

Geoeconomic Pressure

Christopher Clayton* Antonio Coppola[†] Vanessa Hu[‡]
Matteo Maggiori[§] Jesse Schreger[¶] Serdil Tinda^{||}

September 2026

Abstract

The increasing use of geoeconomic pressure is reshaping the global economy. We introduce a methodology using large language models (LLMs) to systematically measure the application of and response to geoeconomic pressure from large textual corpora, classifying which governments apply pressure to which foreign targets, using which instruments, firms, and products. Most importantly, we use LLMs to analyze how and why firms respond to the pressure. We focus on export controls, sanctions, and tariffs as the main instruments of pressure on firms. We find that export controls induce the targeted firms to undertake domestic R&D to build a domestic alternative technology. Western sanctions on Russia cause supply chain re-arrangements, the major beneficiary of which are firms in third-party countries, like China and India, which step into the market vacuum created by the broken business relationships. The threat and impositions of U.S. tariffs cause losses for the majority of U.S. firms, which report absorbing the cost shock on their inputs via their profit margin. Foreign firms are also negatively affected, mostly via the related macroeconomic uncertainty and the cost of relocating production.

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

JEL Codes: C4, F3, F4, F5, G3.

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

[†]University of Chicago Booth School of Business, NBER, and CEPR; acoppola@stanford.edu.

[‡]Columbia Business School; JHu29@gsb.columbia.edu.

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

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

^{||}Stanford University, Department of Economics; serdil@stanford.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, Oliver Wang, 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 in the exertion of geoeconomic pressure 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, the United States has been running a global campaign to restrict the ability of firms around the world to sell cutting-edge semiconductors to China and its allies. China used a threat to suspend rare earth exports to the United States to force U.S. trade concessions. In addition, Russia has repeatedly threatened to suspend natural gas shipments to Europe in response to European sanctions on Russia.

As these events drive major shifts in markets, the real economy, and the international policy environment, academics and policymakers are trying to understand the implications of 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. The immediate actor being pressured might not be the one bearing the cost: for example, the incidence of a tariff might be on the domestic firms, domestic consumers, or foreign producers. In many policy applications, the availability of near real-time information also carries 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.

The key idea of the paper is to use LLMs to measure how and why firms respond to various forms of geoeconomic pressure. A critical advantage of LLMs over classical natural language processing is their ability to extract subtle conditional relationships in the prose. While a keyword-based approach can assess whether a particular document discusses tariffs and whether it also mentions investment, it takes a human or an LLM to determine whether the document is stating that the firm is undertaking this investment *because of* a change in tariff policy and to summarize the rationale for the investment. When a CEO states “we are building a new factory in Vietnam in order to mitigate the costs of tariffs,” human readers interpret this as the firm attributing a causal link between tariff policy and supply chain restructuring. Our LLM pipeline aims to capture this degree of subtlety and reported causality while leaving human readers out of the loop to massively increase the scale of the analysis.

We term the objects we aim to measure the “Subjective Causal Effect” (SCE) to make clear that these are the firm’s top management’s own causal attributions rather than econometrically identified causal effects. As we show, these self-reported mechanisms contain substantial information about firm behavior that is not available from standard data sources.

Similarly, while classic NLP sentiment analysis is effective in assessing who is positive or negative about an issue, it is difficult to use standard techniques to answer *why* a firm is positively or negatively impacted by a given event. We develop a reasoning pipeline in which we use LLMs to categorize the reported reasons behind the firm’s actions. This pipeline has three steps. In the first step the LLM extracts free-text describing the reason behind firms’ actions. In the second step the LLM inductively discovers from the text produced by the first step a set of categorical reasons. The output of the second step is a codebook of categories. In the third step the LLM matches the universe of text extracted in the first step to the codebook. The output is one or more categories that best describe the reason behind each firm action at that point in time. This allows us not only to document that there is heterogeneity in the studied effects across firms, but to associate this heterogeneity with the underlying reasoning in the text. We refer to this approach as the “Reasoning” analysis.

We collect the facts emerging from our analysis of the subjective causal effects and reasons behind firms’ actions around three core instruments of pressure: export controls, sanctions, and tariffs. For export controls, we document a strong reaction of firms targeted by the controls: they undertake domestic R&D and they report the rationale to be creating a domestic alternative that insulates them from foreign geoeconomic pressure. This reaction is particularly prominent among Chinese firms targeted by U.S. export controls in semiconductors and AI. Indeed, in theoretical geoeconomic models, this R&D investment to create a domestic alternative is the major drawback of tightening export controls from the perspective of the country exerting the pressure.

For sanctions, of which those imposed on Russia are the most prominent in our sample, we find two distinct reactions. For firms in the imposing countries, such as the U.S. and Europe, the sanctions impose direct costs due to loss of revenue in Russia and the seizing of assets these companies had in Russia. Increases in energy costs as a result of the sanctions are a secondary source of losses for European firms, but often associated with gains by U.S. firms. This is consistent with European firms largely substituting away from Russian energy and toward energy provided by U.S. firms, which benefit as a result. Firms located in third-party countries, especially India and China, have a markedly different experience. These countries not only did not participate in the sanctions, but actively increased their business ties with Russia. Firms in these countries report benefiting from being able to fill in the market-vacuum created by the sanctions. This includes not only doing business with Russian firms, but also doing new business with Western firms that previously had an established relationship with a Russian firm. Indian firms also specifically mention the benefit of being able to buy Russian crude oil at a discount compared to world market prices for oil.

The use of tariffs, or threats of tariffs, in our sample is largely driven by the U.S. administration’s trade war with China in 2018, and a generalized rethink of tariff policy against all trading partners starting with the “liberation day” announcement in April 2025. Focusing first on U.S. firms confronting U.S. imposed tariffs we find two prominent patterns. 73% of affected firms report being negatively affected by the tariffs. 58% of the negatively affected firms report that tariffs increased their input prices. This is consistent with the tariffs being levied on intermediate inputs purchased

by U.S. firms and the foreigners not cutting their prices sufficiently to offset the tariff. These firms also report increasing their sale prices to pass through some of the tariff cost to their customers, but they also report a compression in profit margins, so that the firms in part absorb the cost shock. The U.S. firms that report being positively affected by the tariffs behave markedly differently: they do not face higher input prices, do not increase their prices (as much), but report increasing their sales and profit margins. This is consistent with these firms being domestic producers that benefit from the tariffs imposed on their competitors, whether a foreign firm or another domestic firm that was sourcing inputs from abroad. As their competitors increase prices in part to offset the cost shock of the tariffs, these firms are able to grab market share in the domestic market. The effects we document are consistent with tariffs being in part absorbed by U.S. importing firms, in part falling on U.S. consumers, and being a subsidy to some domestic producers.

Firms in the rest of the world whose exports to the U.S. are subject to tariffs also report being mostly negatively affected via lower sales, a compression in profit margins, and bearing the costs of re-localizing production to minimize the impact of tariffs. European and Canadian firms mention more frequently the macroeconomic and financial risk associated with the tariffs, a general equilibrium deterioration of the economy, compared to Chinese firms that focus more on the cost of relocating production. These foreign firms also mention tariff-driven price concessions and margin compression, and their U.S. subsidiaries paying the tariffs. The effects are also commingled with these firms attributing costs to retaliatory tariffs imposed by their own country, or other countries, that the firms perceive would not have been there had the U.S. not initiated a tariff war. Overall, we find that the foreign firms are negatively affected, mostly via uncertainty, lower demand, relocation costs, and their local U.S. subsidiaries paying the tariffs.

Our paper brings to bear on geoeconomics the frontier, as it stands, of LLMs usage as a classifier and for open-ended text. This new territory of research is fast expanding and evolving and it brings up new challenges and opportunities for economic research. We discuss in the paper a number of challenges in ongoing work. Many of these challenges are not specific to LLMs, and instead also present in surveys and traditional NLP. For example, when our classifier reports that a CEO said that she took an action in response to export controls, a researcher might be interested in the correctness of the classification at different levels. One level is whether the LLM has correctly classified the text: “did the CEO actually say what the classifier reports?”. This brings up measurement error and bias in LLMs as classifiers. A different level is whether the CEO truthfully reported her beliefs, actions, and attribution of causality. One can reasonably worry about strategic misreporting, cheap talk, or errors in reasoning by the CEO.

In the final section of the paper, and in ongoing work, we tackle these issues. Focusing on LLMs as classifiers, we explore: (i) the connection between our measures and observable firm characteristics (ii) narrow computational reproducibility using open-weights models, greedy sampling, and a local inference infrastructure, (iii) repeating the analysis with different models and (iv) correcting for potential bias in the downstream econometric analysis. Focusing on the earnings calls text and potential for biased reporting by the CEOs, we perform two robustness checks: (i) we assess some

verifiable outcomes versus external hard data, (ii) we repeat the analysis using financial analyst reports on the same firms. Professional analysts are external to the firm and provide a different view of the same events.

The remainder of the paper proceeds as follows. Section 2 describes our textual data and LLM methodology. Section 3 documents the aggregate trends in the three instruments. Section 4 presents our analysis of export controls; Section 5 examines sanctions; Section 6 analyzes tariffs and the ongoing 2025 trade war. Section 7 discusses measurement uncertainty and reproducibility. Section 8 concludes.

Related Literature. Our paper relates methodologically to the literature on natural language processing (NLP) in economics, the literature on open-ended surveys, and in particular to the nascent literature using LLMs as classifiers. The economic focus of the paper relates to the literature on geoeconomics.

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. \(2016\)](#).¹ In particular, we relate to the pioneering work on firm-level uncertainty using NLP by [Hassan et al. \(2019\)](#).² [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. \(2025\)](#), [Fang et al. \(2025\)](#), and [Asirvatham et al. \(2026\)](#).³

Methodologically, our Subjective Causal Effect and Reasoning analysis parallels the large literature on open-ended surveys in economics (see [Stantcheva \(2023\)](#) for an overview). The closest connection is with the literature that aims to extract quantitative measures from open-ended qualitative survey data as discussed in [Haaland et al. \(2025\)](#). As discussed by [Stantcheva \(2023\)](#), LLMs offer the opportunity both to help discover coding schemes and implement the schemes on a large scale. While the survey literature directly asks respondents the questions they want the open-ended text for, we instead have LLMs extract such open-ended text and then adopt the inductive category discovery. Our Reasoning analysis is essentially replacing the human induction step of Thematic Analysis in psychology and other social sciences with LLM discovery of the themes, categories and reasons ([Braun and Clarke \(2006\)](#)). More generally, our paper shares the aim of the firm survey literature in macroeconomics (e.g. [Coibion et al. \(2018\)](#) and [Coibion and Gorodnichenko \(2026\)](#)) while doing it at scale using publicly available corporate text to extract subjective beliefs rather

¹See also [Gentzkow et al. \(2019\)](#) for a review of the literature.

²For additional NLP at the firm-level see [Hassan et al. \(2024b\)](#), [Hassan et al. \(2024a\)](#), and [Flynn and Sastry \(2022\)](#).

³See [Dell \(2024\)](#) for a review of deep learning methods for economists.

than a custom survey.

We contribute to the fast-growing modern literature on geoeconomics, including: [Clayton et al. \(2026\)](#), [Clayton et al. \(2024\)](#), [Thoenig \(2023\)](#), [Fernández-Villaverde et al. \(2024\)](#), [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. \(2025\)](#), [Becko et al. \(2025\)](#), and [Pflueger and Yared \(2024\)](#), [Madsen and Prat \(2026\)](#).⁴

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 ([Drezner \(2003\)](#); [Morgan et al. \(2009\)](#); [Drezner \(2024\)](#)). Recent papers in economics include [Ahn and Ludema \(2020\)](#), [Nigmatulina \(2022\)](#), [Keerati \(2022\)](#), [Itskhoki and Mukhin \(2022\)](#), [Fernandez-Villaverde et al. \(2025\)](#), [Egorov et al. \(2025\)](#), and [Matvos and Neiman \(2026\)](#). [Felbermayr et al. \(2020\)](#) assembled a bilateral database of sanctions around the world. A related literature focuses on export controls: [Flynn et al. \(2025\)](#), [Fan et al. \(2026\)](#), [Crosignani et al. \(2024\)](#), and [Liu et al. \(2025\)](#).

Finally, an voluminous literature focuses on the economic effect of tariffs: [Amiti et al. \(2019\)](#), [Fajgelbaum et al. \(2020\)](#), [Cavallo et al. \(2021\)](#), [Grossman et al. \(2021\)](#), [Cavallo et al. \(2021\)](#), [Amiti et al. \(2026\)](#), [Gopinath and Neiman \(2026\)](#), [Flaaen et al. \(2025\)](#), [Bai et al. \(2025\)](#), [Ganapati and Hottman \(2026\)](#).⁵

2 Data and Methodology

2.1 Textual Data

In order to measure geoeconomic pressure at the firm level, our essential input is firm-level text. We rely primarily on earnings calls and supplement these with analyst reports in robustness checks. As discussed in detail in [Hassan et al. \(2024b\)](#), earnings calls of publicly traded firms are a valuable source of textual information because they allow market participants to ask firm leadership questions of their choosing. While the calls begin with a prepared presentation from management, often the CEO or CFO of the firm, market participants on the call then have the opportunity to ask about the issues that they find most pressing.

One part of our dataset is the global earnings call transcripts from Capital IQ via Wharton Research Data Services (WRDS).⁶ These are the transcripts most often used in academic and industry research, but this corpus of transcripts is dominated by US firms. For example, in 2025

⁴See also [Aiyar et al. \(2024\)](#), [Gopinath et al. \(2025\)](#), [Hakobyan et al. \(2023\)](#), [Ndiaye \(2024\)](#), [Aiyar et al. \(2023\)](#), [Bonadio et al. \(2024\)](#), [Broner et al. \(2025\)](#), and [Clayton et al. \(2025\)](#). [Mohr and Trebesch \(2024\)](#) surveys the burgeoning literature.

⁵See [Ossa and Redding \(2026\)](#) for a literature review.

⁶In ongoing work we are maximizing the coverage of this corpus by adding earnings calls covered globally by Refinitiv and FactSet, and by S&P SCRIPTS Asia Transcripts for Japanese firms.

they contain transcripts from 3253 US firms, 195 Chinese firms, and 5784 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 geoeconomic 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 2008-2026 (approximately 247 thousand documents).

Table 1 provides summary statistics for the panel of transcripts used in this paper. The top panel shows that coverage does increase over time, especially for firms outside of the US. In 2025, the sample includes 59 thousand transcripts, from over 14 thousand firms globally. In total we cover 624 thousand transcripts, from approximately 22 thousand unique firms across all years. The bottom panel provides a cross-section of the country coverage for the year 2024. A well-known pattern in this corpus of text (see for example Hassan et al. (2024a)), is that coverage in each country is skewed towards the largest firms. These firms are both more likely to hold earnings calls, and the datasets are more likely to cover them. For example, while the data covers only 3,397 firms in the US, the firms we match to the equity universe account approximately for 96 percent of the total stock market capitalization for that country. The skew toward larger firms, both in terms of being listed and among the listed firms, is not necessarily problematic for the purpose of this paper since large multinationals have a disproportionate imprint on cross-border activity.⁷

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 large for the target country, are not salient enough to be discussed on the earnings calls of major multinationals. This issue is likely to be particularly important for smaller developing countries. In the event of other important development (such as the start of the COVID pandemic), executives may not discuss and analysts might not ask about geoeconomic pressure even if the firm is affected. Third, CFOs and CEOs might be reluctant to candidly discuss pressure being applied by their own, or even foreign, governments. To address some of these concerns, we turn to analyst reports at both the firm and sector level as a complementary source of text in the robustness tests.⁸

⁷For large firms the country of nationality can differ from the country of incorporation, often because of the parent company being domiciled in a tax-favorable jurisdiction (Coppola et al. (2021)). We assign the country for each listed firm by matching the transcript to a FactSet entity ID for the firm. We then employ a simplified version of the algorithm in Coppola et al. (2021) to determine nationality: we use country of domicile in FactSet, if that country is on the tax-haven list in Coppola et al. (2021), we then chose the country of business nationality from Dealogic as long as that country is not on the same list. This provides a nationality view of the equity market. Importantly, we do not focus on where the firm is listed. Alibaba, for example, is a Chinese firm in our data, despite being listed in New York and Hong Kong and incorporated in the Cayman Islands.

⁸We use research reports from professional financial analysts at the industry-country, or individual firm level. One limitation is that the sample of single-firm analyst reports, like the sample of earnings calls from WRDS, is heavily skewed towards large publicly traded firms from Western markets.

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

(a) Corpus by year

Year	Transcripts			Firms		
	USA	CHN	RoW	USA	CHN	RoW
2010	11,113	916	2,329	3,274	715	1,101
2015	11,940	10,879	5,627	3,193	2,321	1,768
2020	12,520	13,875	13,571	3,354	3,643	4,423
2025	12,228	29,359	17,868	3,253	5,243	5,786
Total (2008–2026)	217,898	246,797	159,149	7,469	5,865	8,485

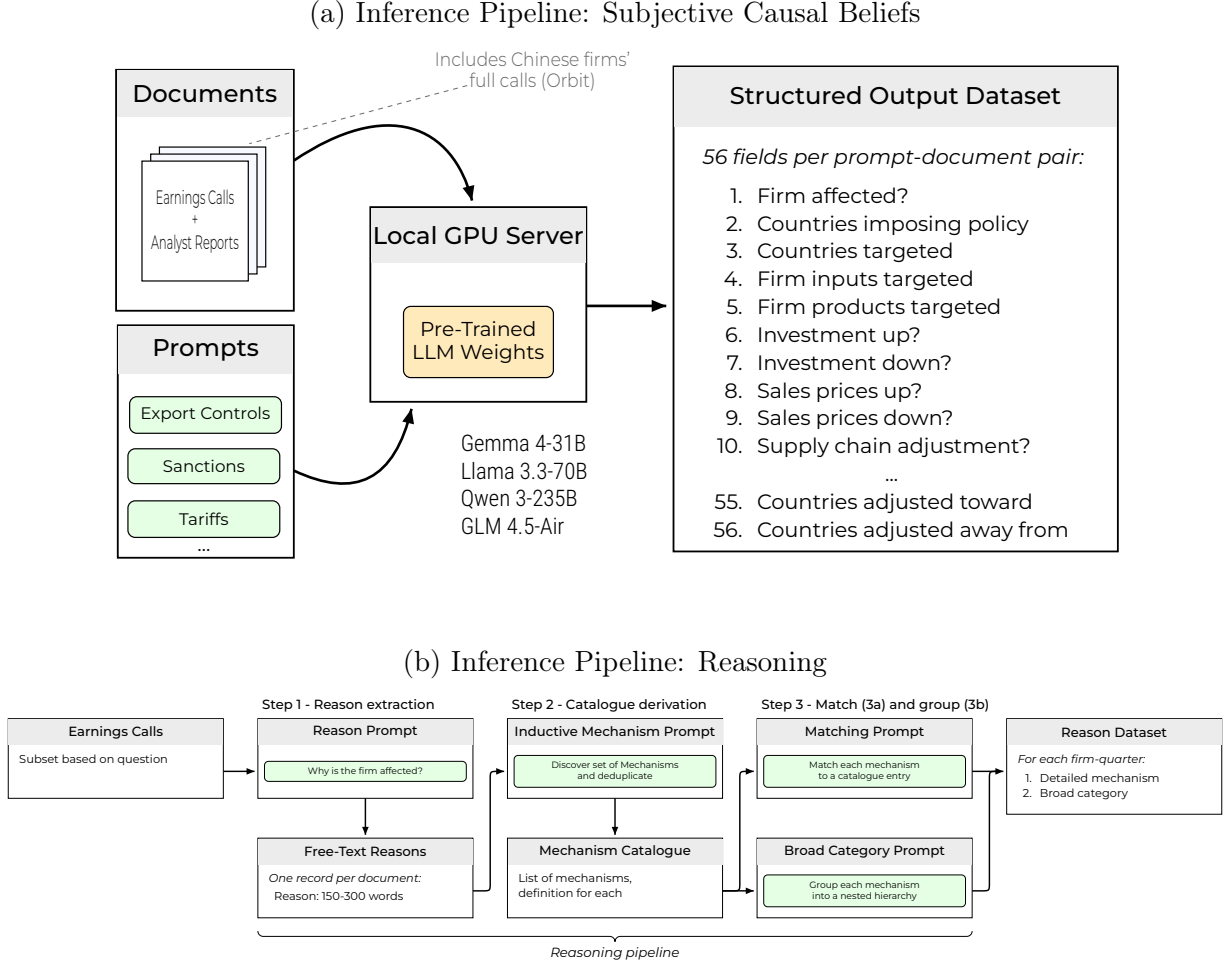
(b) Coverage of the Equity Market universe, 2024

Country	Firms	% of market cap.
United States	3,397	96.3
China	5,145	90.1
Japan	157	47.3
India	931	84.4
United Kingdom	620	89.8
France	144	90.7
Canada	538	91.5
Saudi Arabia	13	78.0
Taiwan	50	55.1
Germany	221	93.9
Switzerland	135	93.8
Australia	445	86.4
South Korea	57	47.2
Brazil	208	93.4
All countries	14,001	86.0

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. Panel (b) focuses on firm coverage in selected countries in terms of stock market capitalization.

2.2 Methodology

Figure 1: Illustration of LLM-based Classification Methodology



Notes: This figure illustrates our inferences pipelines using LLMs. In Panel (a) illustrates the inference pipeline used to classify whether firms were affected by the pressure and how they reacted. Panel (b) illustrates the inference pipeline used to classify why the firms were positively or negatively affected.

Our methodology performs large-scale inference on the textual corpora using pre-trained, open-source large language models (LLMs). All the baseline results shown in the paper use the 31 billion parameter version of Gemma 4 released by Google. In addition, we have repeated the analysis with Meta’s Llama 3.3 70 billion parameter model, GLM 4.5 Air from Z.ai, Alibaba’s 235 billion parameter Qwen 3 235B.

Our inference pipeline has two main branches. The first branch takes the earnings calls text and applies an LLM to classify each text into a predetermined set of questions. This branch aims to extract whether the management reports taking certain action in response to the geoeconomic

pressure. The second branch takes the same text and applies LLM to classify the reasons why firms took those actions. The reasoning classification is inductive, without pre-specified categories. We describe in detail each inference pipeline below.

2.3 Subjective Causal Effect Pipeline

We structure this inference task using a two-part prompt design. First, we run a broad analysis in which for each earnings call we extract 5 booleans for whether a firm is affected by each of 5 forms of geoeconomic pressure: export controls, sanctions, tariffs, boycotts, and investment screening. Second, if any of these instruments flag at the document level in the broad screening, we pass the document through our detailed instrument-specific prompt. At present, we focus on the three most common tools of geoeconomic pressure: export controls, sanctions, and tariffs. In this second stage, we extract details such as the country imposing the policy, the target country of the policy, whether the firm is affected because of the policy targeting their inputs or products, and other characteristics. In addition, we extract a range of information about how the firm is responding to the pressure. For instance, we ask whether the firm is responding by raising its investment or cutting its investment, and if so, which countries? We extract similar information on supply chain adjustment, research and development, price changes, and a range of other responses.

Figure 1 Panel (a) illustrates the methodology. The LLM is a classifier that takes as input the prompt, the text, and returns a dataset with the classifications. Appendix A provides a list of the fields measured in our pipeline and the full text of all prompts is reproduced in Appendix B. The broad prompt is used to reduce the computational demand. For example, an LLM broad pass can quickly determine that a particular earnings call never talks about export controls, saving us the computational cost of passing that document through a detailed prompt about export controls.

2.4 Reasoning Pipeline

This pipeline is illustrated in Figure 1 Panel (b) and has three steps. In the first step, we begin by extracting open-ended responses using an LLM to the question of why a firm reports being affected by a policy (like tariffs) separately for firms that report being positively or negatively affected. Appendices B.5, B.6 and B.7 contain all the prompts used. The output of this first step consists of as many free-text answers as there are earnings calls selected as input. In the second step, we feed the output produced in the first step to an LLM that is prompted to produce a catalog of detailed mechanisms for why firms report particular behavior. A first pass of the LLM produces a list of reasons with no upper bound. A second pass consolidates the list by de-duplicating and avoiding mechanisms that are too similar. We set a limit to the number of mechanisms, but we do not specify the mechanisms themselves. Appendices B.8 and B.9 contain all the prompts used. The output of this second step is a catalog of mechanisms. In the third step, we ask the LLM to match the universe of open-ended text produced in stage 1 to the detailed mechanism from stage 2 that it most closely corresponds to. For ease of presentation, a final step groups together our detailed

mechanisms into broader hierarchical categories (i.e. the mechanism "Trade War Macroeconomic Growth Fear" is part of the broad category "Macroeconomic and Financial Risk"). Appendices B.10 and B.11 contain all the prompts used. The output of this third stage is a dataset that for each earning call in the reasoning analysis provides the matched entry in the fine reasoning catalog, and the fine reasoning are associated with coarser reasons that group together similar economic mechanisms.

