal of Economics*, 2019, *134* (4), 2135–2202.

Hassan, Tarek Alexander, Stephan Hollander, Aakash Kalyani, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun, "Economic Surveillance using Corporate Text," Technical Report, National Bureau of Economic Research 2024.

Hirschman, Albert O, *National power and the structure of foreign trade*, Vol. 105, Univ of California Press, 1945.

Ignatenko, Anna, Ahmad Lashkaripour, Luca Macedoni, and Ina Simonovska, "Making America Great Again? The Economic Impacts of Liberation Day Tariffs," Technical Report, National Bureau of Economic Research 2025.

Itskhoki, Oleg and Dmitry Mukhin, "Sanctions and the exchange rate," Technical Report, National Bureau of Economic Research 2022.

Juhász, Réka, Nathan Lane, Emily Oehlsen, and Verónica C Pérez, "The Who, What, When, and How of Industrial Policy: A Text-Based Approach," *What, When, and How of Industrial Policy: A Text-Based Approach (August 15, 2022)*, 2022.

Kalemli-Özcan, Şebnem, Can Soylu, and Muhammed A Yildirim, "Global Networks, Monetary Policy and Trade," Technical Report, National Bureau of Economic Research 2025.

Kapfer, Craig, Kurt Stine, Balasubramanian Narasimhan, Christopher Mentzel, and Emmanuel Candes, "Marlowe: Stanford's GPU-based Computational Instrument," January 2025.

Keerati, Ritt, "The unintended consequences of financial sanctions," *Available at SSRN 4049281*, 2022.

Keohane, R.O. and J.S. Nye, *Power and Interdependence*, Little Brown, 1977.

Kindleberger, Charles Poor, *The world in depression, 1929-1939*, Univ of California, 1973.

Kleinman, Benny, Ernest Liu, and Stephen Redding, "International Friends and Enemies," *American Economic Journal: Macroeconomics*, 2024.

Kooi, Oliver, "Power and Resilience: An Economic Approach to National Security Policy," *working paper*, 2024.

Lagakos, David, Stelios Michalopoulos, and Hans-Joachim Voth, "American Life Histories," Technical Report, National Bureau of Economic Research 2025.

Liu, Ernest and David Yang, "International Power," 2024.

Ludwig, Jens, Sendhil Mullainathan, and Ashesh Rambachan, "Large language models: An applied econometric framework," *Working Paper*, 2025.

Mattoo, Aaditya, Michele Ruta, and Robert W. Staiger, "Geopolitics and the World Trading System," 2024.

Mayer, Thierry, Isabelle Mejean, and Mathias Thoenig, "The fragmentation paradox: De-risking trade and global safety," Technical Report, working paper 2025.

Mohr, Cathrin and Christoph Trebesch, "Geoeconomics," 2024.

Morgan, T Clifton, Navin Bapat, and Valentin Krustev, "The threat and imposition of economic sanctions, 1971—2000," *Conflict Management and Peace Science*, 2009, *26* (1), 92–110.

Nigmatulina, Dzhamilya, "Sanctions and misallocation. How sanctioned firms won and Russia lost," 2022.

Ottonello, Pablo, Wenting Song, and Sebastian Sotelo, "An Anatomy of Firms' Political Speech," Technical Report, National Bureau of Economic Research 2024.

Pflueger, Carolin and Pierre Yared, "Global Hegemony and Exorbitant Privilege," 2024.

Pischke, Jörn-Steffen, "Measurement error and earnings dynamics: Some estimates from the PSID validation study," *Journal of Business & Economic Statistics*, 1995, *13* (3), 305–314.

Rodríguez-Clare, Andrés, Mauricio Ulate, and Jose P Vasquez, "The 2025 trade war: Dynamic impacts across US States and the global economy," Technical Report, National Bureau of Economic Research 2025.

Sarkar, Suproteem K., "Economic Representations," jan 2025.

Schmillen, Achim, Matthias Umkehrer, and Till von Wachter, "Measurement error in longitudinal earnings data: evidence from Germany," *Journal for Labour Market Research*, 2024, *58* (1), 8.

Thoenig, Mathias, "Trade Policy in the Shadow of War: Quantitative Tools for Geoeconomics," 2023.

Utar, Hâle, Luis Bernardo Torres Ruiz, and Alfonso Cebreros Zurita, "The US-China trade war and the relocation of global value chains to Mexico," 2023.