This reasoning pipeline is used to dig deeper into particular episodes of geoeconomic pressure. For example, technology export controls, the US tariffs after April 2025, or the sanctions on Russia after 2022. As such, the questions and the sample of earnings calls are set to match the event, with details described in the corresponding analysis.

2.5 Computational Infrastructure

To ensure computational reproducibility of our analysis and respect data confidentiality, 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 also use Stanford’s Marlowe, a new GPU-based cluster (Kapfer et al. 2025). In order to minimize arbitrary noise in the LLM outputs, we set the sampling temperature to zero (i.e. greedy sampling). Finally, for the larger dense models, we deploy 4-bit quantized variants of the models using AWQ. This reduces the model’s inference memory needs, allowing us to preserve model quality and speed.

3 Aggregate Trends

We begin our analysis by documenting aggregate trends in the use of the three common forms of geoeconomic pressure: tariffs, sanctions, and export controls. In each case, we plot the share of global firms that report being affected by the instrument in question. We then decompose aggregate pressure into a bilateral measure in terms of which country is imposing the pressure and which one is receiving the pressure.⁹

Figure 2a plots the share of firms in earnings calls reporting that they are affected by export controls in a given quarter, decomposed by the sender and receiver of the pressure. Most export controls are imposed by the US and China targeting each other. The Euro Area countries mostly impose export controls toward Russia, likely as part of sanction packages. China also imposes export controls toward Japan. Overall, the number of firms worldwide that report being affected by export controls is around 1% of our sample, but we observe a striking increase in the share of firms

⁹Because not every instance of a firm classified as being affected by pressure reports a source and target of the pressure, for visual clarity we scale up proportionally the observations that we do observe to reconcile the bilateral with the aggregate. However, none of our analysis uses this scaling.

reporting being affected by export controls over time, with roughly a 4 fold increase since the 2010s. The low level is consistent with the notion of export controls targeting specific firms, materials, or countries. The time trend is consistent with the more widespread use of export controls in recent years.

Figure 2b plots the share of firms reporting being affected by sanctions. Two notable spikes occur in 2014 and 2022, both related to sanctions imposed on Russia by Western countries, most prominently the US and the European Union. These sanctions occurred in response to Russia’s invasion of Crimea in 2014 and broader invasion of Ukraine in 2022. Two other bilateral pairs stand out. First, US sanctions targeting China in the period 2019–2021, which are mostly sanctions applied on specific companies such as Huawei and ZTE. Second, the 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 withdrawal from the JCPOA. An earlier episode of sanctions targeting Iran in the buildup to the JCPOA in 2012 is also visible, but smaller in magnitude.

Figure 2c plots the share of firms reporting being affected by tariffs. The spike in tariff discussion in the first and second quarters of 2025 is the largest event (by share of global firms affected) in our sample. Prior to 2025, the largest spike in tariffs occurred during the US-China Trade War of 2018-2019. Looking at the combinations of countries imposing and receiving the tariffs, the first event in 2018–2019 is quantitatively driven by US tariffs imposed on China, with some retaliation of China imposing tariffs on the US. The second episode in 2025 is driven by the US (threatening to) imposing tariffs on China and a substantially broad set of other countries, including Canada and the Euro Area. Mentions of retaliatory tariffs in the 2025 are relatively muted.

The aggregate trends in Figure 2 offer a first-pass assessment of the LLM methods. The classifiers are indeed picking up well known examples of geoeconomic pressure. They also remind us of interpretational issues when assessing evidence extracted from text. CEOs and analysts might be more likely to discuss new rather than ongoing events. For example, the decrease in the number of firms that report being affected by tariffs after the peak in Q2 of 2025 reflects in part that CEOs were more likely to discuss tariffs in April 2025 than towards the end of the year (when the topic had already been covered in a number of earlier earnings calls).


In the next three sections we use our firm-level measurement to document how these three different forms of pressure caused firms to change their behavior.

4 Export Controls


We begin our instrument-by-instrument analysis with export controls. Figure 3 plots the firms response to export controls. We focus on the sample of firm-quarter observations in which firms report being affected by export control. Then, for each categorical response, we compute the fraction of firms that report an increase and or, separately, a decrease in that category. For example, for “profit margin” we compute the fraction of firms that report experiencing an increase in their profit margin in response to export controls. We plot the increases as positive bars in yellow. We separately

Figure 2: **Aggregate Trends**


(a) Export Controls



(b) Sanctions



(c) Tariffs



Notes: This figure plots the share of firms in earnings calls reporting that they are affected by each instrument in a given quarter. 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.

compute the fraction of firms that experience a decrease and report it in blue. The yellow and blue bars do not mechanically add up to 100 percent, since some firms might report neither or both of the actions for some categories. Figure 3 has 6 panels, with the left column summarizing the responses of affected firms based in the country imposing the export controls and the right column summarizing the response of affected firms in the country targeted by the export controls.

We begin by describing the top row of the figure that covers all global firms. As an example, consider U.S. export controls toward China. If NVIDIA reported being affected by these controls, it would be included in the left panel as a firm based in the country imposing the export controls. By contrast, if SMIC (Semiconductor Manufacturing International Corporation) were affected by the same controls, then it would be included in the right panel as a firm based in the country (China) targeted by the controls. Here, third-country companies (such as ASML, a Dutch firm, in the example) are excluded.


The difference between how firms in the source and target country respond to export controls is striking. In the imposing countries, firms are predominantly negatively affected and generally report lower sales. The response is a quantity adjustment: lower sales leading to a lower profit margin. This is consistent with export controls being predominantly quantity restrictions that negatively affect the private value of the firms producing the items that are being controlled.

We find a wider range of responses from firms in the target countries. First, in a mirror to the firms losing their ability to export, the firms in the targeted country are losing their ability to import the restricted items. 30% of firms correspondingly report reducing their inputs as a result of the export control. More than 20% of affected firms in the targeted countries report raising their investment in response to the controls, with nearly all of that increased investment reported to be domestic. Strikingly, nearly 40% of these affected firms report increasing their research and development when they are the target of export controls. This is consistent with the findings in Flynn et al. (2025) that sectors more exposed to foreign political risk conduct more innovative activity.

In the next two panels, we consider separately the response of U.S. firms and Chinese firms. U.S. and Chinese firms report reacting in similar ways to export controls imposed by their own government: lower sales, lower profit margin, and overall negatively affected. When firms are on the receiving end of the pressure, the most prominent pattern is a strong R&D response of Chinese firms. In both the U.S. and China, firms react to being the target of pressure with an increase in domestic investment and R&D, but the Chinese reaction is more prominent. In China, it is also notable that a substantial fraction of the firms report being positively affected, and we show below that this is because the absence of U.S. competitors is creating domestic market opportunities for some Chinese firms.

The export controls that China and the U.S. impose on each other are using and targeting different industries. Figure 4 is a Sankey chart of who applies export control pressure to whom using which sectors. For export controls, very specific products are being used, but which products


Figure 3: **Export controls: firm responses**



Notes: This figure plots the share of firm-quarters affected by export controls that report an increase (bars above zero) or a decrease (bars below zero) in each outcome. The sample is the firm-quarters in which a firm discusses the effect of export controls on its business. Panel (a) pools all affected firms; panels (b) and (c) restrict to U.S. and Chinese firms, respectively. Within each row, the left chart shows firms headquartered in a country imposing the export controls and the right chart firms headquartered in a country receiving them. Error bars denote 95 percent confidence intervals based on the standard error of each share across firm-quarters.

differ substantially based on who is the sender and receiver.¹⁰ 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.¹¹ The results in Figure 4 are consistent with what Jake Sullivan, the National Security Advisory to President Biden, called a “small yard, high fence” approach: severe export controls that targeted few countries, mostly China, in selected sectors, mostly advanced semiconductors and AI.

Figure 4: **Export controls: from who to whom using what**



Notes: This figure plots the flow of export controls from the imposing country (left) through the sector of the controlled product (center) to the receiving country (right). The sample is the firm-quarters where a firm discusses the effect of export controls on its business and names at least one imposing and one receiving country. Each firm-quarter contributes one unit for every imposing-receiving country pair it mentions, so flow widths are counts of country-pair mentions. The sector is the two-digit SIC industry of the product the firm describes as controlled, assigned by matching the product description to SIC industry names.

¹⁰To assign an industry to the product information extracted from our long prompts for export controls, we embed each product description and each official SIC4 industry description with the embeddings model “bge-large-en-v1.5” and take the SIC4 code with the highest cosine similarity, with keyword rules determining the match for common cases such as “chip” and “rare earth”. We then truncate the codes to two digits and group similar codes. For example, we group machinery, electronics, instruments, and glass to “Semiconductor Supply Chain”; we group mining and metals into “Metals and Minerals”.

¹¹See [Alfaro et al. \(2025\)](#) for a study of directed technical change in rare earths in response to shortages and export controls.

4.1 Understanding the R&D Response

While the results on the differential response of firms targeted by export controls are important in and of themselves, our LLM reasoning pipeline lets us go beyond extracting these subjective causal effects. ASML, the Dutch lithography machines manufacturer, said in its January 2024 earnings call that “we do not expect to get export licenses for our most advanced immersion systems, the NXT:2000 and up for China in 2024,” and argued the effect is at “10% to 15% of our 2023 China system revenue”. We first ask the LLM to determine if the firm says it is positively or negatively affected and write a short reason for why the firm says it is positively or negatively affected, respectively. It writes the reason “the firm expects to be denied export licenses for specific high-end tools for Chinese customers and quantifies the lost revenue” for being negatively affected. We then ask the LLM to classify that specific reason against a taxonomy of about 100 mechanisms for being negatively affected, which we build from the model’s summary over the full set of such reasons. It assigns the mechanism “Export License Denial Loss.” Finally, we group the mechanisms into six economic categories, and this one falls under “Direct Market Access Loss.”¹²

Figure 5 reports the reasoning outcomes for firms exposed to export controls, separating firms that report being positively or negatively affected and those located in the imposing or target countries.¹³ In the imposing countries the reasons behind the losers are mostly driven by “Direct Market Access Loss” followed by “Operational & Compliance Costs”. These firms are losing customers in the banned market-product combinations. The export control rules are also sufficiently opaque, and the penalties for violations sufficiently high, that the cost of complying with the regulations features prominently.

The bottom panel of Figure 5 focuses on firm-quarter combination in the target countries. The winners here mention explicitly facing new demand and opportunities because their foreign competitors have been shut-off from the market. The losers mention the loss of high quality inputs and bottlenecks in supply chains. This figure highlights both the intended policy outcome from the country imposing the controls, deteriorating the production function of the targeted country in that specific sector, and the unwanted target response, the ability to substitute from other foreign sources as well as the increased effort to produce domestically.


The Reasoning Behind the Chinese Firms R&D Response to U.S. Export Controls. The Chinese firms response to U.S. export controls in the semiconductor industry has been the subject of intense policy debates. We zero on the sample of Chinese firms that report raising R&D in response to export controls, and apply the reasoning pipeline to this specific action.

Figure 6 shows that “Core Technology Indigenization” dominates at 66 percent of the flagged transcripts, with the next two motivations focusing on risk-mitigation and product redesign. The

¹²We let the LLMs decide the final number of categories to aggregate into from all the mechanisms.

¹³We filter the firms in two ways to clean up the sample. First, we exclude firm-quarters that feature both negative and positive outcomes (both winners and losers). Second, we restrict to firm-quarters for which the positively or negatively affected classification in the SCE analysis agrees with the winner-loser classification in the Reasoning analysis. These filters also apply to the reasoning analysis of tariffs and sanctions.


Figure 5: Why firms report winning or losing from export controls



Notes: This figure plots the mechanisms through which firms report winning or losing from export controls, grouped into broad categories. The top panel shows firms headquartered in a country imposing the controls and the bottom panel firms headquartered in a targeted country. The unit is the firm-quarter; firm-quarters reporting both gains and losses are dropped. The sample is restricted to transcripts in which the reason extraction and the main classification agree on the sign of the impact. Areas are proportional to shares of firm-quarters within each panel.

Chinese firms are specifically reporting doing R&D in response to U.S. export controls targeting China, and the reason behind the R&D being the creation of a domestic alternative.¹⁴

Figure 6: **Why Chinese firms increase R&D in response to export controls**



Notes: This figure plots the stated motivation for raising R&D among Chinese firms that report increasing R&D in response to export controls targeting China. The unit is the transcript from the Orbit Chinese firms’ earnings calls. The motivation is extracted from the transcript’s free text explanation and grouped into categories discovered from the explanations themselves. Transcripts whose explanation does not state a motivation are shown as “Motivation Unclear”. The sample keeps only transcripts in which the reason extraction and the main classification agree on the sign of the impact.

5 Sanctions

We next turn to studying how firms around the world respond to sanctions. Our analysis focuses on how firms outside of Russia were affected. While we have some coverage of Russian firms, they systematically stop appearing in our data after the invasion and ensuing sanctions.¹⁵ We group firms into those based in countries that are part of the sanctioning coalition, and those located in other countries.¹⁶


¹⁴Digging into the underlying data, we observe some firms that report that they expect their innovation success to offset the effects of the controls, while others emphasize the cost of pursuing the innovation.

¹⁵Appendix Figure A.1 provides an extension using country-sector reports on Russia to cover what happened in Russia as a result of the sanctions. However, these reports are written by Western analysts with limited access to ground truth of economic developments at the firm level in Russia.

¹⁶In the case of sanctions on Russia, we assign a firm to the sanctioning coalition when its own country is among the countries imposing the sanctions. Where a transcript names Russia as the target but leaves the imposer blank, we assign the firm to the coalition if it is domiciled in the EU-27 or in Australia, Canada, Iceland, Japan, Liechtenstein, Monaco, New Zealand, Norway, Singapore, South Korea, Switzerland, Taiwan, Ukraine, the United Kingdom or the United States, and to the third-country group otherwise. Russian firms are excluded throughout.

Figure 7 reports the aggregate response of firms to the sanctions. For firms based in the countries imposing the sanctions, the response is overwhelmingly negative. The firms lose sales, have lower profit margins, and lower investment. Firms located in third-party countries have a more balanced reaction. On the negative side, their reaction is similar to the firms based in the sanctioning countries, but there is a mirror positive reaction. These are firms that gain new business opportunities from the business relationships that the sanctions have broken, or firms that benefit from the price movements (e.g. oil prices) induced by the sanctions.

Figure 7: **2022 Russian Sanctions: firm responses**



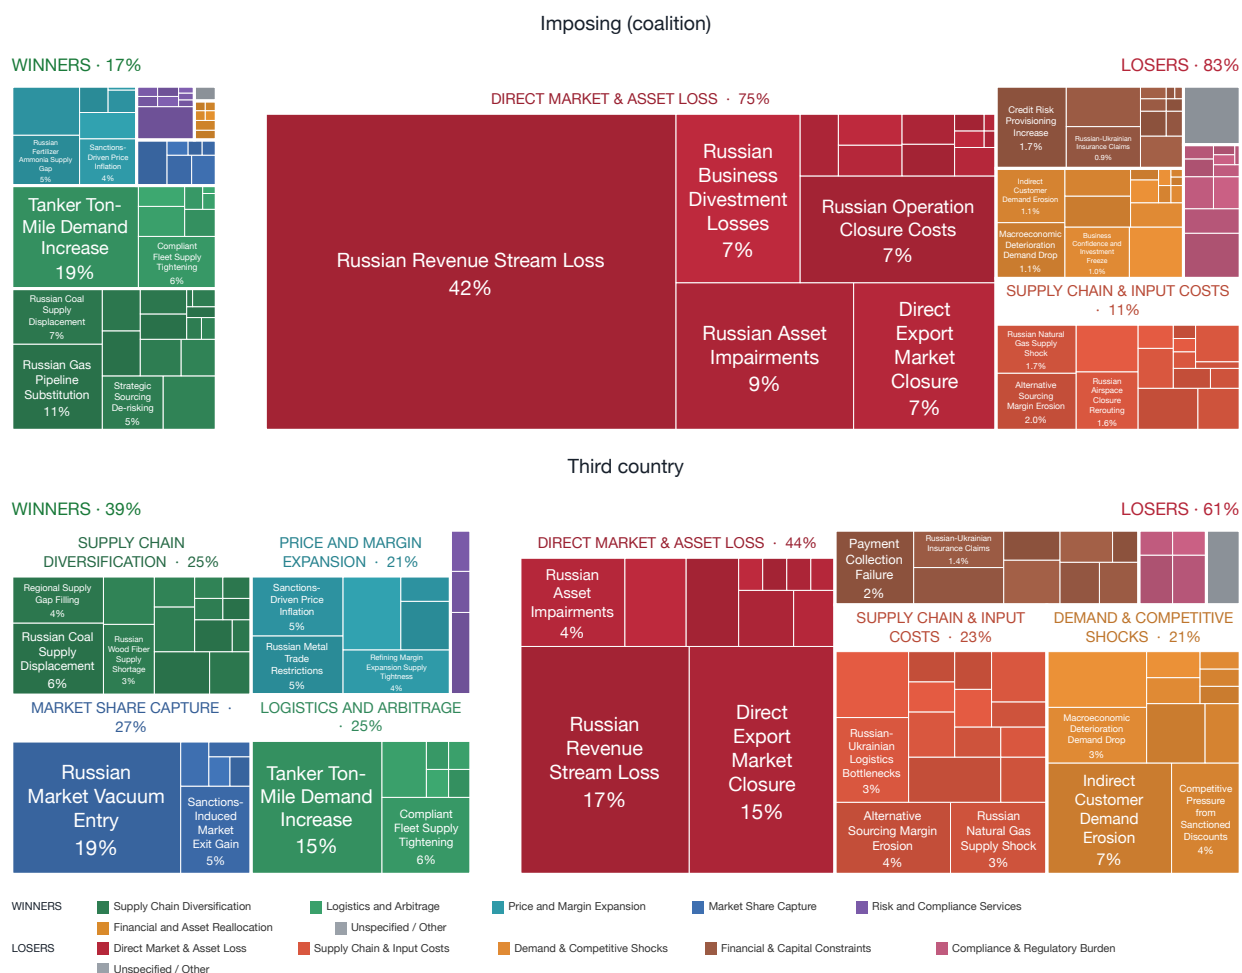
Notes: This figure plots the share of firm-quarters affected by sanctions on Russia that report an increase (+1) or a decrease (−1) in each outcome. The sample is the firm-quarters from 2022 onward where a firm discusses the effect of sanctions naming Russia as the target. Panel (A) shows firms headquartered in the 2022 sanctioning coalition and panel (B) firms headquartered elsewhere, excluding Russia. When the transcript does not name the sanctioning country, the firm is assigned by its headquarters. Error bars denote 95% confidence intervals.

5.1 Reasoning Analysis

In order to understand the effects of the sanctions on firms around the world, we begin with the set of firms-transcripts that flagged as firms being affected by sanctions with Russia as the receiving country. We then apply the reasoning pipeline to this sample.

Figure 8 reports the decomposition of winners and losers around the world focusing on two groups: firms located in countries that were part of the sanctioning coalition, and firms located in third-party countries. Starting with the firms located in the coalition countries, we find that among losing firm-quarters 75% report losing because of a "Direct Market and Asset Loss", of which the largest categories are "Russian Revenue Stream Loss" and "Russian Asset Impairments". This category picks up both the direct effect of sanctions, being unable to sell goods and services to Russia, and the retaliation with Russia seizing local assets of firms from the sanctioning coalition.

Figure 8: Why firms report winning or losing from the 2022 Russia sanctions



Notes: This figure plots the mechanisms through which firms report winning or losing from the 2022 sanctions on Russia, grouped into broad categories. The top panel shows firms headquartered in a country of the sanctioning coalition and the bottom panel firms headquartered in third countries, excluding Russia. The sample covers 2022 onward firm-quarters, drops observations reporting both gains and losses, and keeps only transcripts in which the reason extraction and the main classification agree on the sign of the impact. Areas are proportional to shares of firm-quarters within each panel.

United States

WINNERS · 18%

- Russian Fertilizer Supply Gap: 5%
- Sanctions-Driven Price Inflation: 5%
- Tanker Ton-Mile Demand Increase: 4%

SUPPLY CHAIN DIVERSIFICATION · 58%

- Russian Coal Supply Displacement: 12%
- Russian Gas Pipeline Substitution: 16%
- Russian Uranium Supply Replacement: 5%
- Strategic Sourcing De-risking: 8%

LOSERS · 82%

DIRECT MARKET & ASSET LOSS · 83%

- Russian Revenue Stream Loss: 57%
- Russian Business Divestment Losses: 5%
- Direct Export Market Closure: 7%
- Russian Operation Closure Costs: 6%
- Russian-Ukrainian Insurance Claims: 1.2%
- Russian-Ukrainian Travel Refunds: 1.0%
- Macroeconomic Penetration Demand Drop: 1.1%
- Business Confiscation and Investment Pressure: 1.1%
- Supply Chain Disruption for Components: 1.1%
- Alternative Sourcing Margin Erosion: 1.2%
- Forced Service Termination Bad Debt: 1.0%

Europe

WINNERS · 13%

- Logistics and Arbitrage: 52%
- Tanker Ton-Mile Demand Increase: 34%
- Compliant Fleet Supply Tightening: 10%

LOSERS · 87%

DIRECT MARKET & ASSET LOSS · 71%

- Russian Revenue Stream Loss: 35%
- Russian Operation Closure Costs: 7%
- Direct Export Market Closure: 6%
- Russian Asset Impairments: 10%
- Russian Business Divestment Losses: 8%

DEMAND & COMPETITIVE SHOCKS · 6%

- Tourism and Travel Volume Decline: 1.2%

SUPPLY CHAIN & INPUT COSTS · 14%


- Russian Natural Gas Supply Shock: 2%
- Russian Aerospace Closure Rerouting: 3%
- Alternative Sourcing Margin Erosion: 2.5%
- Credit Risk Provisioning Increase: 2%
- Payment Collection Failure: 0.8%
- Russian-Ukrainian Insurance Claims: 0.8%

Legend:

- WINNERS**
 - Supply Chain Diversification
 - Logistics and Arbitrage
 - Price and Margin Expansion
 - Market Share Capture
 - Risk and Compliance Services
 - Financial and Asset Reallocation
 - Unspecified / Other
- LOSERS**
 - Direct Market & Asset Loss
 - Supply Chain & Input Costs
 - Demand & Competitive Shocks
 - Financial & Capital Constraints
 - Compliance & Regulatory Burden

21

Figure 10: Sanctions mechanisms outside the coalition: China and India



Notes: The figure plots the mechanisms through which firms report winning or losing from the 2022 sanctions on Russia, grouped into broad categories. The top panel shows firms headquartered in China and the bottom panel firms headquartered in India. The sample covers 2022 onward firm-quarters, drops observations reporting both gains and losses, and keeps only transcripts in which the reason extraction and the main classification agree on the sign of the impact. Areas are proportional to shares of firm-quarters within each panel.

This is followed by “Supply Chain & Input Costs”, in which firms report facing costs to diversify their sourcing, facing higher input costs as a result of disruption, or directly paying higher logistics and routing costs. Winners in the coalition are rare and they are mostly concentrated in supplying energy that was previously coming from Russia, as well as increased demand for shipping-tanker services.

Switching to the firms located in third-party countries, the bottom panel of Figure 8 provides a meaningfully different picture. The balance between Winner and Loser is much more even than in the Coalition. In particular, the biggest reason for winning is the new market opportunities created by the loss of Russian competitors, especially in the energy market. Before the sanctions, a Russian firm was the dominant competitor, now that firm can no longer export, and a firm from a third-party country picks up the business. Other major reasons for winning include logistical services and bypassing the sanctions.

The Geography of Winners and Losers from the Sanctions on Russia. While the aggregate patterns above give a snapshot of the effects of these sanctions around the world, more information can be gleaned by studying the prevalence of these categories across different countries. We compare the effect on American, European, Chinese, and Indian firms. The US and Europe led the sanctioning effort on Russia, so the firms located in those countries are part of the coalition imposing the sanctions. On the contrary, India and China not only did not participate in the sanctions but actively increased business ties with Russia. The split, therefore, offers a reasoning-view across geographies that played a very different role in applying geoeconomic pressure to Russia.

Figure 9 contains the results for firms in the U.S. and Europe. The two panels are broadly similar, with losers dominating the picture and providing “Direct Market & Asset Loss” as the underlying reason. What is more interesting are the differences between the U.S. and European firms. European losers mention more often the losses arising from “Supply Chain & Input Costs” compared to U.S. firms. U.S. firms that are positively affected are more likely to mention supplying energy that was previously sources in Russia. This is consistent with Europe reallocating its energy sourcing away from Russia and in part toward the U.S. as a result of the sanctions. European firms suffer more, and some U.S. firms benefit from the increased European demand for energy.

Both the US and European losers contrast sharply with Chinese and Indian firms that report losing as shown in Figure 10. Chinese and Indian firms are much more likely to report losing because of “Demand and Competitive Shocks” or “Supply Chain & Input Costs”, essentially categories that capture disruptions to world economic activity arising from the sanctions. They also report “Direct Market & Asset loss”, but much less prominently than European and U.S. firms do.

The heterogeneity on the winner side is even starker. The single largest reason for Chinese and Indian firms to report winning is that these firms fill the vacuum of those broken business ties between the Coalition and Russia. These new opportunities arising in a two sided vacuum: Chinese and Indian firms are both supplying-buying from the Russian entities, and supplying-buying from firms in the Coalition. Indian firms specifically mention “Discounted Russian Crude Procurement”

as a reason for being positively affected. This is consistent with India buying Russian oil at a discount compared to world market prices.

Figure 11 provides an expanded analysis for countries around the world. The graph is ordered with at the top the countries the firms of which mostly report being positively affected. At the top of the graph: Monaco, Greece, and Bermuda have firms that overwhelmingly benefit from the logistical, insurance, and supply chain adjustments to the sanctions. Other notable winners are firms in commodity and oil exporting countries: Australia, Canada, Norway, South Africa. These firms are benefiting from higher commodity and energy prices resulting from the sanctions and Russia’s own export restrictions in retaliation. The losers are geographically concentrated in the countries that are both part of the sanctioning coalition and geographically close to Russia. The gravity structure of trade means that firms located in these countries tended to have closer business ties with Russia before the sanctions and indeed report losing access to that market as the main reason for the losses.

6 Tariffs


The final instrument that we explore in detail is tariffs. Figure 12 plots the firm response to tariffs. In the imposing country, we find that firms that report being affected by tariffs are likely to report facing higher input prices and are likely to report raising the prices they charge for their products. In addition, firms that report being affected by tariffs also report a reduction in their profit margin. This pattern is consistent with the consumer-tax mechanism of tariffs and the tariffs hitting intermediate inputs: firms in the importing country face higher input prices, including the tariff, and pass through some of that into higher sales prices, and absorb part of the tariff in their profit margin. We explore this mechanism further below. Firms located in countries the export of which are being subjected to tariffs, e.g. a European firm in the case of a U.S. tariffs on imports from Europe, are likely to report lower sales, lower profit margins, and being negatively affected. In the aggregate, we do not find the firms responding that they lower their prices in response to the tariffs, a mechanism that would have been direct evidence of terms of trade manipulation at the firm level. The firms are negatively affected by losing sales, consistent with the lower input acquired by firms in the imposing country in Panel A.

6.1 Winners and Losers from the Tariffs

Having documented how firms adjust their behavior in response to tariffs in the aggregate, we now zoom in to ask why firms report winning or losing from tariffs by applying our reasoning pipeline.


Figure 13 decomposes these firm-quarters by mechanism. In the countries imposing the tariffs the reasoning behind the losses comes from three main components. The largest is the direct effect of the importer paying the tax, inputs become more expensive. This effect is consistent with exporters in other countries not cutting their price sufficiently in response to the tariff, so that the importer

Figure 11: Country mechanisms in the 2022 Russia sanctions



Notes: This figure plots, by headquarters country, the share of firm-quarters reporting a gain (right) or a loss (left) from the 2022 sanctions on Russia, with each bar split by the mechanism category the firm cites. Countries with at least 10 firm-quarters, excluding Russia, are shown and sorted by the share of winners. The sample covers 2022 onward, drops all firm-quarters reporting both gains and losses, and it keeps only transcripts in which the reason extraction and the main classification agree on the sign of the impact. A firm-quarter with several transcripts takes its most frequent mechanism.

Figure 12: **Tariffs: firm responses**



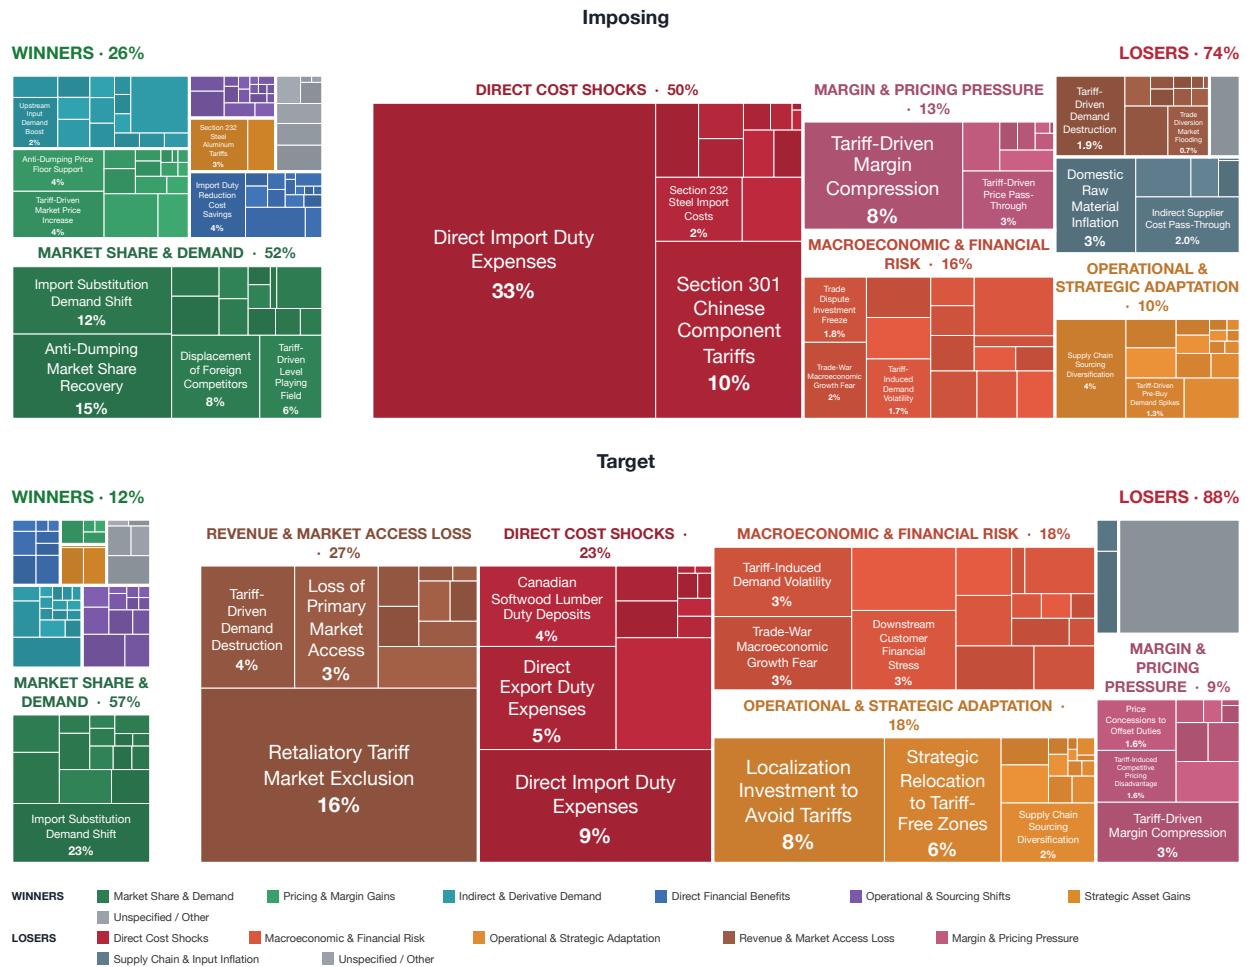
Notes: This figure plots the share of firm-quarters affected by tariffs that report an increase (+1) or a decrease (−1) in each outcome. The sample is the firm-quarters where a firm discusses the effect of tariffs on itself. Panel (A) shows firms headquartered in a country imposing the tariffs and panel (B) firms headquartered in a country receiving them. Error bars denote 95% confidence intervals.

faces a higher price inclusive of the tariff. A related reason for the losses is “Margin & Pricing Pressure”, unless the importer can pass the higher input cost into its sales price, it faces a decrease in the profit margin. A broader category of “Macroeconomic & Financial Risk” speaks to the general equilibrium negative effects of tariffs and the uncertainty related to the trade war. For example, a bank CEO might say that the effect of the tariffs is negative on the bank, despite none of its financial services being directly exposed to the tariffs, because the bank’s industrial customers or the consumers are negatively affected.

The winners in the imposing countries are benefiting because they are gaining demand and market share. These are producers who were facing foreign competition in their domestic market. The tariffs put the foreign competitors at a disadvantage and provide a subsidy to the domestic producers. These producers specifically mention “recovering” their market share thanks to Anti-Dumping tariffs.

The bottom panel of Figure 13 focuses on firms in the countries that are targeted by the tariffs. The losses are coming into two main forms. First, the tariffs imposed by the foreign country are causing a loss in demand and market access, second their are imposing costs as the firms reshape their supply chains to avoid the tariffs. For example, firms might localize production or relocate it to third countries that face lower tariffs. A second distinct channel is that firms mention losing when their own countries impose retaliatory tariffs in response, in which case this losses are of the “imposing country” type. Finally, just like in the top panel, firms report losing because of the uncertainty and general equilibrium effects of the tariff war. We do find evidence of “Price Concessions to Offset Duties”, consistent with terms of trade manipulation, but quantitatively it is

Figure 13: Why firms report winning or losing from tariffs




Notes: This figure plots the mechanisms through which firms report winning or losing from tariffs, grouped into broad categories. The top panel shows firms headquartered in a country imposing the tariffs and the bottom panel firms headquartered in a targeted country. The unit is the firm-quarter, and we drop observations reporting both gains and losses. The sample is restricted to transcripts in which the reason extraction and the main classification agree on the sign of the impact. Areas are proportional to shares of firm-quarters within each panel.

the primary reason for only 1.6% of the losers. We also find evidence that these firms pay the tariffs when the importer of record is their own subsidiary in the country imposing the tariff, or when they have a subsidiary in that country that is importing from other foreign firms.

Since the effects induced by uncertainty or threats about future tariffs feature prominently in the reasoning analysis, we further investigate the split between current and future tariffs in affecting the firms. Figure 14 decomposes the fraction of firms that report being affected by current versus future tariffs in each quarter. The vertical dotted lines are the fourth quarter of 2016 and 2024: in both instances, President Trump won a presidential election running on a campaign that made tariffs a major feature. 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. In 2025, we observe firms being affected by both current and future tariffs, consistent with the Trump administration both immediately imposing tariffs and announcing further tariffs to be imposed at uncertain future dates. By 2026 the effects are mostly arising from current tariffs, despite continued uncertainty about U.S. tariff policy.

Figure 14: **The effect of current and future tariffs**



Notes: This figure plots the composition of firm affected by tariffs into those reporting an effect of current tariffs only, of future tariffs only, and of both by quarter. The sample is the firm-quarters in which a firm discusses the effect of tariffs on its business. Vertical dashed lines mark the fourth quarters of the 2016 and 2024.

6.2 The 2025 Trade War

In this subsection, we apply our methodology to zoom in on the trade war that started in the first quarter of 2025. We analyze text from the first quarter of 2025 until the first quarter of 2026, capturing the April 2nd announcement of tariffs and the ensuing trade war.¹⁷

6.2.1 Changing Global Supply Chains

One of the stated goals of the U.S. government tariffs was to reshape global supply chains toward the United States. In this subsection we explore the extent to which we observe this reallocation in the data. In Figure 15, we separately consider the response of U.S. (top panel) and Chinese (bottom panel) firms to U.S. tariffs from the first quarter of 2025 until the first quarter of 2026. We condition on the set of firms that report changing their supply chains and naming both the country they are moving away from and the country they are moving towards.¹⁸ In the case of the United States, we find that the primary location firms leave is China, with 69% moving away from China. Of those leaving China, the primary market they are moving their supply chain towards is the United States (at 40%), followed by Vietnam and Mexico. This is largely consistent with a China +1 strategy, tariff tunneling via a third country, or friendshoring. While we observe some exits from other markets by U.S. firms, the movement out of China is the dominant pattern.

We find a more nuanced pattern for Chinese firms. While a similar majority of 76% report leaving China, the rest report leaving the United States (23%). Of those leaving China, we find far less movement towards the United States than we do for U.S. firms. Of those Chinese firms reporting moving their supply chains away from China, a substantial fraction report moving to Vietnam, followed by Thailand. In stark contrast, of the 23% of identified Chinese firms leaving the United States, more than a third report moving their supply chains back towards China, a response almost completely absent for the U.S. firms.

6.2.2 2025 Tariff War: Winners and Losers


Having documented the supply chain reallocation of the U.S.-led Trade War, we now turn to understanding the self-reported winners and losers to tariffs around the world. Figure 16 displays the share of both American and non-American firms that report being negatively or positively affected by tariffs. 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. We find that a far higher share, by count, of American firms report being negatively affected by tariffs than do foreign firms.

¹⁷For 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. \(2025\)](#).


¹⁸Among tariff-affected firm-quarters that name the U.S. as the imposing country, 61% of U.S. firms and 49% of Chinese firms report a supply chain adjustment. Of those, we consider only the firms we are able to identify the country they move away from and towards. Flow widths count origin-destination pairs mentioned, so a firm naming several countries contributes several pairs.

Figure 15: **Supply chain reshuffling**

(a) U.S. firms




(b) Chinese firms




Notes: This figure plots the supply chain relocations that firms report in response to U.S. tariffs, from the country they move away from (left) to the country they move towards (right). The sample is the firm-quarters from 2025Q1 to 2026Q1 where a firm discusses the effect of U.S. tariffs and names both an origin and a destination country. Panel (a) shows U.S. firms and panel (b) Chinese firms. Flow widths are origin-destination pairs.

Figure 16: **The self-reported effect of tariffs on American and 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.

Figure 17: Responses of American firms in 2025 to U.S. tariffs, by self-reported impact




Notes: This figure shows how American firms describe adjusting their business in 2025 in response to U.S. imposed tariffs, grouping the responses by whether the firm reports being overall positively or negatively affected. The sample is firm-quarters in which a U.S.-headquartered firm discusses tariffs and names the United States among the imposing countries. Firm-quarters reporting both a positive and a negative impact, or neither, are excluded. For each outcome, the bar is the share of firms reporting an increase minus the share reporting a decrease. Error bars denote 95% confidence intervals.

Figure 17 explores how the responses of American firms differ conditional on whether they report being positively or negatively affected. Among firms that report being negatively affected, at least 75% report facing higher input prices and close to half report raising the prices they charge. The positively affected are close to the mirror image: their net input-price response is near zero (+0.08, against +0.75), and it is smaller than the net response on their own sales prices (+0.12), so what price movement they report falls on what they sell rather than on what they buy. The sharpest contrasts are in quantities and margins. At least 17% of negatively affected firms report falling sales quantities, against more than a third of positively affected firms reporting rising quantities; and net profit margins are -0.63 for the negatively affected against $+0.27$ for the positively affected.

In Figure 18, we explore the price reaction of U.S. firms in 2025 to tariffs more systematically. In particular, we zoom in on only American firms that report being affected by tariffs from 2025 to the present and ask how likely it is that an affected firm is classified as being positively affected by tariffs. In Panel (a) we examine firms that report being positively but not negatively affected, and in Panel (b) we consider any firm that is positively affected, even if they also report some negative effects. We sort firms into 4 possible buckets: (1) firms that do not report facing higher input prices or raising their sales price because of tariffs ("Neither"), (2) firms that report raising their sales price because of tariffs but did not see their input prices rise because of tariffs ("Sales price up only"), (3) firms that saw their input prices increase because of tariffs but did not raise their sales price ("Input price up only"), and (4) firm that report both that their input prices increased because of tariffs and raised their sales prices because of tariffs ("Both input and sales price up"). Each bar then reports that firms within each of those 4 buckets are classified as being positively affected. In Panel (a), we see that 11% of firms in the Neither bucket are classified as being positively affected. Firms that only face higher input prices or those that face higher input prices and raise their sales prices have only a 1% chance of only being positively affected. This is unsurprising as the higher input prices themselves is likely the reason the firm is classified as being negatively affected by tariffs. The striking result is that 22% of firms who raise their sales prices because of tariffs but do not see their input prices rise are classified as being only positively affected by tariffs. In panel (b), we see that 41% of firms in the sales-price-up-only bucket are classified as having any positive effect. This is consistent with a channel where the winners from tariffs are those firms who do not rely heavily on inputs who prices increase rise because of tariffs, but who are able to raise their prices as their import-reliant competitors raise their prices in response to rising costs.

What is behind these results? 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 American firms or foreign firms that sell in the US, 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 raise their prices. It is the domestic producer subsidy component of a tariff. As an example, the CEO and CFO of Cleveland-Cliffs Inc, an Ohio-based steel manufacturer,

Figure 18: Price Responses to Tariffs Among U.S. Firms



Notes: This figure plots the share of U.S. firm-quarters affected by tariffs that report a positive impact, by whether the firm attributes higher input prices, higher sales prices, both, or neither to tariffs. The sample is firm-quarters from 2025Q1 onward in which a U.S.-headquartered firm discusses the effect of tariffs on itself, whichever country imposed them. Panel (b) uses all such firm-quarters; panel (a) drops those reporting both a positive and a negative impact. Firm-quarters reporting neither sign remain in the denominator of both panels. Error bars are 95% confidence intervals from firm-clustered standard errors.

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 imports from all countries. These tariffs are critical to addressing the problem... 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.”

As it turns out, in March 2025 Cleveland-Cliffs announced layoffs of 1,200 workers in their Michigan and Minnesota operations, but the management maintained an upbeat assessment for the long-run impact of the tariffs. The fact that the manager’s forecast of positive effects did not necessarily materialize is captured in our idea of the Subjective Causal Effect as the object of measurement.

6.3 Reasoning Behind the Effects of Tariffs

We next apply our reasoning pipeline to the tariffs. Figure 19 decomposes winner and loser firm-quarters in the 2025 trade war for tariffs imposed by the United States. The panels focus on the reasoning for the reaction to U.S. imposed tariffs by firms in the U.S., China, Europe, and Canada.

The U.S., being the imposing country, has a distinctive pattern with firms mostly losing because of the “Direct Cost Shocks.” On the receiving end of the tariffs there is heterogeneity between

Chinese, European and Canadian firms. China was in practice the target of the highest (potentially) tariffs. Consequently, the losers in China focus on the cost of “Operational & Strategic Adaptation” to minimize the direct exposure to U.S. tariffs, and to the extent that adaptation was not successful the second reason for losses is “Revenue & Market Access Loss”. European and Canadian firms, instead, focus most on “Macroeconomic & Financial Risk”. For Europe and Canada, the U.S. tariffs were a prominent example of geoeconomic pressure that led to a major reevaluation of the overall economic relationship with the US. This major rupture in economic relations can have negative consequences for the firms well beyond the direct effect of tariffs.

Firms in these foreign countries also report paying the tariffs in different ways. The most obvious is that they mention “Price Concessions to Offset Duties”. This is the cleanest translation in the data of the textbook terms of trade manipulation: the exporter cuts its price and absorbs part of the tariffs. However, in the data we find more often related mechanisms. For example, the firms might mention paying for the tariffs because the importer of record in the U.S. is their own subsidiary.

We broaden our analysis beyond these countries in Figure 20. For each country, a pair of stacked bars gives the broad-mechanism composition of its winners and of its losers. The graph is sorted with at the top the countries with the highest net share of winners. Losers around the world are likely to report Macroeconomic and Financial Risk or a loss of Revenue and Market Access as a primary reason for the losses, but direct cost shocks also feature prominently. Countries with many winners, like Greece, benefited by providing services that allowed other countries to reshape their supply chains in response to the tariffs.


7 Validation

Given that our paper and analysis relies on LLM-based classification of text, the natural question arises of whether our results are actually capturing the underlying reality. We find it useful to conceptualize the issues around two broad categories. One category is about the relationship between the text and the actual real outcomes at the firm level. Issues in this category include, for example: strategic and cheap talk from the CEOs and CFOs, skew in the coverage of earnings calls, and differential propensity to talk about specific issues in general or over time. Intuitively, we think of the first category as related to what the text covers and how to interpret it compared to the underlying reality. The second category is about how the classifier extracts information from the text. Issues in this category include, for example: replicability, and noise and bias in measurement error. Some of these issues apply to traditional NLP, and surveys more generally. Some issues are more specific or exacerbated by LLMs.

7.1 Observables


In this section, for each of the three forms of geoeconomic pressure, we present evidence that our measures line up with observable firm outcomes.

Figure 19: Reasoning Analysis for Reaction to U.S. Imposed Tariffs in 2025 Trade War



Notes: This figure plots the mechanisms through which firms report winning or losing from the 2025 U.S. tariffs, grouped into broad categories. The four panels show firms headquartered in the United States, China, Europe, and Canada. The sample covers 2025 transcripts that name the United States as the imposing country, uses the firm-quarter as unit, drops firm-quarters reporting both gains and losses, and keeps only transcripts in which the reason extraction and the main classification agree on the sign of the impact.

Figure 20: Country mechanisms in the 2025 trade war




Notes: This figure plots, by headquarters country, the share of firm-quarters reporting a loss (left of zero) or a gain (right) from the 2025 trade war, with each bar split by the mechanism category the firm cites; the two shares sum to 100 for each country. The sample covers 2025 transcripts that name the U.S. as an imposing or targeted country, drops firm-quarters reporting both gains and losses, and keeps only transcripts in which reason extraction and main classification agree on the sign of the impact. A firm-quarter with several transcripts takes its most frequent mechanism. Countries with at least 10 firm-quarters are shown, sorted by the share of winners.

7.1.1 Export Controls Validation: BIS Entity List

To validate our export control classifications, we construct a panel of Chinese firms placed on BIS export control lists (Entity List, Unverified List, Military End-User List) using the eCFR API and Federal Register API. We follow the methodology of [Crosignani et al. \(2024\)](#) to identify Chinese firms on the Entity list and find 1,550 unique Chinese entities across lists¹⁹. We divide the firms in the data into three groups: Chinese firms on the U.S. export control lists, all other Chinese firms, and firms in the Rest of the World (excluding the U.S.). Across these three groups, we measure the share of firms that report being affected by export control during the period from 2008 to 2026 in our LLM classifications. We find that Chinese firms on the Entity List have a 53% chance of ever reporting being affected by export controls, compared to 13.5% of other Chinese firms, and only 7% of firms from the rest of the world (excluding the U.S.). The fact that Chinese firms on the entity list are four times more likely than other Chinese firms and seven times more likely than firms from the rest of the world to report being affected by export controls provides strong evidence that our flag for being affected by export controls is associated with the underlying reality.

Figure 21: Chinese firm Entity List inclusion and LLM classification




Notes: The figure compares how often our export control flags are positive across three groups of firms. The unit of observation is the firm and the sample is every firm with at least one earnings call transcript in the corpus, from 2008 to 2026. A firm counts as flagged if the second-stage export-control prompt ever records an effect on the firm in a quarter. Firms are split by headquarters country and by whether they appear on the US export control lists: Chinese firms on the lists, Chinese firms not on them, and rest-of-world (RoW) firms not on them (excluding U.S. firms). Each bar is the share of firms in its group that are ever flagged for export controls.

¹⁹Among these, we can match 98 to a company in our corpus, which is the group the first bar of Figure 21 describes. Most listed entities are unlisted and have no earnings calls.

7.1.2 Validation: Sanctions

We next turn to validating the sanctions classifications against externally observable behavior. We focus on the probability that a firm, excluding Russia firms, reports being negatively affected by sanctions on Russia depending on how reliant it was on selling products to Russia before the sanctions regime was put in place. In order to do this, we match each of our firms to the FactSet Geographic Revenue (GeoRev) dataset. This dataset estimates the share of each firm's revenue earned by sales to different countries and geographic areas. In Figure 22, we report a binned scatter plot of the relationship between revenue from Russia prior to the invasion of Ukraine and the subsequent probability of being negatively affected. We exclude the Russian firms from the sample. We also report separately Indian and Chinese firms as a placebo, since those countries did not participate in the sanctions. We document a strong positive relationship among all firms, and a much flatter relationship for Indian and Chinese firms. This provides evidence that the sanctions flag are picking up economically meaningful channels.

Figure 22: **Sales to Russia and Probability of Being Negatively Affected by Sanctions in 2022**




Notes: The figure plots the share of firm-quarters reporting a negative sanctions impact against the firm's share of sales to Russia, separately for firms headquartered in China and India and for firms headquartered elsewhere. The sample covers firm-quarters in 2022 for non-Russian firms with a FactSet GeoRev revenue share for Russia, measured from the latest report before the invasion. Firm-quarters that mention sanctions not targeting Russia are dropped. Within each group, firms with no sales to Russia form one bin and the rest are split into 19 equal-sized bins; each point shows the bin average.

7.1.3 Validation: Tariffs

We finally turn to validating the measurement of tariffs. Figure 23 provides a similar analysis as Figure 22 does for the exposure to sanctions on Russia. It bin-scatters the probability that a firm reports being negatively affected by US tariffs during 2025 against the firm's sales share to the United States in the firm's most recent FactSet GeoRev report. We split firms into whether they sell goods or services, as tariffs are on goods but not services, allowing services to serve as a placebo group, where we capture exposure to the US during the trade war, but presumably without the direct effect of the tariffs themselves.²⁰ We find that the probability of being negatively affected by tariffs increases far more sharply with the US sales share for goods providers rather than services providers. At a broad level, this analysis confirms that LLM classifications are related to the underlying real outcomes at the firm level.

Figure 23: Sales to USA and probability of being negatively affected by tariffs in 2025



Notes: This figure plots the share of firm-quarters reporting a negative tariff impact against the firm's share of sales to the United States, separately for goods-producing and service firms, classified by two-digit SIC. The sample covers firm-quarters in 2025 for non-US firms with a FactSet GeoRev revenue share for the United States, taken from the firm's latest report. Firm-quarters that mention tariffs not imposed by the United States are dropped. Firms with no US sales form one bin and the rest are split into equal-sized bins.

²⁰Service firms can of course still report being affected by the tariffs because their clients are and because of the macroeconomic effect of tariffs.

7.2 Reproducibility and Inter-Model Agreement

Reproducibility in the narrow sense asks whether the same text, submitted to the same prompt, generates the same classification across multiple runs. We found that in practice a stable compute architecture and greedy sampling of open-weight models performed well in this respect. The same, of course, cannot be guaranteed with closed-weight models that might change over time without the researcher’s knowledge or control.

We focus instead on the sensitivity of the classification to the choice of model, with Gemma 4 as our benchmark. Figure 24 provides confusion matrices as well as summary statistics for four different model runs. The models are: Gemma 4 31B, Qwen 3 235B-A22B, GLM 4.5 Air, and Llama 3.3 70B. The classifications being inspected are the flags for the firm-transcript being affected by tariffs in Panel (a), sanctions in Panel (b), and export controls in Panel (c). These flags are produced by the broad prompt in Appendix B.1 applied to all earnings calls in our sample. Starting from Panel (a), the block matrices in the lower left half are confusion matrices. Consider for example the confusion matrix comparing Gemma (our benchmark) and Qwen. For 579,744 transcripts the two model both flag a 0, and for 35,794 transcripts both models flag a 1. These are the model agreements. For 7,087 transcripts Gemma flags 1, while Qwen flags 0; and for 2144 the reverse happens. These are the model disagreements. The upper left of the same panel reports common statistical metrics to assess the degree to which the two classifiers agree. We focus on the Gwet’s AC1 statistic that measures model agreement, while penalizing for random chance of agreement. An AC1 of 0.983 means the two models tend to have high agreement. The next reported statistic is Cohen’s κ , which tends to be lower since our underlying data is naturally skewed (in most quarters firms are not affected by tariffs in the ground truth). Finally, the pos statistic is the Dice’s coefficient on positive agreement.

Inspecting the three panels together highlights several patterns. First, the level of confusion in these classifiers is not trivial, especially for export controls. Disagreements can be substantial. At the same level, all models are clearly picking up a common component of the underlying ground truth. Llama appear to be the most conservative model in terms of flagging a positive exposure to the instruments.

Figure 25 present the Gwet’s AC1 statistics for classifications done by the long prompts that are instrument specific. These are the prompts in Appendices B.2 – B.4. For any two models, we restrict to the sample in which the both agree on the first question (the firm is affected), and then compute the AC1 statistic for each of the subsequent questions. The picture that emerges is one of common underlying ground truth, but with substantial levels of errors in the classifiers.

The possibility for substantial noise and bias in LLM classification is a noted property and a challenge as well as an opportunity offered for researchers using these new methods. In ongoing work, we are adding comparisons for subsamples to frontier closed weight models, and potentially to human readers. Indeed, there is an active and fast moving literature on debiasing LLM classifications and how to adjust point estimates and standard errors in econometric analysis that uses LLM generated data. The literature in statistics on Prediction-Powered-Inference (PPI) (Angelopoulos et al. (2023, 2024)) is being imported and adapted in economics (Ludwig et al. (2025); Carlson and Dell (2025)).

In ongoing work, we are experimenting with these methods. A broader aim of paper is to eventually make all the runs of different LLMs on a large text corpora available to other researchers in the hope that they can become a laboratory to measure LLMs performance as classifiers in a large scale real-world economic application.

7.3 Strategic Disclosure and the Text Source

One potential concern with using earnings calls to explore politically sensitive topics such as geoeconomic pressure is that the executives on the call may be reluctant to admit how they are affected or even tilt their responses to please governments listening in around the world. Indeed, this is the type of issue that cross-model validation or prompt perturbation cannot address because it would be an issue with the text itself and not the model’s reading of it.

In order to assess the extent of this issue, we turn to a second source of text: sell-side analyst reports. We redo our main analysis with analyst reports instead of earnings calls. Intuitively, external financial analysts are more independent (compared to firm insiders) and are paid and compensated to explain to their readers what the firm is doing and why. Figure 26 reports the net effects estimated via earnings calls in Panel (a) and analyst reports in Panel (b). Comparing across the two panels, we find strikingly similar patterns.²¹ Figure 27 echoes the same finding by making the comparison at detailed question level. We take average over observations for each firm-quarter-instrument pair shared in both corpora.

7.4 Comparison with Bag-of-Words Approach

One natural question about our analysis is how it would compare to a more traditional NLP approach relying on bag of words. While there is no straightforward way to measure things like the subjective causal effects using those earlier methods, the most straightforward comparison is for a simple object like whether or not a firm reports being affected by tariffs. Here, we compare the share of firms that we flag as being affected by tariffs based on our prompt with the share flagged by Bag-of-Words:

*“effect_any”: 1 if the firm discusses tariffs at any point in the call, and 0 otherwise.
Even if the term “tariffs” is not explicitly used throughout the call, you should return a
1 if the firm discusses impacts on its business that clearly relate to tariffs.*

The start of this prompt includes the definition, saying “Tariffs are defined as taxes imposed on imported foreign goods by the importing country. Tariffs are NOT export restrictions, quotas, embargoes, financial sanctions, boycotts, or non-tariff barriers.” For the bag of words approach, we simply flag any transcript in which the word “tariff(s)” appears. The comparison is plotted in Figure 28 . We see that they are very strongly correlated, with the dictionary approach flagging

²¹To make the two samples as comparable as possible, we restrict the sample period to 2011Q1 to 2025Q4 that both earnings calls and analyst reports have coverage for and exclude Chinese firms from Orbit because we do not have sufficient coverage of those from analyst reports compared to earnings calls.

Figure 24: **Inter-model agreement: confusion matrices**

(a) Tariffs

		Gemma		Llama		Qwen		GLM	
		0	1	0	1	0	1	0	1
Gemma	0			AC ₁	0.970	AC ₁	0.983	AC ₁	0.975
	1			κ	0.754	κ	0.878	κ	0.829
Llama	0	580,407	15,202			AC ₁	0.977	AC ₁	0.967
	1	1,522	27,687			κ	0.801	κ	0.732
Qwen	0	579,744	7,087	584,862	1,960			AC ₁	0.976
	1	2,144	35,794	10,697	27,241			κ	0.827
GLM	0	574,244	6,164	578,703	1,696	576,957	3,405		
	1	7,690	36,723	16,900	27,513	9,870	34,531		

(b) Sanctions

		Gemma		Llama		Qwen		GLM	
		0	1	0	1	0	1	0	1
Gemma	0			AC ₁	0.988	AC ₁	0.989	AC ₁	0.988
	1			κ	0.666	κ	0.718	κ	0.672
Llama	0	610,287	5,070			AC ₁	0.989	AC ₁	0.989
	1	2,106	7,355			κ	0.688	κ	0.646
Qwen	0	607,400	3,761	608,958	2,194			AC ₁	0.988
	1	2,790	8,556	4,164	7,182			κ	0.659
GLM	0	609,876	4,767	611,567	3,067	608,193	4,217		
	1	2,521	7,658	3,785	6,394	2,963	7,129		

(c) Export Controls

		Gemma		Llama		Qwen		GLM	
		0	1	0	1	0	1	0	1
Gemma	0			AC ₁	0.991	AC ₁	0.992	AC ₁	0.982
	1			κ	0.453	κ	0.633	κ	0.411
Llama	0	616,689	4,187			AC ₁	0.991	AC ₁	0.982
	1	1,534	2,408			κ	0.435	κ	0.306
Qwen	0	615,934	2,459	616,713	1,671			AC ₁	0.983
	1	2,247	4,134	4,112	2,269			κ	0.416
GLM	0	610,355	2,765	611,618	1,493	610,527	2,549		
	1	7,873	3,829	9,254	2,448	7,862	3,831		

Notes: The figure crosses pairwise four models' answers to the first step broad question on whether the transcript discusses export control, sanctions, tariffs, respectively. Each model pair is one 2×2 block. Below the diagonal are the four counts, with the shaded cell the number of transcripts both models flag (a = both flag, b and c = one only, d = neither). Above the diagonal are that block's agreement statistics: Gwet's AC1 and Cohen's κ , which both correct raw agreement for chance but differ in what they take chance to be, and positive agreement $2a/(2a + b + c)$. The diagonal is a model against itself and is left blank.

Figure 25: **Inter-model agreement: Gwet's AC_1 on flags in the long prompt**






Figure 26: **Firms' responses to geoeconomic pressure: calls vs. 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' earnings calls, while Panel (b) is based on single-firm analysts reports. We restrict the sample period of 2011Q1 to 2025Q4 that both earnings calls and analyst reports have coverage on and drop Chinese firms from Orbit because of insufficient analyst coverage.


Figure 27: Earnings Calls vs Analyst Reports



Notes: This figure plots, for each policy instrument and firm response variable, the average firm response measured on analyst reports (x-axis) against the same average measured on earnings calls (y-axis), with the 45-degree line for reference. Both corpora are classified with the same prompts. The sample is restricted to firm-quarters covered by both (2011Q1 to 2025Q4), excluding Orbit firms which we do not have analyst reports for. Each outcome is the share of firms reporting an increase minus the share reporting a decrease.

a larger share in most quarters. The difference is proportionally largest before 2018 and shrinks during the 2025 trade war. When we explore the differences, we find that firms, particularly energy and utility firms, use the term “tariff” in relation to their pricing schedule for fees. Our LLM-based methodology where the word tariff is defined as an import tax successfully filtered those uses of the word out, providing support for our approach.

Figure 28: **Comparing LLM and bag-of-words measures of tariff exposure**



Notes: This figure compares the share of firms classified as affected by tariffs using our LLM-based methodology versus a standard bag-of-words (dictionary) approach.

8 Conclusion

This paper develops a systematic methodology leveraging LLMs to measure and analyze the growing role of geoeconomic pressure in global economic relations. By analyzing large textual corpora of corporate earnings calls, we construct a comprehensive database of coercive economic actions and link them to firm-level economic responses.

First, export controls drive targeted firms to increase their research and development, with this effect particularly prominent for Chinese firms receiving US semiconductor export controls. These firms report increasing their R&D and domestic investment to develop substitutes for the blocked products. Second, sanctions (on Russia) have heterogeneous impact with firms in the sanctioning coalition loosening out, and firms in China and India benefiting by stepping into the vacuum. Third, tariffs, or the threat of tariffs, impact negatively the vast majority of firms in the imposing country, mostly the U.S. firms. Firms in the rest of the world are also negatively affected, mostly via the macroeconomic uncertainty, but also because of the cost of relocating production or because their own subsidiary is the importer of record.

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, 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, Chris Flanagan, Sebastian Heise, and David E Weinstein**, “Who is paying for the 2025 US tariffs?,” *Liberty Street Economics*, 2026, *12*.
- 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.
- Angelopoulos, Anastasios N., Stephen Bates, Clara Fannjian, Michael I. Jordan, and Tijana Zrnic**, “Prediction-powered inference,” *Science*, 2023, *382* (6671), eadf6056.
- Angelopoulos, Anastasios N., Stephen Bates, Michael I. Jordan, and Tijana Zrnic**, “PPI++: Efficient Prediction-Powered Inference,” *arXiv preprint arXiv:2311.01453*, 2024.
- Asirvatham, Hemanth, Elliott Moksiki, and Andrei Shleifer**, “GPT as a Measurement Tool,” Technical Report, National Bureau of Economic Research 2026.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub**, “The Macroeconomics of Tariff Shocks,” Technical Report, National Bureau of Economic Research 2025.
- Bai, Liang, Torsten Jaccard, and Sebastian Stumpner**, “Tariff Pass-through into Retail Prices: Measurement, Replacement, and Shrink-flation,” *Replacement, and Shrink-flation (December 19, 2025)*, 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.
- Becko, John S, Gene M Grossman, and Elhanan Helpman**, “Optimal tariffs with geopolitical

- alignment,” Technical Report, National Bureau of Economic Research 2025.
- Becko, John Sturm and Dangel G. O’Connor**, “Strategic (Dis)Integration,” 2024.
- Bonadio, Barthél  my, Zhen Huo, Elliot Kang, Andrei A. Levchenko, Nitya Pandalai-Nayar, Hiroshi Toma, and Petia Topalova**, “Playing with blocs: Quantifying decoupling,” 2024.
- Braun, Virginia and Victoria Clarke**, “Using thematic analysis in psychology,” *Qualitative research in psychology*, 2006, 3 (2), 77–101.
- 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.
- Carlson, Jacob and Melissa Dell**, “A Unifying Framework for Robust and Efficient Inference with Unstructured Data,” *Working Paper*, 2025.
- 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 Theory of Economic Coercion and Fragmentation,” *Accepted at the Journal of Political Economy*, 2024.
- Clayton, Christopher, Matteo Maggiori, and Jesse Schreger**, “The Political Economy of Geoeconomic Power,” 2025.
- Clayton, Christopher, Matteo Maggiori, and Jesse Schreger**, “A Framework for Geoeconomics,” *Econometrica*, 2026, 94 (1), 105–136.
- Coibion, Olivier and Yuri Gorodnichenko**, *Expectations matter: The new causal macroeco-*

- nomics of surveys and experiments*, Princeton University Press, 2026.
- Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar**, “How do firms form their expectations? New survey evidence,” *American Economic Review*, 2018, 108 (9), 2671–2713.
- Coppola, Antonio, Matteo Maggiori, Brent Neiman, and Jesse Schreger**, “Redrawing the Map of Global Capital Flows: The Role of Cross-Border Financing and Tax Havens,” *The Quarterly Journal of Economics*, 2021, 136 (3), 1499–1556.
- 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.
- 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.
- Egorov, Konstantin, Vasily Korovkin, Alexey Makarin, and Dzhamilya Nigmatulina**, “Export Sanctions: Rerouting, Disruption, and the Reach of Economic Warfare,” 2025.
- 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.
- Fan, Jingting, Kyle Handley, and Nuno Limão**, “Multilateral Export Control Regimes and Aggregate Exports,” in “AEA Papers and Proceedings,” Vol. 116 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203-2425 2026, pp. 109–113.
- Fang, Hanming, Ming Li, and Guangli Lu**, “Decoding China’s Industrial Policies,” Technical Report, National Bureau of Economic Research 2025.
- 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.
- Flaen, Aaron B, Ali Hortaçsu, Felix Tintelnot, Nicolás Urdaneta, and Daniel Xu**, “Who

- pays for tariffs along the supply chain? Evidence from European wine tariffs,” Technical Report, National Bureau of Economic Research 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.
- Ganapati, Sharat and Colin J Hottman**, “Did Foreigners Pay America’s Tariffs? Quantity Discounts, Scale Economies and Incomplete Pass-Through,” Technical Report, National Bureau of Economic Research 2026.
- 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 and Brent Neiman**, “The incidence of tariffs: Rates and reality,” Technical Report, National Bureau of Economic Research 2026.
- 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, Elhanan Helpman, and Stephen Redding**, “When tariffs disrupt global supply chains,” Technical Report 2021.
- Haaland, Ingar, Christopher Roth, Stefanie Stantcheva, and Johannes Wohlfart**, “Understanding economic behavior using open-ended survey data,” *Journal of Economic Literature*, 2025, 63 (4), 1244–1280.
- 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.
- 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.
- 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.
- Liu, Xueyue, Yu Liu, Alexey Makarin et al.**, “Export controls and innovation in sanctioned countries,” 2025.
- Ludwig, Jens, Sendhil Mullainathan, and Ashesh Rambachan**, “Large language models: An applied econometric framework,” *Working Paper*, 2025.
- Madsen, Erik and Andrea Prat**, “Competing Powers,” *Working paper*, 2026.

- Mattoo, Aaditya, Michele Ruta, and Robert W. Staiger**, “Geopolitics and the World Trading System,” 2024.
- Matvos, Gregor and Brent Neiman**, “Financial Sanctions and the Global Payments Network,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2026, (2026-49).
- Mayer, Thierry, Isabelle Mejean, and Mathias Thoenig**, “The fragmentation paradox: Derisking trade and global safety,” Technical Report, working paper 2025.
- Mohr, Cathrin and Christoph Trebesch**, “Geeconomics,” 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.
- Ndiaye, Abdoulaye**, “A Theory of International Boycotts,” Technical Report, CESifo Working Paper 2024.
- Nigmatulina, Dzhamilya**, “Sanctions and misallocation. How sanctioned firms won and Russia lost,” 2022.
- Ossa, Ralph and Stephen J Redding**, “The economics of tariffs,” Technical Report, National Bureau of Economic Research 2026.
- 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.
- 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.
- Stantcheva, Stefanie**, “How to run surveys: A guide to creating your own identifying variation and revealing the invisible,” *Annual Review of Economics*, 2023, 15 (1), 205–234.
- Thoenig, Mathias**, “Trade Policy in the Shadow of War: Quantitative Tools for Geoeconomics,” 2023.

ONLINE APPENDIX FOR “GEOECONOMIC PRESSURE”

Christopher Clayton Antonio Coppola Vanessa Hu
Matteo Maggiori Jesse Schreger Serdil Tinda

September 2026

A Dataset Fields

This appendix documents every field returned by the second-step, instrument-specific long prompts. Fields are classified by all three instrument prompts—tariffs, export controls, and sanctions—except those marked “Sanctions,” which appear only in the sanctions prompt.

Table A.1: **Fields extracted by the instrument-specific prompts**

Field	Type	Definition	Domain
<i>Policy detection</i>			
effect_any	Binary	Firm discusses the policy at any point in the call.	All
effect_{current,future}	Binary	Discusses policy currently in effect / possibly imposed in the future that impacts its business.	All
<i>Policy direction</i>			
policy_{up,down}	Binary	Discusses new or intensified / removed or eased policy.	All
<i>Geographic and product scope</i>			
countries_imposing	Countries	Countries whose policy the firm discusses.	All
countries_receiving	Countries	Countries or entities targeted by the policy discussed.	All
product_receiving	Text	Goods or services the firm sells that are targeted by the policy.	All
input_receiving	Text	Goods or services the firm buys that are targeted by the policy.	All
<i>Investment</i>			
investment_{up,down}	Binary	Increasing / decreasing investment.	All
countries_investment_{up,down}	Countries	Countries where investment is rising / falling.	All
domestic_investment_{up,down}	Binary	Increasing / decreasing investment in its headquarter country.	All

continued on next page

Table A.1 (continued)

Field	Type	Definition	Domain
<i>Employment</i>			
labor_{up,down}	Binary	Increasing / decreasing workers or hours worked.	All
countries_labor_{up,down}	Countries	Countries where employment is rising / falling.	All
domestic_labor_{up,down}	Binary	Increasing / decreasing employment in its headquarter country.	All
<i>Prices</i>			
sales_price_{up,down}	Binary	Raising / lowering the price it charges for products sold.	All
countries_sales_price_{up,down}	Countries	Countries where sales prices are rising / falling.	All
input_price_{up,down}	Binary	Facing / expecting higher / lower prices on inputs purchased.	All
countries_input_price_{up,down}	Countries	Countries where input prices are rising / falling.	All
<i>Quantities</i>			
sales_{up,down}	Binary	Increasing / decreasing quantities of products sold.	All
countries_sales_{up,down}	Countries	Countries where sales quantities are rising / falling.	All
input_{up,down}	Binary	Increasing / decreasing quantities of inputs purchased.	All
countries_input_{up,down}	Countries	Countries where input quantities are rising / falling.	All
<i>Inventory</i>			
inventory_product_{up,down}	Binary	Increasing / decreasing inventory of products it sells.	All
inventory_input_{up,down}	Binary	Increasing / decreasing inventory of inputs it purchases.	All
<i>Research and development</i>			
rd_{up,down}	Binary	Increasing / decreasing research and development.	All
<i>Strategic responses</i>			
delay	Binary	Delaying a decision in response to the policy.	All
supply_chain_adj	Binary	Adjusting / plans to adjust its supply chain.	All
countries_supply_chain_adj_{towards,away}	Countries	Countries toward which supply-chain exposure rises / away from which it falls.	All

continued on next page

Table A.1 (continued)

Field	Type	Definition	Domain
lobbying_{up,down}	Binary	Increasing / decreasing lobbying activity.	All
financing_adj	Binary	Adjusting / plans to adjust financing decisions.	All
payments_adj	Binary	Adjusting how it processes payments or financial transactions.	Sanctions
<i>Profitability and overall impact</i>			
profit_margin_{up,down}	Binary	Profit margins increased / decreased or expected to.	All
{negative,positive}_impact	Binary	Reports any negative / positive impact on its business.	All
<i>Sanctions mechanism</i>			
trade_sanctions	Binary	Sanctions on the import/export of goods or their transport (ships, ports, planes).	Sanctions
financial_sanctions	Binary	Sanctions on financial assets, transactions, or services (e.g. asset freezes, payment restrictions).	Sanctions
<i>Narrative output</i>			
summary	Text	≤100-word summary of how the firm is (or may be) affected; “Not affected” if none.	All
evaluation	Text	≤100-word self-assessment of how well the summary agrees with the model’s analysis.	All

Notes: Binary fields equal 1 when the firm explicitly or implicitly attributes the action to current or future policy, and 0 otherwise. “Countries” fields list the relevant countries discussed, or are missing when none are specified. Field names and definitions follow the version-8 prompts.

B Full LLM Prompts

This appendix reproduces the full text of the prompts used in our analysis. Section B.1 presents the first-step broad prompt, which is used to flag whether a document discusses various forms of geoeconomic pressure. Section B.2 presents the second-step long prompt for tariffs, Section B.3 for sanctions, and B.4 for export controls. Sections B.5 and onward present the prompts for the reasoning analysis.

B.1 First-Step Broad Prompt

You are assisting me in analyzing companies' earnings calls transcript. The transcript is provided below.

<transcript>

{}

</transcript>

Task Overview

Your goal is to determine whether the company discusses its business being affected by various economic policies. We define each in turn.

- 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.
- Sanctions are government-imposed restrictions on trade or financial transactions designed to coerce, punish, or deter targeted firms or governments.
- 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. This is in contrast to import tariffs which are instead taxes imposed by the importing country.
- 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.
- Investment screening reviews for national security concerns. Examples of investment screening reviews by each country are provided below:
 - * United States Committee on Foreign Investment in the United States (CFIUS)
 - * Canada Investment Canada Act (ICA) and Foreign Investment Review and Economic Security (FIRES)
 - * European Union EU Screening Mechanism (2019)
 - * Germany Foreign Trade and Payments Act (AWG) and Ordinance (AWV)
 - * France Code Monétaire et Financier
 - * Israel - Advisory Committee for National Security Aspects of Foreign Investments
 - * Italy Golden Power Rules
 - * United Kingdom National Security and Investment Act (NSIA, 2021)
 - * Norway Security Act screening, chapter 10
 - * Australia Foreign Investment Review Board (FIRB)
 - * Japan Foreign Exchange and Foreign Trade Act (FEFTA)
 - * China National Security Review of Foreign Investment
 - * India Press Note 3 (2020) and FDI Policy
 - * South Korea Foreign Investment Promotion Act (FIPA) & related security screening rules
 - * Taiwan - Department of Investment Review (DIR) or Investment Commission Review
 - * Russia Strategic Sectors Law
 - * United Arab Emirates FDI Law & Sector-Specific Reviews
 - * Netherlands - Act on Security Screening of Investments, Mergers and Acquisitions (Wet veiligheidsstoets investeringen, fusies en overnames, or Vifo Act)
 - * Belgium - Interfederal Screening Commission (Comité de filtrage interfederaal/Interfederaale Screeningscommissie)

```

* Poland - Act on Control of Certain Investments
* Sweden - Protective Security Act (S kerhetsskyddslagen)
* Singapore - Significant Investments Review Act
* Saudi Arabia - Permanent Ministerial Committee for Examining Foreign Investments

## Output Requirements ##
You must provide your response as a JSON object containing exactly 7 fields as specified
  below. Each field must be included with the exact field name provided.

## JSON Schema ##
Your response must be a valid JSON object with the following fields:

- "analysis": 300 words or fewer of analysis to cover the following points: whether
  tariffs, sanctions, export controls, boycotts, or investment screening are discussed
  explicitly or implicitly. The nature and details of the policies that are discussed
  , and any potential impact on the firm.
- "tariffs_any": 1 if the firm explicitly or implicitly discusses tariffs at any point
  in the call, and 0 otherwise.
- "sanctions_any": 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.
- "export_controls_any": 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.
- "boycotts_any": 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.
- "investment_screening_any": 1 if the firm explicitly or implicitly at any point in the
  call discusses investment screening reviews, and 0 otherwise.
- "summary": summary of 200 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, and investment screening. Omit all unrelated content. If the
  firm is not affected by any of these policies, write "Not affected."

## 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.

## 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: ##

- 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.
- Only generate valid JSON object, no markdown or other formats.

/no_think

```

B.2 Second-Step Long Prompt: Tariffs

```

You will analyze a company earnings call transcript to identify and extract information
about how tariff policies affect the firm's decisions and business operations.

<transcript>
{}
</transcript>

## Task Overview ##
Analyze the transcript above to determine whether the firm reports that its decisions
are being affected by tariff policies. Tariffs are defined as taxes imposed on
imported foreign goods by the importing country. Tariffs are NOT export restrictions
, quotas, embargoes, financial sanctions, boycotts, or non-tariff barriers.

## Response Instructions: Part 1 (Analysis) ##
The first part of your response should be an analysis of whether the firm is reporting
that its 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 firm 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 firm discusses tariffs on the goods it sells, or on those that it 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 firms' current profits.
- Any potential impact on the firm's future profits.

```

- Any impacts on the firms' 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 firms' behavior (e.g., if the firm reports importing less, report from which destination).

Response Instructions: Part 2 (Structured JSON Output)

You must provide your response as a JSON object containing exactly 56 fields as specified below. Each field must be included with the exact field name provided.

- "effect_any": 1 if the firm discusses tariffs at any point in the call, and 0 otherwise. Even if the term "tariffs" is not explicitly used throughout the call, you should return a 1 if the firm discusses impacts on its business that clearly relate to tariffs.
- effect_current : 1 if the firm explicitly or implicitly discusses tariffs that are currently in effect and that impact the firm business, and 0 otherwise. Please do not classify tariffs that might be imposed in the future as a 1 in this flag.
- effect_future : 1 if the firm explicitly or implicitly discusses tariffs that might be imposed in the future but are not currently in effect and that might impact the firm business, and 0 otherwise. Please do not classify tariffs that are currently in effect as a 1 in this flag.
- policy_up : 1 if the firm explicitly or implicitly discusses new tariffs or increases in the range or magnitude of existing tariffs that it faces or might face in the future, and 0 otherwise. Please do not classify decreases or removal of tariffs as a 1 in this flag.
- policy_down : 1 if the firm explicitly or implicitly discusses removal of tariffs or decreases in the range or magnitude of existing tariffs that it faces or might face in the future, and 0 otherwise. Please do not classify increases or new tariffs as a 1 in this flag.
- "countries_imposing": countries whose tariffs policy the firm discusses, if any. For example, if the firm reports concerns about tariffs imposed by the US government on goods imported from China, this field should say "USA"
- "countries_receiving": countries targeted by the tariffs discussed by the firm, if any. For example, if the firm reports concerns about tariffs imposed by the US government on goods imported from China, this field should say "China".
- product_receiving : goods that the firm sells that are targeted by the tariffs, if any. For example, if the firm reports current or future tariffs imposed by the US government on steel that the firm sells to its US customers, this field should say steel .
- input_receiving : goods that the firm buys that are targeted by tariffs, if any. For example, if the firm reports current or future tariffs imposed by the US government on steel that the firm buys from Mexico, this field should say steel .
- "investment_up": 1 if the firm says it 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 company explicitly attributes a change in its investment decisions to tariffs.
- "investment_down": 1 if the firm says it 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 company explicitly attributes a change in its investment decisions to tariffs.
- "countries_investment_up": countries where the firm reports 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".
- "countries_investment_down": countries where the firm reports 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".

- "domestic_investment_up": 1 if the firm says it 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 company explicitly attributes a change in its investment decisions to tariffs.
- "domestic_investment_down": 1 if the firm says it 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 company explicitly attributes a change in its investment decisions to tariffs.
- "labor_up": 1 if the firm says it 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 company explicitly attributes a change in its employment decisions to tariffs.
- "labor_down": 1 if the firm says it 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 company explicitly attributes a change in its employment decisions to tariffs.
- "countries_labor_up": countries where the firm reports 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".
- "countries_labor_down": countries where the firm reports 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".
- "domestic_labor_up": 1 if the firm says it 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 company explicitly attributes a change in its employment decisions to tariffs.
- "domestic_labor_down": 1 if the firm says it 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 company explicitly attributes a change in its employment decisions to tariffs.
- "sales_price_up": 1 if the firm says it 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 company explicitly attributes a change in its product pricing decisions to tariffs.
- "sales_price_down": 1 if the firm says it 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 company explicitly attributes a change in its product pricing decisions to tariffs.
- "countries_sales_price_up": countries where the firm reports 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".
- "countries_sales_price_down": countries where the firm reports 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".
- "input_price_up": 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 company explicitly attributes a change in its input prices to tariffs.
- "input_price_down": 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 company explicitly attributes a change in its input prices to tariffs.
- "countries_input_price_up": countries where the firm reports 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".

- "countries_input_price_down": countries where the firm reports 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".
- "sales_up": 1 if the firm says it 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 company explicitly attributes a change in its product sales to tariffs.
- "sales_down": 1 if the firm says it 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 company explicitly attributes a change in its product sales to tariffs.
- "countries_sales_up": countries where the firm reports 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".
- "countries_sales_down": countries where the firm reports 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".
- "input_up": 1 if the firm says it 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 company explicitly attributes a change in its input quantities to tariffs.
- "input_down": 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 company explicitly attributes a change in its input quantities to tariffs.
- "countries_input_up": countries where the firm reports 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".
- "countries_input_down": countries where the firm reports 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".
- "inventory_product_up": 1 if the firm says it 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 company explicitly attributes a change in its inventory decisions to tariffs.
- "inventory_product_down": 1 if the firm says it 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 company explicitly attributes a change in its inventory decisions to tariffs.
- "inventory_input_up": 1 if the firm says it 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 company explicitly attributes a change in its inventory decisions to tariffs.
- "inventory_input_down": 1 if the firm says it 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 company explicitly attributes a change in its inventory decisions to tariffs.
- "rd_up": 1 if the firm says it 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 company explicitly attributes a change in its research and development decisions to tariffs.
- "rd_down": 1 if the firm says it 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 company explicitly attributes a change in its research and development decisions to tariffs.

```

- "delay": 1 if the firm says it is delaying a decision as a result of current or future
  tariffs. Only set the flag to 1 if the company explicitly attributes the delay to
  tariffs.
- "supply_chain_adj": 1 if the firm says it is adjusting its supply chain as a result of
  current or future tariffs. Only set the flag to 1 if the company is explicitly
  adjusting supply chains due to tariffs.
- "countries_supply_chain_adj_towards": countries where the firm reports 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".
- "countries_supply_chain_adj_away": countries where the firm reports 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".
- "lobbying_up": 1 if the firm says it 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 company explicitly attributes a change in its lobbying policy
  to tariffs.
- "lobbying_down": 1 if the firm says it 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 company explicitly attributes a change in its lobbying policy
  to tariffs.
- "profit_margin_up": 1 if the firm says its 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 company explicitly attributes a change in its profit
  margins to tariffs.
- "profit_margin_down": 1 if the firm says its 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 company explicitly attributes a change in its profit
  margins to tariffs.
- "financing_adj": 1 if the firm says it is adjusting its financing decisions in
  response to current or future tariffs. Only set the flag to 1 if the company is
  explicitly adjusting its financing decisions due to tariffs.
- "negative_impact": 1 if the firm reports any negative impact on its business as a
  result of current or future tariffs and 0 otherwise.
- "positive_impact": 1 if the firm reports any positive impact on its business as a
  result of current or future tariffs and 0 otherwise.
- "summary": 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."
- "evaluation": evaluate how well the summary agrees with your initial analysis. Keep
  the evaluation to 100 words or less.

## Important Requirements ##
- Only set flags to 1 when the company EXPLICITLY attributes changes to tariffs
- Use "NaN" for country fields when no country is specified or unclear
- 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.
- Tariffs must be taxes imposed by the importing country
- First write your analysis in <ANALYSIS> tags, then output the JSON object.

/no_think

```

B.3 Second-Step Long Prompt: Sanctions

You will analyze a company earnings call transcript to identify and extract information about how sanctions policies affect the firm's decisions and business operations.

```
<transcript>
{}
</transcript>
```

Task Overview

Analyze the transcript above to determine whether the firm reports that its decisions are being affected by sanctions. Sanctions are government-imposed restrictions on trade or financial transactions designed to coerce, punish, or deter targeted firms or governments. Tariffs, which are taxes on imports, are not sanctions.

Response Instructions: Part 1 (Analysis)

The first part of your response should be an analysis of whether the firm is reporting that its decisions are being affected by sanctions. Sanctions are government-imposed restrictions on trade or financial transactions designed to coerce, punish, or deter targeted firms or governments. Tariffs, which are taxes on imports, are not sanctions.

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 sanctions are discussed explicitly or implicitly (e.g., by not using the phrase sanctions but referencing impacts on the firm's business that clearly relate to sanctions).
- Whether the firm discusses current sanctions (i.e., sanctions that have already been imposed) or the potential of future sanctions (i.e., sanctions that have not yet been imposed).
- Whether the firm discusses sanctions on the goods or services it sells or the goods or services it buys.
- The nature and details of the sanctions policies that are discussed.
- The countries that are imposing the sanctions and those that are subject to the sanctions.
- Any impacts on the firms' current profits.
- Any potential impact on the firm's future profits.
- Any impacts on the firms' 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 firms' behavior (e.g., if the firm reports importing less, report from which destination).

Response Instructions: Part 2 (Structured JSON Output)

You must provide your response as a JSON object containing exactly 59 fields as specified below. Each field must be included with the exact field name provided.

- "effect_any": 1 if the firm discusses sanctions at any point in the call, and 0 otherwise. Even if the term sanctions is not explicitly used throughout the call, you should return a 1 if the firm discusses impacts on its business that clearly relate to sanctions.
- effect_current : 1 if the firm explicitly or implicitly discusses sanctions that are currently in effect and that impact the firm business, and 0 otherwise. Please do not classify sanctions that might imposed in the future as a 1 in this flag.

- `effect_future` : 1 if the firm explicitly or implicitly discusses sanctions that might be imposed in the future but are not currently in effect and that might impact the firm business, and 0 otherwise. Please do not classify sanctions that are currently in effect as a 1 in this flag.
- `policy_up` : 1 if the firm explicitly or implicitly discusses being affected by new sanctions being imposed or the strengthening of existing sanctions now or in the future, and 0 otherwise. Please do not classify decreases or lifting of sanctions as a 1 in this flag.
- `policy_down` : 1 if the firm explicitly or implicitly discusses being affected by existing sanctions being lifted or relaxed now or in the future, and 0 otherwise. Please do not classify the introduction of new sanctions or strengthening of existing ones as a 1 in this flag.
- `"countries_imposing"`: countries whose sanctions policy the firm discusses, if any. For example, if the firm reports current or future sanctions imposed by the US government on goods or services exported by Russia, this field should say "USA".
- `"countries_receiving"`: entities (countries or firms) targeted by the sanctions discussed by the firm, if any. For example, if the firm reports current or future sanctions imposed by the US government on goods or services exported by Russia, this field should say "Russia".
- `product_receiving` : goods or services that the firm sells that are targeted by the sanctions, if any. For example, if the firm reports current or future sanctions imposed by the US government on military equipment that it exports to Russia, this field should say "military equipment".
- `input_receiving` : goods or services that the firm buys that are targeted by the sanctions, if any. For example, if the firm reports current or future sanctions imposed by the US government on oil that it imports from Russia, this field should say "oil".
- `"investment_up"`: 1 if the firm says it is increasing or planning to increase its investment as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to sanctions.
- `"investment_down"`: 1 if the firm says it is decreasing or planning to decrease its investment as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to sanctions.
- `"countries_investment_up"`: countries where the firm reports increasing or planning to increase investment as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- `"countries_investment_down"`: countries where the firm reports decreasing or planning to decrease investment as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- `"domestic_investment_up"`: 1 if the firm says it is increasing or planning to increase its investment in its headquarter country as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to sanctions.
- `"domestic_investment_down"`: 1 if the firm says it is decreasing or planning to decrease its investment in its headquarter country as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to sanctions.
- `"labor_up"`: 1 if the firm says it is increasing or planning to increase its number of workers or their hours worked as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to sanctions.
- `"labor_down"`: 1 if the firm says it is decreasing or planning to decrease its number of workers or their hours worked as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in

its employment decisions to sanctions.

- "countries_labor_up": countries where the firm reports increasing or planning to increase its number of workers or their hours worked as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "countries_labor_down": countries where the firm reports decreasing or planning to decrease its number of workers or their hours worked as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "domestic_labor_up": 1 if the firm says it 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 sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to sanctions.
- "domestic_labor_down": 1 if the firm says it 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 sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to sanctions.
- "sales_price_up": 1 if the firm says it is increasing or planning to increase the price it charges for the products it sells as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product pricing decisions to sanctions.
- "sales_price_down": 1 if the firm says it is decreasing or planning to decrease the price it charges for the products it sells as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product pricing decisions to sanctions.
- "countries_sales_price_up": countries where the firm reports increasing or planning to increase the price it charges for the products it sells as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "countries_sales_price_down": countries where the firm reports decreasing or planning to decrease the price it charges for the products it sells as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "input_price_up": 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 sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input prices to sanctions.
- "input_price_down": 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 sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input prices to sanctions.
- "countries_input_price_up": countries where the firm reports facing or expecting to face higher prices on the inputs it purchases as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "countries_input_price_down": countries where the firm reports facing or expecting to face lower prices on the inputs it purchases as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "sales_up": 1 if the firm says it is increasing or planning to increase the quantities of the products that it sells as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product sales to sanctions.
- "sales_down": 1 if the firm says it is decreasing or planning to decrease the quantities of the products that it sells as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product sales to sanctions.
- "countries_sales_up": countries where the firm reports increasing or planning to increase the quantities of the products that it sells as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "countries_sales_down": countries where the firm reports decreasing or planning to decrease the quantities of the products that it sells as a result of current or

future sanctions. If no country is specified or it is unclear, report "NaN".

- "input_up": 1 if the firm says it is increasing or planning to increase the quantities of the inputs that it is purchasing as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input quantities to sanctions.
- "input_down": 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 sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input quantities to sanctions.
- "countries_input_up": countries where the firm reports increasing or planning to increase the quantities of the inputs that it is purchasing as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "countries_input_down": countries where the firm reports decreasing or planning to decrease the quantities of the products as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "inventory_product_up": 1 if the firm says it is increasing or expecting to increase the amount of inventory of the products it sells as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to sanctions.
- "inventory_product_down": 1 if the firm says it is decreasing or expecting to decrease the amount of inventory of the products it sells as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to sanctions.
- "inventory_input_up": 1 if the firm says it is increasing or expecting to increase the amount of inventory of inputs it purchases as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to sanctions.
- "inventory_input_down": 1 if the firm says it is decreasing or expecting to decrease the amount of inventory of inputs it purchases as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to sanctions.
- "rd_up": 1 if the firm says it is increasing or expecting to increase the amount of research and development it undertakes as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its research and development decisions to sanctions.
- "rd_down": 1 if the firm says it is decreasing or expecting to decrease the amount of research and development it undertakes as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its research and development decisions to sanctions.
- "delay": 1 if the firm says it is delaying a decision as a result of current or future sanctions. Only set the flag to 1 if the company explicitly attributes the delay to sanctions.
- "supply_chain_adj": 1 if the firm says it is adjusting or planning to adjust its supply chain as a result of current or future sanctions. Only set the flag to 1 if the company explicitly attributes adjusting its supply chains to sanctions.
- "countries_supply_chain_adj_towards": countries where the firm reports increasing or planning to increase the exposure of its supply chain as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "countries_supply_chain_adj_away": countries where the firm reports reducing or planning to reduce the exposure of its supply chain as a result of current or future sanctions. If no country is specified or it is unclear, report "NaN".
- "lobbying_up": 1 if the firm says it is increasing or expecting to increase its lobbying activity as a result of current or future sanctions and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its lobbying policy to sanctions.

```

- "lobbying_down": 1 if the firm says it is decreasing or expecting to decrease its
  lobbying activity as a result of current or future sanctions and 0 otherwise. Only
  set the flag to 1 if the company explicitly attributes a change in its lobbying
  policy to sanctions.
- "profit_margin_up": 1 if the firm says its profit margins have increased or are
  expected to increase as a result of current or future sanctions and 0 otherwise.
  Only set the flag to 1 if the company explicitly attributes a change in its profit
  margins to sanctions.
- "profit_margin_down": 1 if the firm says its profit margins have decreased or are
  expected to decreased as a result of current or future sanctions and 0 otherwise.
  Only set the flag to 1 if the company explicitly attributes a change in its profit
  margins to sanctions.
- "financing_adj": 1 if the firm says it is adjusting or planning to adjust its
  financing decisions in response to current or future sanctions. Only set the flag to
  1 if the company explicitly attributes the adjustment of its financing decisions to
  sanctions.
- "negative_impact": 1 if the firm reports any negative impact on its business as a
  result of current or future sanctions and 0 otherwise.
- "positive_impact": 1 if the firm reports any positive impact on its business as a
  result of current or future sanctions and 0 otherwise.
-   trade_sanctions       : 1 if the firm discusses sanctions at any point in the call that
  are related to the import or export of goods or means of transport of those goods (
  restrictions on ships, ports, planes), and 0 otherwise.
- "financial_sanctions": 1 if the firm discusses sanctions at any point in the call
  that are related to financial assets, financial transactions, or financial services
  (e.g. asset freezes, or restrictions to international payments), and 0 otherwise.
-   payments_adj        ": 1 if the firm says it is adjusting or planning to adjust how it
  processes payments or financial transactions in response to current or future
  sanctions. Only set the flag to 1 if the company explicitly attributes the
  adjustment of its processing of payments or financial transactions to sanctions.
- "summary": summary of 100 words or fewer that captures only how the firm is (or may be
  ) affected by sanctions. Omit all unrelated content. If the firm is not affected by
  any tariffs, write "Not affected."
- "evaluation": evaluate how well the summary agrees with your initial analysis. Keep
  the evaluation to 100 words or less.

## Important Notes ##
- Do not consider tariffs, taxes imposed by the importing country, as indications of
  sanctions.
- First write your analysis in <ANALYSIS> tags, then output the JSON object.

/no_think

```

B.4 Second-Step Long Prompt: Export Controls

```

You will analyze a company earnings call transcript to identify and extract information
about how export control policies affect the firm's decisions and business
operations.

<transcript>
{}
</transcript>

```

Task Overview

Analyze the transcript above to determine whether the firm reports that its decisions are being affected by export control policies. Export controls are defined as restrictions on the sale of certain goods, technologies, or services. These must be imposed by the exporting country in order to be classified as export controls. Export controls are not tariffs, financial sanctions, or boycotts.

Response Instructions: Part 1 (Analysis)

The first part of your response should be an analysis of whether the firm is reporting that its decisions are being affected by export control policies. Export controls are defined as restrictions on the sale of certain goods, technologies, or services. These must be imposed by the exporting country in order to be classified as export controls. Export controls are not tariffs, financial sanctions, or boycotts.

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 export controls are discussed explicitly or implicitly (e.g., by not using the phrase export controls but referencing impacts on the firm's business that clearly relate to export controls).
- Whether the firm discusses current export controls (i.e., export controls that have already been imposed) or the potential of future export controls (i.e., export controls that have not yet been imposed).
- Whether the firm discusses export controls on the goods or services it sells or the goods or services it buys.
- The nature and details of the export control policies that are discussed.
- The countries that are imposing the export controls and those that are subject to the export controls.
- Any impacts on the firms' current profits.
- Any potential impact on the firm's future profits.
- Any impacts on the firms' 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 firms' behavior (e.g., if the firm reports importing less, report from which destination).

Response Instructions: Part 2 (Structured JSON Output)

You must provide your response as a JSON object containing exactly 56 fields as specified below. Each field must be included with the exact field name provided.

- "effect_any": 1 if the firm discusses export controls at any point in the call, and 0 otherwise. Even if the term export controls is not explicitly used throughout the call, you should return a 1 if the firm discusses impacts on its business that clearly relate to export controls.
- effect_current : 1 if the firm explicitly or implicitly discusses export controls that are currently in effect and that impact the firm business, and 0 otherwise. Please do not classify export controls that might be imposed in the future as a 1 in this flag.
- effect_future : 1 if the firm explicitly or implicitly discusses export controls that might be imposed in the future but are not currently in effect and that might impact the firm business, and 0 otherwise. Please do not classify export controls that are currently in effect as a 1 in this flag.
- policy_up : 1 if the firm explicitly or implicitly discusses new export controls or increases in the range or magnitude of existing export controls that it faces or

- might face in the future, and 0 otherwise. Please do not classify decreases or removal of export controls as a 1 in this flag.
- `policy_down` : 1 if the firm explicitly or implicitly discusses removal of export controls or decreases in the range or magnitude of existing export controls that it faces or might face in the future, and 0 otherwise. Please do not classify increases or new export controls as a 1 in this flag.
 - `"countries_imposing"`: countries whose export control policy the firm discusses, if any. For example, if the firm reports current or future export controls imposed by the US government on goods or services exported to China, this field should say "USA"
 - `"countries_receiving"`: entities (countries or firms) targeted by the export controls discussed by the firm, if any. For example, if the firm reports current or future export controls imposed by the US government on goods or services exported to China, this field should say "China".
 - `product_receiving` : goods or services that the firm sells that are targeted by the export controls, if any. For example, if the firm reports current or future export controls imposed by the US government on semiconductors that it exports to China, this field should say `semiconductors` .
 - `input_receiving` : goods or services that the firm buys that are targeted by the export controls, if any. For example, if the firm reports current or future export controls imposed by the US government on semiconductors that it buys from the US, this field should say `semiconductors` .
 - `"investment_up"`: 1 if the firm says it is increasing or planning to increase its investment as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to export controls.
 - `"investment_down"`: 1 if the firm says it is decreasing or planning to decrease its investment as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to export controls.
 - `"countries_investment_up"`: countries where the firm reports increasing or planning to increase investment as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
 - `"countries_investment_down"`: countries where the firm reports decreasing or planning to decrease investment as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
 - `"domestic_investment_up"`: 1 if the firm says it is increasing or planning to increase its investment in its headquarter country as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to export controls.
 - `"domestic_investment_down"`: 1 if the firm says it is decreasing or planning to decrease its investment in its headquarter country as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its investment decisions to export controls.
 - `"labor_up"`: 1 if the firm says it is increasing or planning to increase its number of workers or their hours worked as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to export controls.
 - `"labor_down"`: 1 if the firm says it is decreasing or planning to decrease its number of workers or their hours worked as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to export controls.
 - `"countries_labor_up"`: countries where the firm reports increasing or planning to increase its number of workers or their hours worked as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".

- "countries_labor_down": countries where the firm reports decreasing or planning to decrease its number of workers or their hours worked as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "domestic_labor_up": 1 if the firm says it 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 export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to export controls.
- "domestic_labor_down": 1 if the firm says it 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 export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its employment decisions to export controls.
- "sales_price_up": 1 if the firm says it is increasing or planning to increase the price it charges for the products it sells as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product pricing decisions to export controls.
- "sales_price_down": 1 if the firm says it is decreasing or planning to decrease the price it charges for the products it sells as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product pricing decisions to export controls.
- "countries_sales_price_up": countries where the firm reports increasing or planning to increase the price it charges for the products it sells as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "countries_sales_price_down": countries where the firm reports decreasing or planning to decrease the price it charges for the products it sells as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "input_price_up": 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 export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input prices to export controls.
- "input_price_down": 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 export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input prices to export controls.
- "countries_input_price_up": countries where the firm reports facing or expecting to face higher prices on the inputs it purchases as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "countries_input_price_down": countries where the firm reports facing or expecting to face lower prices on the inputs it purchases as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "sales_up": 1 if the firm says it is increasing or planning to increase the quantities of the products that it sells as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product sales to export controls.
- "sales_down": 1 if the firm says it is decreasing or planning to decrease the quantities of the products that it sells as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its product sales to export controls.
- "countries_sales_up": countries where the firm reports increasing or planning to increase the quantities of the products that it sells as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "countries_sales_down": countries where the firm reports decreasing or planning to decrease the quantities of the products that it sells as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".

- "input_up": 1 if the firm says it is increasing or planning to increase the quantities of the inputs that it is purchasing as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input quantities to export controls.
- "input_down": 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 export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its input quantities to export controls.
- "countries_input_up": countries where the firm reports increasing or planning to increase the quantities of the inputs that it is purchasing as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "countries_input_down": countries where the firm reports decreasing or planning to decrease the quantities of the products as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "inventory_product_up": 1 if the firm says it is increasing or expecting to increase the amount of inventory of the products it sells as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to export controls.
- "inventory_product_down": 1 if the firm says it is decreasing or expecting to decrease the amount of inventory of the products it sells as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to export controls.
- "inventory_input_up": 1 if the firm says it is increasing or expecting to increase the amount of inventory of inputs it purchases as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to export controls.
- "inventory_input_down": 1 if the firm says it is decreasing or expecting to decrease the amount of inventory of inputs it purchases as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its inventory decisions to export controls.
- "rd_up": 1 if the firm says it is increasing or expecting to increase the amount of research and development it undertakes as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its research and development decisions to export controls.
- "rd_down": 1 if the firm says it is decreasing or expecting to decrease the amount of research and development it undertakes as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its research and development decisions to export controls.
- "delay": 1 if the firm says it is delaying a decision as a result of current or future export controls. Only set the flag to 1 if the company explicitly attributes the delay to export controls.
- "supply_chain_adj": 1 if the firm says it is adjusting or planning to adjust its supply chain as a result of current or future export controls. Only set the flag to 1 if the company explicitly attributes adjusting its supply chains to export controls.
- "countries_supply_chain_adj_towards": countries where the firm reports increasing or planning to increase the exposure of its supply chain as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "countries_supply_chain_adj_away": countries where the firm reports reducing or planning to reduce the exposure of its supply chain as a result of current or future export controls. If no country is specified or it is unclear, report "NaN".
- "lobbying_up": 1 if the firm says it is increasing or expecting to increase its lobbying activity as a result of current or future export controls and 0 otherwise. Only set the flag to 1 if the company explicitly attributes a change in its lobbying policy to export controls.

```

- "lobbying_down": 1 if the firm says it is decreasing or expecting to decrease its
  lobbying activity as a result of current or future export controls and 0 otherwise.
  Only set the flag to 1 if the company explicitly attributes a change in its lobbying
  policy to export controls.
- "profit_margin_up": 1 if the firm says its profit margins have increased or are
  expected to increase as a result of current or future export controls and 0
  otherwise. Only set the flag to 1 if the company explicitly attributes a change in
  its profit margins to export controls.
- "profit_margin_down": 1 if the firm says its profit margins have decreased or are
  expected to decreased as a result of current or future export controls and 0
  otherwise. Only set the flag to 1 if the company explicitly attributes a change in
  its profit margins to export controls.
- "financing_adj": 1 if the firm says it is adjusting or planning to adjust its
  financing decisions in response to current or future export controls. Only set the
  flag to 1 if the company explicitly attributes the adjustment of its financing
  decisions to export controls.
- "negative_impact": 1 if the firm reports any negative impact on its business as a
  result of current or future export controls and 0 otherwise.
- "positive_impact": 1 if the firm reports any positive impact on its business as a
  result of current or future export controls and 0 otherwise.
- "summary": summary of 100 words or fewer that captures only how the firm is or may be
  affected by export controls. Omit all unrelated content. If the firm is not affected
  by any export controls, write "Not affected".
- "evaluation": evaluate how well the summary agrees with your initial analysis. Keep
  the evaluation to 100 words or less.

## Important Notes ##
- Do not consider tariffs (taxes imposed by the importing country) or generic financial
  sanctions as indications of export controls. Export controls must be restrictions
  imposed by the exporting country.
- First write your analysis in <ANALYSIS> tags, then output the JSON object.

/no_think

```

B.5 Reason Analysis Prompt: Step 1 Reason Extraction (Tariffs)

```

You are an expert analyst studying how firms are affected by tariffs. You are reading
the full transcript of a firm's earnings call (or investor-relations activity record
).

## Scenario context ##
Tariffs are taxes imposed on imported foreign goods by the importing country. Tariffs
are NOT export restrictions, quotas, embargoes, financial sanctions, boycotts, or
non-tariff barriers.

Your job: read this transcript and decide, grounded ONLY in what the firm itself says,
whether the firm describes a meaningful impact on its own operations attributable (
directly or indirectly) to tariffs.

<transcript>
{}
</transcript>

```

```

## What counts ##
For both flags below, the impact must be one the firm itself ties to a tariff context.
Do not count export restrictions, financial sanctions, boycotts, or other non-tariff
measures as tariffs. Mere boilerplate that the firm does not connect to its own
operations should NOT trigger either flag.

The two flags are INDEPENDENT and you must judge each on its own evidence      neither
implies the other. A firm may be affected in exactly one of four ways: only
negatively, only positively, both, or neither. Do not assume both; most firms are
affected in only one direction. Set both flags true ONLY when the firm separately
describes a real benefit AND a real harm. If neither side applies, leave both flags
false.

- `positive_impact` = true`      the firm describes a concrete way its operations
  BENEFITED from tariffs.
- `negative_impact` = true`      the firm describes a concrete way its operations were
  HURT by tariffs.

## The two reason fields ##
You will write the positive and negative explanations SEPARATELY. Each field must
contain only its own side      never mention, hedge against, or net out the other
side inside a field. Even when one side mechanically offsets the other, describe
each side's mechanism on its own; do not phrase one field as "X more than offset Y".

`reason_positive_impact` (write only if `positive_impact` is true; otherwise null): a
single prose explanation of the BENEFIT (at most 250 words), grounded only in this
transcript. It must:
- Identify the SPECIFIC products, markets, geographies, or capabilities involved
- Describe what changed for the firm and why, citing the tariff context the firm itself
  names
- Identify the underlying MECHANISM that actually fits this transcript
- Include 1-2 DIRECT QUOTES from the transcript that best illustrate the benefit (use
  quotation marks)
- If the firm cites government support, subsidies, procurement preference, or other
  policy tailwinds, note this explicitly

`reason_negative_impact` (write only if `negative_impact` is true; otherwise null): a
single prose explanation of the HARM (at most 250 words), grounded only in this
transcript, following the same requirements as above but for the negative side, with
1-2 DIRECT QUOTES illustrating the harm.

Do NOT speculate beyond what the transcript says. Do NOT use generic language about "
navigating challenges"      be specific. Do NOT manufacture a benefit to fill `
reason_positive_impact`, or a harm to fill `reason_negative_impact`; if a side's
flag is false, its field is null.

When BOTH flags are false, both reason fields are null and you instead write one short
sentence in `reason_none` stating why.

## Required JSON Output ##
Respond with ONLY a valid JSON object with exactly these fields:
- `"positive_impact"`: boolean
- `"negative_impact"`: boolean
- `"reason_positive_impact"`: string (250 words maximum) when `positive_impact` is true,
  else null

```

```
- `reason_negative_impact`: string (250 words maximum) when `negative_impact` is true,
    else null
- `reason_none`: string (one short sentence) when both flags are false, else null

Example shape:
{{
    "positive_impact": true,
    "negative_impact": false,
    "reason_positive_impact": "The firm's core segment benefited from ... \"...quoted
        phrase...\" ...",
    "reason_negative_impact": null,
    "reason_none": null
}}
```

/no_think

B.6 Reason Analysis Prompt: Step 1 Reason Extraction (Sanctions)

You are an expert analyst studying how firms are affected by sanctions. You are reading the full transcript of a firm's earnings call (or investor-relations activity record).

Scenario context

Sanctions are government-imposed restrictions on trade or financial transactions designed to coerce, punish, or deter targeted firms or governments. Tariffs, which are taxes on imports, are not sanctions.

Your job: read this transcript and decide, grounded ONLY in what the firm itself says, whether the firm describes a meaningful impact on its own operations attributable (directly or indirectly) to sanctions.

<transcript>

{}

</transcript>

What counts

For both flags below, the impact must be one the firm itself ties to a sanctions context. Do not count tariffs (taxes imposed by the importing country) as sanctions. Mere boilerplate that the firm does not connect to its own operations should NOT trigger either flag.

The two flags are INDEPENDENT and you must judge each on its own evidence. Neither implies the other. A firm may be affected in exactly one of four ways: only negatively, only positively, both, or neither. Do not assume both; most firms are affected in only one direction. Set both flags true ONLY when the firm separately describes a real benefit AND a real harm. If neither side applies, leave both flags false.

```
- `positive_impact = true`      the firm describes a concrete way its operations
    BENEFITED from sanctions.
```

```

- `negative_impact` = true`      the firm describes a concrete way its operations were
    HURT by sanctions.

## The two reason fields ##
You will write the positive and negative explanations SEPARATELY. Each field must
    contain only its own side      never mention, hedge against, or net out the other
    side inside a field. Even when one side mechanically offsets the other, describe
    each side's mechanism on its own; do not phrase one field as "X more than offset Y".

`reason_positive_impact` (write only if `positive_impact` is true; otherwise null): a
    single prose explanation of the BENEFIT (at most 250 words), grounded only in this
    transcript. It must:
- Identify the SPECIFIC products, markets, geographies, or capabilities involved
- Describe what changed for the firm and why, citing the sanctions context the firm
    itself names
- Identify the underlying MECHANISM that actually fits this transcript
- Include 1-2 DIRECT QUOTES from the transcript that best illustrate the benefit (use
    quotation marks)
- If the firm cites government support, subsidies, sanctioning-country procurement
    preference, or other policy tailwinds, note this explicitly

`reason_negative_impact` (write only if `negative_impact` is true; otherwise null): a
    single prose explanation of the HARM (at most 250 words), grounded only in this
    transcript, following the same requirements as above but for the negative side, with
    1-2 DIRECT QUOTES illustrating the harm.

Do NOT speculate beyond what the transcript says. Do NOT use generic language about "
    navigating challenges"      be specific. Do NOT manufacture a benefit to fill `
    reason_positive_impact`, or a harm to fill `reason_negative_impact`; if a side's
    flag is false, its field is null.

When BOTH flags are false, both reason fields are null and you instead write one short
    sentence in `reason_none` stating why.

## Required JSON Output ##
Respond with ONLY a valid JSON object with exactly these fields:
- `"positive_impact"`: boolean
- `"negative_impact"`: boolean
- `"reason_positive_impact"`: string (250 words maximum) when `positive_impact` is true,
    else null
- `"reason_negative_impact"`: string (250 words maximum) when `negative_impact` is true,
    else null
- `"reason_none"`: string (one short sentence) when both flags are false, else null

Example shape:
{{
    "positive_impact": true,
    "negative_impact": false,
    "reason_positive_impact": "The firm's core segment benefited from ... \"...quoted
        phrase...\" ...",
    "reason_negative_impact": null,
    "reason_none": null
}}

/no_think

```

B.7 Reason Analysis Prompt: Step 1 Reason Extraction (Export Controls)

```
You are an expert analyst studying how firms are affected by export controls. You are
reading the full transcript of a firm's earnings call (or investor-relations
activity record).

## Scenario context ##
Export controls are restrictions on the sale of certain goods, technologies, or services
. These must be imposed by the exporting country in order to be classified as export
controls. Export controls are not tariffs, financial sanctions, or boycotts.

Your job: read this transcript and decide, grounded ONLY in what the firm itself says,
whether the firm describes a meaningful impact on its own operations attributable (
directly or indirectly) to export controls.

<transcript>
{}
</transcript>

## What counts ##
For both flags below, the impact must be one the firm itself ties to an export-control
context. Do not count tariffs (taxes imposed by the importing country) or generic
financial sanctions as export controls; export controls must be restrictions imposed
by the exporting country. Mere boilerplate that the firm does not connect to its
own operations should NOT trigger either flag.

The two flags are INDEPENDENT and you must judge each on its own evidence      neither
implies the other. A firm may be affected in exactly one of four ways: only
negatively, only positively, both, or neither. Do not assume both; most firms are
affected in only one direction. Set both flags true ONLY when the firm separately
describes a real benefit AND a real harm. If neither side applies, leave both flags
false.

- `positive_impact = true`      the firm describes a concrete way its operations
  BENEFITED from export controls.
- `negative_impact = true`      the firm describes a concrete way its operations were
  HURT by export controls.

## The two reason fields ##
You will write the positive and negative explanations SEPARATELY. Each field must
contain only its own side      never mention, hedge against, or net out the other
side inside a field. Even when one side mechanically offsets the other, describe
each side's mechanism on its own; do not phrase one field as "X more than offset Y".

`reason_positive_impact` (write only if `positive_impact` is true; otherwise null): a
single prose explanation of the BENEFIT (at most 250 words), grounded only in this
transcript. It must:
- Identify the SPECIFIC products, markets, geographies, or capabilities involved
- Describe what changed for the firm and why, citing the export-control context the firm
  itself names
- Identify the underlying MECHANISM that actually fits this transcript
- Include 1-2 DIRECT QUOTES from the transcript that best illustrate the benefit (use
  quotation marks)
```

```

- If the firm cites government support, subsidies, procurement preference, or other
  policy tailwinds, note this explicitly

`reason_negative_impact` (write only if `negative_impact` is true; otherwise null): a
  single prose explanation of the HARM (at most 250 words), grounded only in this
  transcript, following the same requirements as above but for the negative side, with
  1-2 DIRECT QUOTES illustrating the harm.

Do NOT speculate beyond what the transcript says. Do NOT use generic language about "
  navigating challenges"      be specific. Do NOT manufacture a benefit to fill `
  reason_positive_impact`, or a harm to fill `reason_negative_impact`; if a side's
  flag is false, its field is null.

When BOTH flags are false, both reason fields are null and you instead write one short
  sentence in `reason_none` stating why.

## Required JSON Output ##
Respond with ONLY a valid JSON object with exactly these fields:
- `"positive_impact"`: boolean
- `"negative_impact"`: boolean
- `"reason_positive_impact"`: string (250 words maximum) when `positive_impact` is true,
  else null
- `"reason_negative_impact"`: string (250 words maximum) when `negative_impact` is true,
  else null
- `"reason_none"`: string (one short sentence) when both flags are false, else null

Example shape:
{{
  "positive_impact": true,
  "negative_impact": false,
  "reason_positive_impact": "The firm's core segment benefited from ... \"...quoted
    phrase...\" ...",
  "reason_negative_impact": null,
  "reason_none": null
}}

/no_think

```

B.8 Reason Analysis Prompt: Step 2.1 Taxonomy Discovery

```

You are an expert economist analyzing patterns in a batch of firm-quarter observations.
  Each observation is a 150-300 word analyst explanation describing how a specific
  firm was affected by a particular policy / event / scenario.

Below are sampled observations. Read them carefully and let the patterns in the firms'
  own explanations drive the mechanisms you propose. The scenario (what policy / event
  the observations are about) and the direction of impact (e.g. positive, negative)
  are not stated separately      infer them from the observations themselves.

<observations>
{}
</observations>

```

```

## Task
Discover the distinct MECHANISMS through which the underlying policy/event drove the
observed firm-level impacts. Output a flat list      no broad categories or hierarchy
; categories are derived downstream by clustering the consolidated mechanism set.

Requirements:
1. List EVERY meaningfully distinct mechanism you observe in this bundle. There is NO
upper limit and NO target count      if you see 5 distinct channels, list 5; if you
see 30, list 30. A downstream consolidation step merges near-duplicates ACROSS
bundles, so it's much better to over-list distinct channels here than to compress
them. Err on the side of listing more, not fewer.
2. Each mechanism is one concrete channel grounded in the observations (e.g. specific
input cost shock, specific demand shift, specific compliance burden)      not generic
risk-management language.
3. Mechanisms must be MUTUALLY EXCLUSIVE: no two entries describe the same channel. Two
mechanisms count as distinct when they operate via different transmission paths,
target different products / markets / capabilities, or describe different firm
responses      even if the headline impact looks similar. When in doubt between
merging and keeping separate, KEEP THEM SEPARATE.
4. Mechanisms should be COLLECTIVELY EXHAUSTIVE over this bundle      every observation
in the bundle should map cleanly to one of your mechanisms.
5. Order from most specific / distinctive to most generic.
6. Each mechanism needs a short name (3-7 words) suitable for chart legends.
7. Do NOT include an "Unspecified / Other" catch-all in this list      the consolidation
step adds it.

Output format      respond with ONLY valid JSON. Keep descriptions SHORT (under 15 words
each):
{{
  "mechanisms": [
    [{"name": "Short mechanism name", "description": "Under 15 words"}]
  ]
}}

/no_think

```

B.9 Reason Analysis Prompt: Step 2.2 Taxonomy Consolidation

```

You are an expert economist consolidating a draft mechanism list.

Below is a draft mechanism list produced by independently analyzing different bundles of
firm-quarter observations. It contains some near-duplicates (singular/plural
variants, slightly different phrasings of the exact same channel) that you should
merge, AND many genuinely distinct channels that you should PRESERVE so the
downstream taxonomy can distinguish them.

<draft_mechanisms>
$mechanisms_text
</draft_mechanisms>

## Task
Produce a consolidated list of EXACTLY $max_mechanisms mechanisms total: (
  $max_mechanisms      1) substantive entries + the mandatory `Unspecified / Other`

```

last entry. This is a firm target, not an upper bound.

Requirements:

1. Target \$max_mechanisms mechanisms total: (\$max_mechanisms 1) substantive entries plus the mandatory `Unspecified / Other` last entry. \$max_mechanisms is the intended maximum do not exceed it unless the draft genuinely contains more distinct channels that you cannot honestly merge. Over-consolidation destroys downstream signal and makes the clustering step (which builds broader categories at multiple K levels) less informative; under-consolidation pads the taxonomy with synthetic distinctions. Only fall short of \$max_mechanisms when the draft genuinely has fewer than (\$max_mechanisms 1) distinct channels after merging ONLY true near-duplicates never invent or split mechanisms artificially to hit the count.
2. Only merge entries that are GENUINE near-duplicates of each other: same channel in different words (e.g. "Lost Russian sales" and "Loss of Russia revenue"), singular/plural variants, or trivial paraphrases.
3. Keep mechanisms SEPARATE when they describe different transmission paths, different products / markets / capabilities, or different firm responses even when the headline result looks similar. Examples that MUST stay separate:
 - "Lost direct Russian sales" vs "Secondary-sanctions risk constraining counterparties" different paths to revenue loss.
 - "Higher LNG demand" vs "Higher oil-price uplift" both energy upside, distinct channels.
 - "Asset write-down on Russian subsidiary" vs "Payment / FX friction with Russian counterparties" different impacts even within the same exposure.When in doubt between merging and keeping separate, KEEP THEM SEPARATE.

4. MUTUALLY EXCLUSIVE AND NAME-DISTINGUISHABLE: no two entries describe the SAME channel, AND no two names should be so close that a reader could plausibly mistake one for the other. After producing your draft, scan every pair of names if two could refer to the same channel based on name alone, either MERGE them (if truly the same) or RENAME one to disambiguate (if genuinely different). Avoid families of names that differ only by a vague adjective ("Sanctions Impact" vs "Sanctions Effect" vs "Sanctions Disruption" these are interchangeable, not distinctive). (The list does NOT need to be a maximally compact partition granular distinct channels are the goal.)
5. DISTINCTIVE, CONCRETE CHANNELS not generic super-buckets. Both name and description must anchor in something specific drawn from the draft above: a named commodity, market, geography, sector, firm response, transmission path, counterparty type, or product class. Generic risk-management labels are forbidden.

Anti-patterns (DO NOT use names or descriptions like these any mechanism could be tagged this way):

"Cost Effects"	"Demand Effects"	"Supply Chain Disruption"	"Operational Impacts"
"Financial Effects"	"Compliance Issues"	"Market Disruption"	"Risk Exposure"
"Revenue Loss"	"Operating Headwinds"	"Geopolitical Pressure"	"Macro Uncertainty"

If the channel really IS about (say) input costs, name the SPECIFIC input: "Russian Palladium Cost Surge" not "Cost Effects".

Good patterns:

- "Russian Energy Embargo Costs" (commodity + scenario)
- "Defense Contractor Backlog Boost" (sector + direction)
- "Russia-Linked Asset Write-downs" (geography + firm action)
- "Secondary Sanctions Compliance Burden" (transmission path)
- "Western Firms Withdraw From Russia" (action + market)
- "Aerospace Titanium Sourcing Shock" (sector + commodity + dynamic)

```

6. NAME format: 3-7 words, Title Case throughout, consistent punctuation across the list
   . DESCRIPTION format: one line under 15 words, must mention at least one literal
   commodity / market / counterparty / action / sector by name      paraphrasing the
   same generic idea across multiple descriptions is forbidden. Two descriptions should
   never be substitutable; each must say WHAT IS SPECIFICALLY HAPPENING in that
   channel, in literal terms drawn from the draft.
7. Order from most specific / distinctive to most generic; the `Unspecified / Other`
   fallback is always last.
8. The very last entry MUST be {"name": "Unspecified / Other", "description": "Catch-all
   when no listed mechanism applies."}.

## Self-check before responding
- Scan every pair of mechanism names: can you tell them apart on name alone? If not,
  merge or rename.
- Scan every description: does each one name a literal commodity / market / sector /
  action drawn from the draft? If a description could be swapped onto a different
  mechanism without losing meaning, rewrite it to be specific.
- No name in your output appears on the anti-pattern list above (or any close cousin of
  it).

Output format      respond with ONLY valid JSON, keep descriptions SHORT (under 15 words
each):
{
  "mechanisms": [
    {"name": "Short Mechanism Name", "description": "Under 15 words"},
    ...,
    {"name": "Unspecified / Other", "description": "Catch-all when no listed
      mechanism applies."}
  ]
}

/no_think

```

B.10 Reason Analysis Prompt: Step 3.1 Mechanism Classification

```

You are an expert analyst classifying a single firm-quarter observation against a fixed
flat taxonomy of mechanisms.

The case below contains an analyst's reasoning about why a specific firm was affected by
a particular policy / event / scenario. The header lines (REASON / TOPIC / COMPANY
/ SOURCE / PERIOD) name the scenario and the impact direction this case belongs to;
the REASONING section is the analyst's prose.

Classify the PRIMARY mechanism described in that reasoning. There is no hierarchy
pick one mechanism name from the list below.

## CRITICAL: Allowed mechanism list
Each bullet below is one allowed mechanism: a QUOTED name followed by a short
description for context. You MUST pick `primary_mechanism` by copying the QUOTED
NAME VERBATIM (exact spelling, capitalization, punctuation) from one of these
bullets. The description after the em-dash is for disambiguation only      do NOT
include it in your output, do NOT paraphrase the name, do NOT abbreviate, pluralize,
or invent a new label.

```

```
--ENUM_LIST--
```

If none of the listed mechanisms fits, you MUST use `"Unspecified / Other"`. Do NOT invent a new label.

`secondary` entries must also be copied verbatim (quoted-name-only) from the same list (or be `"Unspecified / Other"`); the list may be empty.

```
## Case to classify
{}
```

```
## Rules
```

1. Choose the SINGLE most important mechanism described in the analyst's reasoning. Set `primary\_mechanism` to one mechanism name copied VERBATIM from the list above.
2. Other applicable mechanisms (also verbatim from the list) go in `secondary` as a list of strings. May be empty.
3. Use `"Unspecified / Other"` ONLY when no listed mechanism applies.
4. `confidence`: `"high"` if clear, `"medium"` if implied, `"low"` if ambiguous.

```
## Required JSON Output
```

Respond with a JSON object containing exactly these fields:

- `"primary\_mechanism"`: one mechanism name copied VERBATIM from the allowed list
- `"secondary"`: array of mechanism-name strings, each copied VERBATIM from the allowed list (may be empty)
- `"confidence"`: one of `"high"`, `"medium"`, `"low"``
- `"reasoning"`: brief explanation (1-2 sentences)

```
## Final check before responding
```

Re-read your output and confirm that `primary\_mechanism` is one of the QUOTED names from the allowed mechanism list above (NOT a description, NOT a paraphrase, NOT a combined name+description string). Every entry in `secondary` must also be one of those quoted names. If anything is off, fix it before responding fall back to `"Unspecified / Other"` if no listed mechanism fits.

```
/no_think
```

B.11 Reason Analysis Prompt: Step 3.2 Category Grouping

You are an expert international economist building a mechanism taxonomy for an empirical paper on geoeconomic pressure.

Below are {len(mechs)} fine-grained mechanisms describing why firms report {sidew} {INSTRUMENT_LABEL[instrument]}. Each is a QUOTED name followed by a short description .

```
{lines}
```

```
{k_line}
```

```
Rules:
```

- MECE: assign every mechanism to EXACTLY ONE category.
- Names economically meaningful and figure-legible (<= 4 words). No Other/Unspecified/Misc/General category.

- Group by economic channel, not by geography or severity.
- Every quoted mechanism name above must appear in exactly one category's members.

Output ONLY JSON: `{{"categories": [{{"name": "...", "members": ["<exact mechanism name>", ...]}}]}}`

B.12 Reason Analysis Prompt: Step 1 Reason Extraction (Why Firms Increase R&D)

You are an expert analyst studying how firms shift their R&D investment in response to export controls. You are reading the full transcript of a firm's earnings call (or investor-relations activity record).

Scenario context

Export controls are restrictions on the sale of certain goods, technologies, or services. These must be imposed by the exporting country in order to be classified as export controls. Export controls are not tariffs, financial sanctions, or boycotts.

Your job: read this transcript and decide, grounded ONLY in what the firm itself says, whether it is increasing or expecting to increase the amount of research and development it undertakes as a result of current or future export controls.

`<transcript>`

`{}`

`</transcript>`

What counts

Set ``rd_up = true`` when the firm describes a concrete increase in its own R&D (spend / headcount / programs / intensity / roadmap acceleration) and explicitly attributes that increase to export controls.

Do NOT set ``rd_up = true`` for:

- Generic statements that R&D is "important" or "a priority" without a specific increase.
- R&D increases that the firm attributes to unrelated drivers (new product cycle, AI demand boom, general competitive pressure) with no link to export controls.
- Vague references to "geopolitical headwinds" or "trade tensions" that are not connected to the firm's own R&D decisions.
- Tariff-only narratives: tariffs are a separate instrument. Only count export-control-linked R&D responses. If the firm only discusses tariffs, set the flag false.
- Statements that the firm WILL consider R&D changes if controls tighten further (hypothetical / contingent), unless the firm describes an action already taken.
- R&D CUTS, or R&D held flat, attributed to export controls: those are not what this flag captures.

The two reason fields

The schema below uses a separate field per outcome: the explanation when the flag is true never shares a slot with the reject-sentence when it's false.

```
`reason_rd_up` (write only if `rd_up` is true; otherwise null): a single 150-300 word
  prose explanation, grounded only in this transcript, of why the flag fires. It must:
- Identify the SPECIFIC R&D action taken (e.g. "raised R&D spending from X% to Y% of
  revenue", "added N engineers", "opened a research center in Z", "accelerated next-
  gen GPU roadmap", "launched in-house EDA program").
- Identify the SPECIFIC products, technologies, capabilities, or markets the R&D is
  aimed at (e.g. "advanced-node logic chips", "domestic lithography tools", "AI
  accelerator compliant with the new performance threshold", "alternative to
  restricted EDA software").
- Identify the underlying MECHANISM: substitution for restricted inputs, indigenization
  / self-reliance, designing around performance thresholds, accelerating product
  roadmaps to outpace controls, exploiting competitor weakness created by controls,
  building compliance capability, or other firm-named mechanism.
- Include 1-2 DIRECT QUOTES from the transcript that best illustrate the R&D response (
  use quotation marks).
- Note any government support the firm cites (subsidies, tax credits, "Big Fund", CHIPS
  Act money, regional grants) only if the firm itself connects them to the R&D ramp.
```

Do NOT speculate beyond what the transcript says. Do NOT use generic language about "
 navigating challenges" or "investing for the future" be specific about the R&D
 action, the technology target, and the export-control link. Do NOT manufacture an R&
 D increase to fill `reason\_rd\_up`; if `rd\_up` is false, this field is null.

When `rd\_up` is false, `reason\_rd\_up` is null and you instead write one short sentence
 in `reason\_none` stating why e.g. "Export controls are mentioned only as a macro
 headwind with no firm-specific R&D response described.", or "Firm describes higher
 R&D but attributes it entirely to its AI product cycle, not to export controls.", or
 "Firm discusses tariffs and supply-chain disruption but does not link these to R&D
 investment changes."

Required JSON Output

Respond with ONLY a valid JSON object with exactly these fields:

```
- `rd_up`: boolean
- `reason_rd_up`: string (150-300 words) when `rd_up` is true, else null
- `reason_none`: string (one short sentence) when `rd_up` is false, else null
```

Example shape (rd_up true):

```
{
  "rd_up": true,
  "reason_rd_up": "The firm describes ... \"...quoted phrase from transcript...\"",
  "reason_none": null
}
```

Example shape (rd_up false):

```
{
  "rd_up": false,
  "reason_rd_up": null,
  "reason_none": "Export controls are mentioned only as a macro headwind, with no firm
  -specific R&D response described."
}
```

/no_think

C Inter-Model Agreement Tables

This appendix reports the four-model comparison behind Section 7.2. Each of Gemma 4, Llama 3.3, Qwen 3 and GLM 4.5 Air was run over the same transcripts with the same two-stage prompts, so any two models can be compared answer by answer. Tables A.2 to A.4 then compare our baseline model, Gemma, with Llama field by field on the transcripts that both models flag for the instrument, so they ask whether two models that agree a firm is affected also agree on how. Tables A.5 to A.19 repeat this comparison for the remaining five pairs, ordered by instrument (export controls, sanctions, tariffs). In every table, a , b , c and d count the transcripts on which both models set the flag, only the first does, only the second does, and neither does, and $n = a + b + c + d$. Raw agreement is $p_o = (a + d)/n$. Gwet's AC_1 , our main statistic and shown in bold, is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$, where π is the mean of the two models' positive rates. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$, where p_A and p_B are the two positive rates. Positive agreement is $2a/(2a + b + c)$.

Table A.2: Inter-model agreement on export control effects: Gemma vs. Llama

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	1,611	105	125	100	0.889	0.882	0.852	0.399	0.933
Effect future	313	145	311	1,172	0.279	0.765	0.607	0.421	0.579
Negative impact	1,066	193	192	490	0.648	0.802	0.635	0.565	0.847
Positive impact	338	183	136	1,284	0.256	0.836	0.734	0.569	0.679
Investment up	109	112	170	1,550	0.129	0.855	0.813	0.354	0.436
<i>Investment down</i>	4	19	12	1,906	0.010	0.984	0.984	0.197	0.205
<i>Labor up</i>	2	11	12	1,916	0.007	0.988	0.988	0.142	0.148
<i>Labor down</i>	5	6	9	1,921	0.006	0.992	0.992	0.396	0.400
Sales up	64	154	33	1,690	0.081	0.904	0.887	0.362	0.406
Sales down	271	325	43	1,302	0.234	0.810	0.704	0.487	0.596
Sales price up	23	17	37	1,864	0.026	0.972	0.971	0.446	0.460
<i>Sales price down</i>	1	9	3	1,928	0.004	0.994	0.994	0.140	0.143
<i>Input up</i>	5	31	7	1,898	0.012	0.980	0.980	0.201	0.208
Input down	14	173	27	1,727	0.059	0.897	0.884	0.091	0.123
Input price up	43	14	118	1,766	0.056	0.932	0.924	0.367	0.394
<i>Input price down</i>	4	4	1	1,932	0.003	0.997	0.997	0.614	0.615
Inventory product up	14	54	29	1,844	0.029	0.957	0.955	0.231	0.252
<i>Inventory product down</i>	7	12	14	1,908	0.010	0.987	0.986	0.343	0.350
Inventory input up	23	47	18	1,853	0.029	0.967	0.965	0.398	0.414
<i>Inventory input down</i>	0	3	8	1,930	0.003	0.994	0.994	-0.002	0.000
R&D up	101	115	117	1,608	0.112	0.880	0.851	0.398	0.465
<i>R&D down</i>	0	2	0	1,939	0.001	0.999	0.999	0.000	0.000
Lobbying up	5	68	5	1,863	0.021	0.962	0.961	0.112	0.120
<i>Lobbying down</i>	0	0	0	1,941	0.000	1.000	1.000	–	–
Profit margin up	62	65	55	1,759	0.063	0.938	0.930	0.475	0.508
Profit margin down	212	131	181	1,417	0.190	0.839	0.768	0.477	0.576
Delay	155	129	116	1,541	0.143	0.874	0.833	0.485	0.559
Supply chain adj.	512	160	403	866	0.409	0.710	0.439	0.410	0.645
<i>Financing adj.</i>	1	16	19	1,905	0.010	0.982	0.982	0.045	0.054
Domestic investment up	65	80	93	1,703	0.078	0.911	0.896	0.381	0.429
<i>Domestic investment down</i>	1	6	3	1,931	0.003	0.995	0.995	0.180	0.182
<i>Domestic labor up</i>	1	9	10	1,921	0.005	0.990	0.990	0.090	0.095
<i>Domestic labor down</i>	0	1	7	1,933	0.002	0.996	0.996	-0.001	0.000

Notes: A = Gemma, B = Llama. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by export controls (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.3: Inter-model agreement on sanctions effects: Gemma vs. Llama

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	4,442	182	365	251	0.900	0.896	0.873	0.423	0.942
Effect future	928	498	681	3,133	0.290	0.775	0.618	0.454	0.612
Negative impact	3,381	254	739	866	0.740	0.810	0.692	0.513	0.872
Positive impact	893	574	233	3,540	0.247	0.846	0.755	0.589	0.689
Investment up	133	159	358	4,590	0.075	0.901	0.886	0.290	0.340
Investment down	202	451	162	4,425	0.097	0.883	0.858	0.338	0.397
Labor up	8	41	43	5,148	0.010	0.984	0.984	0.152	0.160
Labor down	32	144	77	4,987	0.027	0.958	0.955	0.204	0.225
Sales up	188	306	150	4,596	0.079	0.913	0.898	0.406	0.452
Sales down	710	1,280	202	3,048	0.277	0.717	0.528	0.329	0.489
Sales price up	62	122	374	4,682	0.059	0.905	0.893	0.158	0.200
<i>Sales price down</i>	0	29	6	5,205	0.003	0.993	0.993	-0.002	0.000
Input up	16	83	49	5,092	0.016	0.975	0.974	0.183	0.195
Input down	90	405	163	4,582	0.071	0.892	0.875	0.189	0.241
Input price up	296	53	805	4,086	0.138	0.836	0.785	0.342	0.408
<i>Input price down</i>	2	30	9	5,199	0.004	0.993	0.992	0.090	0.093
Inventory product up	13	58	84	5,085	0.016	0.973	0.972	0.141	0.155
Inventory product down	7	40	39	5,154	0.009	0.985	0.985	0.143	0.151
Inventory input up	69	61	115	4,995	0.030	0.966	0.964	0.423	0.439
<i>Inventory input down</i>	0	10	19	5,211	0.003	0.994	0.994	-0.003	0.000
R&D up	92	64	153	4,931	0.038	0.959	0.955	0.438	0.459
<i>R&D down</i>	0	19	3	5,218	0.002	0.996	0.996	-0.001	0.000
Lobbying up	1	61	36	5,142	0.009	0.981	0.981	0.011	0.020
<i>Lobbying down</i>	0	0	1	5,239	0.000	1.000	1.000	0.000	0.000
Profit margin up	301	331	161	4,447	0.104	0.906	0.885	0.499	0.550
Profit margin down	970	870	529	2,871	0.319	0.733	0.528	0.388	0.581
Delay	179	224	247	4,590	0.079	0.910	0.895	0.383	0.432
Supply chain adj.	1,390	453	1,370	2,027	0.439	0.652	0.314	0.315	0.604
Financing adj.	167	182	468	4,423	0.094	0.876	0.851	0.277	0.339
Domestic investment up	64	82	161	4,933	0.035	0.954	0.950	0.322	0.345
<i>Domestic investment down</i>	0	10	22	5,208	0.003	0.994	0.994	-0.003	0.000
<i>Domestic labor up</i>	5	9	26	5,200	0.004	0.993	0.993	0.219	0.222
<i>Domestic labor down</i>	0	3	15	5,222	0.002	0.997	0.997	-0.001	0.000

Notes: A = Gemma, B = Llama. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by sanctions (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.4: Inter-model agreement on tariff effects: Gemma vs. Llama

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	17,129	1,697	2,177	3,767	0.770	0.844	0.758	0.559	0.898
Effect future	10,466	3,700	2,205	8,399	0.542	0.762	0.527	0.522	0.780
Negative impact	15,361	2,752	1,982	4,675	0.716	0.809	0.678	0.531	0.866
Positive impact	3,312	2,675	1,479	17,304	0.218	0.832	0.746	0.509	0.615
Investment up	1,337	1,536	1,613	20,284	0.118	0.873	0.840	0.387	0.459
Investment down	53	223	357	24,137	0.014	0.977	0.976	0.143	0.155
Labor up	33	50	188	24,499	0.006	0.990	0.990	0.213	0.217
Labor down	40	73	239	24,418	0.008	0.987	0.987	0.199	0.204
Sales up	509	1,168	517	22,576	0.055	0.932	0.924	0.343	0.377
Sales down	1,687	1,969	957	20,157	0.127	0.882	0.848	0.470	0.536
Sales price up	3,631	1,373	2,280	17,486	0.220	0.853	0.775	0.572	0.665
Sales price down	63	221	66	24,420	0.008	0.988	0.988	0.300	0.305
Input up	39	433	160	24,138	0.014	0.976	0.975	0.106	0.116
Input down	119	1,514	593	22,544	0.047	0.915	0.907	0.064	0.101
Input price up	6,149	1,170	4,513	12,938	0.363	0.771	0.573	0.513	0.684
Input price down	67	249	68	24,386	0.009	0.987	0.987	0.292	0.297
Inventory product up	414	891	336	23,129	0.041	0.950	0.946	0.379	0.403
Inventory product down	87	193	414	24,076	0.016	0.975	0.975	0.211	0.223
Inventory input up	280	411	502	23,577	0.030	0.963	0.961	0.361	0.380
Inventory input down	8	57	260	24,445	0.007	0.987	0.987	0.044	0.048
R&D up	146	162	775	23,687	0.025	0.962	0.960	0.223	0.238
<i>R&D down</i>	0	14	30	24,726	0.001	0.998	0.998	-0.001	0.000
Lobbying up	189	694	254	23,633	0.027	0.962	0.960	0.268	0.285
<i>Lobbying down</i>	0	0	0	24,770	0.000	1.000	1.000	—	—
Profit margin up	711	899	868	22,292	0.064	0.929	0.919	0.408	0.446
Profit margin down	5,790	2,667	2,716	13,597	0.342	0.783	0.605	0.517	0.683
Delay	760	1,181	582	22,247	0.066	0.929	0.919	0.426	0.463
Supply chain adj.	8,793	1,249	4,891	9,837	0.479	0.752	0.505	0.514	0.741
Financing adj.	52	163	496	24,059	0.015	0.973	0.973	0.125	0.136
Domestic investment up	200	220	1,028	23,322	0.033	0.950	0.946	0.223	0.243
Domestic investment down	21	96	190	24,463	0.007	0.988	0.988	0.123	0.128
Domestic labor up	16	15	105	24,634	0.003	0.995	0.995	0.209	0.211
Domestic labor down	32	47	155	24,536	0.005	0.992	0.992	0.237	0.241

Notes: A = Gemma, B = Llama. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by tariffs (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: *a* = both set it, *b* = A only, *c* = B only, *d* = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.5: Inter-model agreement on export control effects: Gemma vs. Qwen

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	2,684	232	168	227	0.871	0.879	0.844	0.463	0.931
Effect future	580	233	198	2,297	0.240	0.870	0.795	0.643	0.729
Negative impact	1,703	444	189	970	0.611	0.809	0.635	0.600	0.843
Positive impact	907	138	572	1,688	0.382	0.785	0.593	0.553	0.719
Investment up	423	45	757	2,085	0.249	0.758	0.613	0.390	0.513
Investment down	25	11	64	3,214	0.019	0.977	0.976	0.391	0.400
Labor up	23	1	784	2,502	0.126	0.763	0.696	0.042	0.055
Labor down	15	2	70	3,224	0.015	0.978	0.978	0.288	0.294
Sales up	400	95	521	2,294	0.214	0.814	0.720	0.460	0.565
Sales down	773	189	219	2,129	0.295	0.877	0.789	0.704	0.791
Sales price up	44	14	97	3,152	0.030	0.966	0.964	0.428	0.442
<i>Sales price down</i>	12	9	27	3,261	0.009	0.989	0.989	0.395	0.400
Input up	36	39	144	3,088	0.039	0.945	0.940	0.259	0.282
Input down	232	176	266	2,636	0.137	0.866	0.825	0.436	0.512
Input price up	71	16	620	2,593	0.118	0.807	0.757	0.142	0.183
<i>Input price down</i>	6	8	15	3,275	0.005	0.993	0.993	0.339	0.343
Inventory product up	69	38	273	2,932	0.068	0.906	0.893	0.271	0.307
Inventory product down	23	8	97	3,181	0.023	0.968	0.967	0.294	0.305
Inventory input up	108	11	593	2,595	0.124	0.817	0.767	0.215	0.263
<i>Inventory input down</i>	2	2	36	3,269	0.006	0.989	0.988	0.093	0.095
R&D up	461	50	679	2,119	0.249	0.780	0.648	0.439	0.558
<i>R&D down</i>	1	4	11	3,297	0.003	0.995	0.995	0.116	0.118
Lobbying up	81	28	69	3,135	0.039	0.971	0.968	0.611	0.625
<i>Lobbying down</i>	0	0	0	3,313	0.000	1.000	1.000	–	–
Profit margin up	171	22	754	2,353	0.169	0.765	0.673	0.232	0.306
Profit margin down	482	83	739	1,999	0.270	0.751	0.589	0.399	0.540
Delay	342	57	442	2,465	0.179	0.849	0.786	0.498	0.578
Supply chain adj.	1,164	212	486	1,445	0.458	0.789	0.581	0.578	0.769
Financing adj.	28	6	311	2,964	0.056	0.904	0.893	0.134	0.150
Domestic investment up	326	26	708	2,249	0.209	0.778	0.668	0.371	0.470
<i>Domestic investment down</i>	5	2	38	3,268	0.008	0.988	0.988	0.197	0.200
Domestic labor up	12	7	697	2,596	0.110	0.787	0.736	0.022	0.033
<i>Domestic labor down</i>	2	0	46	3,265	0.008	0.986	0.986	0.079	0.080

Notes: A = Gemma, B = Qwen. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by export controls (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_{APB} + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.6: Inter-model agreement on export control effects: Gemma vs. GLM

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	2,045	201	93	172	0.873	0.883	0.850	0.474	0.933
Effect future	512	136	352	1,511	0.301	0.806	0.664	0.542	0.677
Negative impact	1,457	326	127	601	0.670	0.820	0.677	0.595	0.865
Positive impact	433	271	30	1,777	0.232	0.880	0.814	0.668	0.742
Investment up	146	168	67	2,130	0.105	0.906	0.885	0.504	0.554
<i>Investment down</i>	8	22	3	2,478	0.008	0.990	0.990	0.386	0.390
<i>Labor up</i>	5	15	3	2,488	0.006	0.993	0.993	0.354	0.357
<i>Labor down</i>	9	7	5	2,490	0.006	0.995	0.995	0.598	0.600
Sales up	15	296	2	2,198	0.065	0.881	0.865	0.080	0.091
Sales down	267	603	14	1,627	0.229	0.754	0.620	0.355	0.464
Sales price up	26	34	4	2,447	0.018	0.985	0.984	0.571	0.578
<i>Sales price down</i>	4	10	1	2,496	0.004	0.996	0.996	0.419	0.421
Input up	21	34	13	2,443	0.018	0.981	0.981	0.463	0.472
Input down	10	259	3	2,239	0.056	0.896	0.883	0.062	0.071
Input price up	59	17	33	2,402	0.033	0.980	0.979	0.692	0.702
<i>Input price down</i>	4	5	0	2,502	0.003	0.998	0.998	0.615	0.615
Inventory product up	24	61	9	2,417	0.023	0.972	0.971	0.395	0.407
<i>Inventory product down</i>	15	10	7	2,479	0.009	0.993	0.993	0.635	0.638
Inventory input up	68	33	18	2,392	0.037	0.980	0.978	0.717	0.727
<i>Inventory input down</i>	0	4	1	2,506	0.001	0.998	0.998	-0.001	0.000
R&D up	156	155	23	2,177	0.098	0.929	0.914	0.601	0.637
<i>R&D down</i>	1	4	1	2,505	0.001	0.998	0.998	0.285	0.286
Lobbying up	25	53	11	2,422	0.023	0.975	0.973	0.427	0.439
<i>Lobbying down</i>	0	0	0	2,511	0.000	1.000	1.000	—	—
Profit margin up	48	109	5	2,349	0.042	0.955	0.951	0.439	0.457
Profit margin down	254	246	83	1,928	0.167	0.869	0.819	0.532	0.607
Delay	161	220	43	2,087	0.116	0.895	0.868	0.497	0.550
Supply chain adj.	774	204	215	1,318	0.392	0.833	0.681	0.650	0.787
<i>Financing adj.</i>	2	20	6	2,483	0.006	0.990	0.990	0.129	0.133
Domestic investment up	109	107	43	2,252	0.073	0.940	0.931	0.561	0.592
<i>Domestic investment down</i>	1	5	2	2,503	0.002	0.997	0.997	0.221	0.222
<i>Domestic labor up</i>	2	15	3	2,491	0.004	0.993	0.993	0.179	0.182
<i>Domestic labor down</i>	0	1	3	2,507	0.001	0.998	0.998	-0.001	0.000

Notes: A = Gemma, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by export controls (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.7: Inter-model agreement on export control effects: Llama vs. Qwen

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	1,542	128	90	85	0.895	0.882	0.854	0.373	0.934
Effect future	304	289	108	1,142	0.273	0.785	0.643	0.463	0.605
Negative impact	933	273	149	485	0.622	0.771	0.567	0.515	0.816
Positive impact	378	85	374	1,001	0.331	0.750	0.552	0.451	0.622
Investment up	209	77	355	1,203	0.230	0.766	0.637	0.360	0.492
Investment down	6	9	49	1,781	0.019	0.969	0.967	0.161	0.171
Labor up	6	5	363	1,470	0.103	0.800	0.755	0.020	0.032
Labor down	7	7	45	1,785	0.018	0.972	0.971	0.203	0.212
Sales up	86	18	364	1,374	0.150	0.793	0.721	0.241	0.310
Sales down	260	48	355	1,181	0.250	0.781	0.650	0.438	0.563
Sales price up	37	20	49	1,739	0.039	0.963	0.960	0.499	0.517
<i>Sales price down</i>	1	1	14	1,829	0.005	0.992	0.992	0.116	0.118
Input up	7	8	84	1,743	0.029	0.950	0.947	0.120	0.132
Input down	19	20	213	1,593	0.073	0.874	0.854	0.108	0.140
Input price up	97	55	234	1,452	0.131	0.843	0.796	0.325	0.402
<i>Input price down</i>	3	1	10	1,828	0.005	0.994	0.994	0.351	0.353
Inventory product up	21	24	145	1,656	0.057	0.908	0.897	0.167	0.199
Inventory product down	11	11	64	1,757	0.026	0.959	0.957	0.212	0.227
Inventory input up	30	11	290	1,512	0.098	0.837	0.802	0.132	0.166
<i>Inventory input down</i>	1	7	26	1,809	0.009	0.982	0.982	0.051	0.057
R&D up	161	61	362	1,261	0.202	0.771	0.662	0.317	0.432
<i>R&D down</i>	0	0	4	1,842	0.001	0.998	0.998	0.000	0.000
Lobbying up	8	2	89	1,747	0.029	0.951	0.948	0.141	0.150
<i>Lobbying down</i>	0	0	0	1,846	0.000	1.000	1.000	–	–
Profit margin up	97	15	365	1,359	0.156	0.793	0.719	0.266	0.338
Profit margin down	311	79	411	1,038	0.302	0.734	0.539	0.392	0.559
Delay	180	84	313	1,262	0.206	0.784	0.679	0.355	0.476
Supply chain adj.	598	291	253	700	0.472	0.705	0.411	0.408	0.687
Financing adj.	10	8	156	1,668	0.050	0.911	0.902	0.093	0.109
Domestic investment up	122	44	350	1,328	0.173	0.786	0.701	0.288	0.382
<i>Domestic investment down</i>	1	3	26	1,816	0.008	0.984	0.984	0.061	0.065
Domestic labor up	4	6	305	1,530	0.086	0.831	0.800	0.015	0.025
<i>Domestic labor down</i>	2	5	25	1,814	0.009	0.984	0.983	0.112	0.118

Notes: A = Llama, B = Qwen. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by export controls (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.8: **Inter-model agreement on export control effects: Llama vs. GLM**

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	1,447	124	57	79	0.901	0.894	0.871	0.410	0.941
Effect future	381	188	176	962	0.330	0.787	0.618	0.518	0.677
Negative impact	954	216	104	433	0.653	0.813	0.657	0.589	0.856
Positive impact	237	184	80	1,206	0.216	0.845	0.766	0.546	0.642
Investment up	88	190	59	1,370	0.124	0.854	0.813	0.340	0.414
<i>Investment down</i>	3	15	11	1,678	0.009	0.985	0.984	0.180	0.188
<i>Labor up</i>	3	13	2	1,689	0.006	0.991	0.991	0.283	0.286
<i>Labor down</i>	5	7	3	1,692	0.006	0.994	0.994	0.497	0.500
Sales up	11	89	4	1,603	0.034	0.946	0.942	0.179	0.191
Sales down	159	162	45	1,341	0.154	0.879	0.836	0.538	0.606
Sales price up	21	37	6	1,643	0.025	0.975	0.974	0.483	0.494
<i>Sales price down</i>	1	1	4	1,701	0.002	0.997	0.997	0.285	0.286
<i>Input up</i>	4	14	22	1,667	0.013	0.979	0.978	0.171	0.182
<i>Input down</i>	4	32	9	1,662	0.014	0.976	0.975	0.154	0.163
Input price up	47	112	14	1,534	0.064	0.926	0.916	0.396	0.427
<i>Input price down</i>	3	0	1	1,703	0.002	0.999	0.999	0.857	0.857
Inventory product up	18	31	15	1,643	0.024	0.973	0.972	0.426	0.439
<i>Inventory product down</i>	8	14	8	1,677	0.011	0.987	0.987	0.415	0.421
Inventory input up	24	15	35	1,633	0.029	0.971	0.969	0.475	0.490
<i>Inventory input down</i>	0	7	2	1,698	0.003	0.995	0.995	-0.002	0.000
R&D up	65	143	44	1,455	0.093	0.890	0.868	0.356	0.410
<i>R&D down</i>	0	0	1	1,706	0.000	0.999	0.999	0.000	0.000
<i>Lobbying up</i>	6	6	25	1,670	0.013	0.982	0.981	0.272	0.279
<i>Lobbying down</i>	0	0	0	1,707	0.000	1.000	1.000	—	—
Profit margin up	28	76	11	1,592	0.042	0.949	0.945	0.371	0.392
Profit margin down	190	204	70	1,243	0.192	0.839	0.767	0.487	0.581
Delay	91	175	58	1,383	0.122	0.864	0.826	0.368	0.439
Supply chain adj.	526	343	143	695	0.450	0.715	0.436	0.433	0.684
<i>Financing adj.</i>	0	19	4	1,684	0.007	0.987	0.986	-0.004	0.000
Domestic investment up	49	112	46	1,500	0.075	0.907	0.893	0.336	0.383
<i>Domestic investment down</i>	0	5	5	1,697	0.003	0.994	0.994	-0.003	0.000
<i>Domestic labor up</i>	2	11	2	1,692	0.005	0.992	0.992	0.233	0.235
<i>Domestic labor down</i>	1	5	2	1,699	0.003	0.996	0.996	0.220	0.222

Notes: A = Llama, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by export controls (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.9: Inter-model agreement on export control effects: Qwen vs. GLM

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	1,957	122	89	211	0.867	0.911	0.885	0.616	0.949
Effect future	481	99	348	1,449	0.296	0.812	0.677	0.555	0.683
Negative impact	1,285	243	216	630	0.638	0.807	0.641	0.582	0.848
Positive impact	436	576	35	1,326	0.312	0.743	0.549	0.435	0.588
Investment up	199	607	16	1,556	0.215	0.738	0.605	0.288	0.390
Investment down	12	63	1	2,305	0.018	0.973	0.972	0.266	0.273
Labor up	8	546	0	1,824	0.118	0.770	0.710	0.022	0.028
Labor down	13	60	1	2,305	0.018	0.974	0.973	0.292	0.299
Sales up	20	608	0	1,749	0.136	0.744	0.665	0.046	0.062
Sales down	259	624	24	1,471	0.245	0.728	0.567	0.322	0.444
Sales price up	24	97	4	2,251	0.031	0.957	0.955	0.309	0.322
<i>Sales price down</i>	5	23	0	2,349	0.007	0.990	0.990	0.301	0.303
Input up	19	120	13	2,224	0.036	0.944	0.940	0.205	0.222
Input down	14	311	1	2,053	0.071	0.869	0.849	0.071	0.082
Input price up	73	393	10	1,896	0.116	0.830	0.786	0.220	0.266
<i>Input price down</i>	3	13	0	2,359	0.004	0.995	0.994	0.314	0.316
Inventory product up	30	216	3	2,131	0.059	0.908	0.897	0.195	0.215
Inventory product down	23	80	4	2,271	0.027	0.965	0.963	0.342	0.354
Inventory input up	81	404	3	1,887	0.120	0.829	0.783	0.239	0.285
<i>Inventory input down</i>	2	32	1	2,342	0.008	0.986	0.986	0.106	0.108
R&D up	165	589	6	1,618	0.194	0.750	0.636	0.271	0.357
<i>R&D down</i>	1	7	1	2,371	0.002	0.997	0.997	0.199	0.200
Lobbying up	34	87	3	2,256	0.033	0.962	0.960	0.416	0.430
<i>Lobbying down</i>	0	0	0	2,380	0.000	1.000	1.000	–	–
Profit margin up	45	576	1	1,750	0.141	0.757	0.679	0.103	0.135
Profit margin down	298	741	47	1,289	0.291	0.668	0.435	0.272	0.431
Delay	177	522	16	1,659	0.188	0.773	0.674	0.309	0.397
Supply chain adj.	821	386	127	1,043	0.453	0.784	0.572	0.570	0.762
Financing adj.	7	250	0	2,120	0.056	0.895	0.883	0.048	0.053
Domestic investment up	139	527	9	1,702	0.171	0.775	0.685	0.267	0.342
<i>Domestic investment down</i>	4	32	1	2,343	0.009	0.986	0.986	0.192	0.195
Domestic labor up	5	449	0	1,925	0.096	0.811	0.771	0.018	0.022
<i>Domestic labor down</i>	4	37	1	2,338	0.010	0.984	0.984	0.171	0.174

Notes: A = Qwen, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by export controls (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.10: **Inter-model agreement on sanctions effects: Gemma vs. Qwen**

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	6,764	436	277	586	0.883	0.912	0.889	0.572	0.950
Effect future	1,233	565	314	5,949	0.207	0.891	0.838	0.669	0.737
Negative impact	4,986	768	330	1,969	0.687	0.864	0.761	0.684	0.901
Positive impact	1,713	294	718	5,329	0.276	0.874	0.791	0.686	0.772
Investment up	329	77	561	7,091	0.080	0.921	0.907	0.471	0.508
Investment down	792	339	287	6,642	0.137	0.922	0.898	0.672	0.717
Labor up	59	16	327	7,656	0.029	0.957	0.955	0.244	0.256
Labor down	235	68	463	7,294	0.062	0.934	0.925	0.440	0.470
Sales up	599	115	643	6,705	0.121	0.906	0.881	0.563	0.612
Sales down	2,900	445	421	4,292	0.414	0.893	0.791	0.778	0.870
Sales price up	184	68	531	7,270	0.060	0.926	0.916	0.351	0.381
Sales price down	25	15	92	7,925	0.010	0.987	0.986	0.313	0.318
Input up	74	73	241	7,675	0.029	0.961	0.959	0.303	0.320
Input down	502	270	624	6,658	0.118	0.889	0.860	0.469	0.529
Input price up	451	63	861	6,672	0.113	0.885	0.856	0.443	0.494
Input price down	33	11	86	7,929	0.010	0.988	0.988	0.400	0.405
Inventory product up	37	40	211	7,771	0.020	0.969	0.968	0.216	0.228
Inventory product down	38	25	356	7,637	0.028	0.953	0.950	0.155	0.166
Inventory input up	148	23	440	7,439	0.047	0.942	0.937	0.369	0.390
Inventory input down	5	9	228	7,812	0.015	0.971	0.970	0.037	0.040
R&D up	165	38	307	7,547	0.042	0.957	0.953	0.470	0.489
R&D down	14	15	22	8,013	0.004	0.995	0.995	0.428	0.431
Lobbying up	43	26	52	7,944	0.010	0.990	0.990	0.520	0.524
<i>Lobbying down</i>	0	0	0	8,065	0.000	1.000	1.000	–	–
Profit margin up	738	106	910	6,308	0.155	0.874	0.829	0.527	0.592
Profit margin down	2,646	340	1,167	3,906	0.422	0.813	0.635	0.621	0.778
Delay	433	114	870	6,638	0.115	0.878	0.847	0.412	0.468
Supply chain adj.	1,974	816	795	4,467	0.345	0.800	0.635	0.557	0.710
Financing adj.	352	87	531	7,087	0.082	0.923	0.910	0.496	0.533
Domestic investment up	187	39	424	7,408	0.052	0.943	0.936	0.423	0.447
Domestic investment down	8	2	63	7,988	0.005	0.992	0.992	0.196	0.198
Domestic labor up	19	4	278	7,755	0.020	0.965	0.964	0.114	0.119
Domestic labor down	3	2	53	8,003	0.004	0.993	0.993	0.097	0.098

Notes: A = Gemma, B = Qwen. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by sanctions (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: *a* = both set it, *b* = A only, *c* = B only, *d* = neither, and *n* = *a* + *b* + *c* + *d*. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.11: Inter-model agreement on sanctions effects: Gemma vs. GLM

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	5,157	178	230	278	0.918	0.930	0.918	0.539	0.962
Effect future	978	412	488	3,965	0.244	0.846	0.756	0.583	0.685
Negative impact	3,813	590	226	1,214	0.722	0.860	0.767	0.653	0.903
Positive impact	887	546	71	4,339	0.205	0.894	0.843	0.679	0.742
Investment up	131	193	42	5,477	0.043	0.960	0.956	0.508	0.527
Investment down	104	812	11	4,916	0.088	0.859	0.832	0.173	0.202
Labor up	13	44	5	5,781	0.006	0.992	0.992	0.344	0.347
Labor down	41	194	16	5,592	0.025	0.964	0.962	0.269	0.281
Sales up	44	478	1	5,320	0.049	0.918	0.910	0.143	0.155
Sales down	623	2,043	27	3,150	0.284	0.646	0.403	0.240	0.376
Sales price up	28	170	21	5,624	0.021	0.967	0.966	0.216	0.227
<i>Sales price down</i>	3	30	1	5,809	0.003	0.995	0.995	0.161	0.162
Input up	28	89	20	5,706	0.014	0.981	0.981	0.332	0.339
Input down	75	540	18	5,210	0.061	0.905	0.892	0.189	0.212
Input price up	259	130	143	5,311	0.068	0.953	0.947	0.630	0.655
<i>Input price down</i>	4	27	1	5,811	0.003	0.995	0.995	0.221	0.222
Inventory product up	13	54	14	5,762	0.008	0.988	0.988	0.272	0.277
Inventory product down	9	48	9	5,777	0.006	0.990	0.990	0.236	0.240
Inventory input up	109	42	60	5,632	0.027	0.983	0.982	0.672	0.681
<i>Inventory input down</i>	3	13	1	5,826	0.002	0.998	0.998	0.299	0.300
R&D up	91	76	20	5,656	0.024	0.984	0.983	0.647	0.655
<i>R&D down</i>	7	19	7	5,810	0.003	0.996	0.996	0.348	0.350
Lobbying up	19	38	9	5,777	0.007	0.992	0.992	0.443	0.447
<i>Lobbying down</i>	0	0	0	5,843	0.000	1.000	1.000	–	–
Profit margin up	259	373	32	5,179	0.079	0.931	0.919	0.529	0.561
Profit margin down	1,185	1,173	185	3,300	0.319	0.768	0.589	0.482	0.636
Delay	194	266	101	5,282	0.065	0.937	0.929	0.482	0.514
Supply chain adj.	1,598	610	564	3,071	0.374	0.799	0.622	0.571	0.731
Financing adj.	162	231	45	5,405	0.051	0.953	0.948	0.518	0.540
Domestic investment up	78	82	33	5,650	0.023	0.980	0.979	0.566	0.576
<i>Domestic investment down</i>	0	8	1	5,834	0.001	0.998	0.998	-0.000	0.000
<i>Domestic labor up</i>	3	11	2	5,827	0.002	0.998	0.998	0.315	0.316
<i>Domestic labor down</i>	1	2	3	5,837	0.001	0.999	0.999	0.285	0.286

Notes: A = Gemma, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by sanctions (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.12: Inter-model agreement on sanctions effects: Llama vs. Qwen

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	4,200	409	143	260	0.893	0.890	0.864	0.428	0.938
Effect future	778	750	356	3,127	0.266	0.779	0.638	0.439	0.585
Negative impact	3,077	899	179	853	0.722	0.785	0.640	0.477	0.851
Positive impact	842	180	789	3,195	0.265	0.806	0.683	0.512	0.635
Investment up	194	237	357	4,219	0.098	0.881	0.856	0.330	0.395
Investment down	196	171	477	4,165	0.104	0.871	0.841	0.312	0.377
Labor up	23	26	206	4,755	0.028	0.954	0.951	0.152	0.165
Labor down	57	56	342	4,556	0.051	0.921	0.912	0.194	0.223
Sales up	228	75	574	4,134	0.110	0.870	0.839	0.356	0.413
Sales down	728	171	1,243	2,867	0.286	0.718	0.522	0.346	0.507
Sales price up	176	242	262	4,328	0.085	0.899	0.881	0.356	0.411
Sales price down	0	6	74	4,929	0.008	0.984	0.984	-0.002	0.000
Input up	27	36	174	4,776	0.026	0.958	0.956	0.189	0.205
Input down	114	138	564	4,193	0.093	0.860	0.831	0.185	0.245
Input price up	526	540	256	3,685	0.185	0.841	0.773	0.475	0.569
Input price down	5	6	74	4,926	0.009	0.984	0.984	0.108	0.111
Inventory product up	25	66	129	4,789	0.024	0.961	0.959	0.185	0.204
Inventory product down	13	29	182	4,783	0.024	0.958	0.956	0.097	0.110
Inventory input up	106	75	251	4,574	0.054	0.935	0.928	0.364	0.394
Inventory input down	2	17	127	4,863	0.015	0.971	0.970	0.021	0.027
R&D up	100	93	198	4,619	0.049	0.942	0.936	0.378	0.407
<i>R&D down</i>	1	3	22	4,987	0.003	0.995	0.995	0.073	0.074
Lobbying up	2	34	62	4,915	0.010	0.981	0.980	0.031	0.040
<i>Lobbying down</i>	0	1	0	5,012	0.000	1.000	1.000	0.000	0.000
Profit margin up	370	71	749	3,821	0.156	0.836	0.778	0.398	0.474
Profit margin down	1,115	366	1,175	2,354	0.376	0.692	0.420	0.362	0.591
Delay	239	173	704	3,896	0.135	0.825	0.772	0.269	0.353
Supply chain adj.	1,275	1,342	388	2,002	0.427	0.654	0.323	0.319	0.596
Financing adj.	235	391	395	3,989	0.125	0.843	0.799	0.285	0.374
Domestic investment up	94	97	275	4,544	0.056	0.926	0.917	0.301	0.336
Domestic investment down	3	24	48	4,936	0.008	0.986	0.985	0.070	0.077
Domestic labor up	18	11	160	4,818	0.021	0.966	0.964	0.166	0.174
Domestic labor down	3	15	33	4,960	0.005	0.990	0.990	0.107	0.111

Notes: A = Llama, B = Qwen. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by sanctions (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: *a* = both set it, *b* = A only, *c* = B only, *d* = neither, and *n* = *a* + *b* + *c* + *d*. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.13: **Inter-model agreement on sanctions effects: Llama vs. GLM**

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	3,946	210	123	155	0.927	0.925	0.913	0.442	0.960
Effect future	779	541	459	2,655	0.288	0.774	0.617	0.451	0.609
Negative impact	2,837	743	138	716	0.739	0.801	0.677	0.497	0.866
Positive impact	610	306	206	3,312	0.195	0.885	0.832	0.633	0.704
Investment up	81	329	48	3,976	0.061	0.915	0.904	0.268	0.301
Investment down	57	288	36	4,053	0.049	0.927	0.919	0.235	0.260
Labor up	7	42	9	4,376	0.007	0.988	0.988	0.211	0.215
Labor down	27	76	25	4,306	0.017	0.977	0.976	0.338	0.348
Sales up	27	258	9	4,140	0.036	0.940	0.935	0.156	0.168
Sales down	342	502	164	3,426	0.152	0.850	0.798	0.425	0.507
Sales price up	25	367	11	4,031	0.048	0.915	0.906	0.103	0.117
<i>Sales price down</i>	0	6	3	4,425	0.001	0.998	0.998	-0.001	0.000
Input up	13	45	26	4,350	0.011	0.984	0.984	0.260	0.268
Input down	36	204	40	4,154	0.036	0.945	0.941	0.207	0.228
Input price up	262	698	42	3,432	0.143	0.833	0.779	0.346	0.415
<i>Input price down</i>	3	5	2	4,424	0.001	0.998	0.998	0.461	0.462
Inventory product up	12	76	15	4,331	0.013	0.979	0.979	0.201	0.209
Inventory product down	3	37	14	4,380	0.006	0.988	0.988	0.100	0.105
Inventory input up	77	94	61	4,202	0.035	0.965	0.963	0.480	0.498
<i>Inventory input down</i>	0	18	2	4,414	0.002	0.995	0.995	-0.001	0.000
R&D up	65	123	18	4,228	0.031	0.968	0.966	0.466	0.480
<i>R&D down</i>	1	3	11	4,419	0.002	0.997	0.997	0.124	0.125
Lobbying up	1	30	20	4,383	0.006	0.989	0.989	0.033	0.038
<i>Lobbying down</i>	0	1	0	4,433	0.000	1.000	1.000	0.000	0.000
Profit margin up	176	224	87	3,947	0.075	0.930	0.919	0.495	0.531
Profit margin down	683	691	406	2,654	0.278	0.753	0.587	0.386	0.555
Delay	128	253	126	3,927	0.072	0.915	0.901	0.359	0.403
Supply chain adj.	1,257	1,173	363	1,641	0.457	0.654	0.312	0.325	0.621
Financing adj.	121	456	69	3,788	0.086	0.882	0.859	0.268	0.316
Domestic investment up	42	134	33	4,225	0.028	0.962	0.960	0.318	0.335
<i>Domestic investment down</i>	1	21	1	4,411	0.003	0.995	0.995	0.083	0.083
<i>Domestic labor up</i>	4	26	3	4,401	0.004	0.993	0.993	0.214	0.216
<i>Domestic labor down</i>	1	13	2	4,418	0.002	0.997	0.997	0.117	0.118

Notes: A = Llama, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by sanctions (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: *a* = both set it, *b* = A only, *c* = B only, *d* = neither, and *n* = *a* + *b* + *c* + *d*. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.14: Inter-model agreement on sanctions effects: Qwen vs. GLM

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	4,812	109	288	304	0.909	0.928	0.914	0.567	0.960
Effect future	793	296	583	3,842	0.224	0.841	0.756	0.543	0.643
Negative impact	3,422	428	403	1,258	0.696	0.849	0.739	0.643	0.892
Positive impact	823	827	65	3,794	0.230	0.838	0.749	0.555	0.649
Investment up	121	489	21	4,876	0.068	0.907	0.894	0.292	0.322
Investment down	102	747	12	4,650	0.087	0.862	0.836	0.182	0.212
Labor up	19	242	0	5,250	0.025	0.956	0.954	0.130	0.136
Labor down	49	471	8	4,984	0.052	0.913	0.904	0.154	0.170
Sales up	41	792	2	4,679	0.079	0.856	0.831	0.080	0.094
Sales down	588	1,927	32	2,964	0.284	0.645	0.400	0.237	0.375
Sales price up	41	455	5	5,007	0.049	0.916	0.908	0.138	0.151
Sales price down	4	75	1	5,430	0.008	0.986	0.986	0.094	0.095
Input up	28	188	17	5,282	0.024	0.963	0.961	0.204	0.215
Input down	71	765	17	4,654	0.084	0.858	0.832	0.128	0.154
Input price up	333	557	48	4,569	0.115	0.890	0.862	0.473	0.524
Input price down	3	74	1	5,434	0.007	0.986	0.986	0.073	0.074
Inventory product up	20	164	5	5,322	0.019	0.969	0.968	0.185	0.191
Inventory product down	14	259	2	5,235	0.026	0.953	0.950	0.092	0.097
Inventory input up	148	279	19	5,060	0.054	0.946	0.940	0.475	0.498
Inventory input down	1	165	2	5,341	0.015	0.970	0.969	0.011	0.012
R&D up	70	245	13	5,184	0.036	0.953	0.950	0.336	0.352
<i>R&D down</i>	9	20	6	5,480	0.004	0.995	0.995	0.407	0.409
Lobbying up	16	57	8	5,434	0.009	0.988	0.988	0.325	0.330
<i>Lobbying down</i>	0	0	0	5,515	0.000	1.000	1.000	–	–
Profit margin up	268	867	10	4,367	0.128	0.841	0.795	0.325	0.379
Profit margin down	1,214	1,605	123	2,571	0.377	0.687	0.409	0.380	0.584
Delay	236	797	45	4,431	0.119	0.847	0.807	0.303	0.359
Supply chain adj.	1,389	589	621	2,910	0.362	0.780	0.592	0.524	0.697
Financing adj.	176	564	25	4,746	0.085	0.893	0.873	0.336	0.374
Domestic investment up	76	319	10	5,106	0.044	0.940	0.935	0.298	0.316
Domestic investment down	1	53	1	5,457	0.005	0.990	0.990	0.035	0.036
Domestic labor up	6	184	0	5,315	0.018	0.967	0.965	0.059	0.061
<i>Domestic labor down</i>	4	38	0	5,470	0.004	0.993	0.993	0.173	0.174

Notes: A = Qwen, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by sanctions (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.15: Inter-model agreement on tariff effects: Gemma vs. Qwen

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	19,842	3,348	1,873	7,428	0.691	0.839	0.720	0.625	0.884
Effect future	14,072	3,672	1,692	13,002	0.516	0.835	0.670	0.670	0.840
Negative impact	17,535	4,567	1,689	8,648	0.637	0.807	0.641	0.587	0.849
Positive impact	5,814	1,691	2,328	22,604	0.241	0.876	0.805	0.662	0.743
Investment up	2,568	842	1,103	27,960	0.109	0.940	0.926	0.692	0.725
Investment down	198	136	370	31,785	0.014	0.984	0.984	0.432	0.439
Labor up	61	33	897	31,475	0.016	0.971	0.970	0.111	0.116
Labor down	87	40	291	32,065	0.008	0.990	0.990	0.341	0.345
Sales up	1,282	736	1,339	29,072	0.072	0.936	0.926	0.519	0.553
Sales down	2,966	1,370	1,383	26,739	0.134	0.915	0.890	0.634	0.683
Sales price up	4,206	1,031	819	26,306	0.159	0.943	0.922	0.786	0.820
Sales price down	193	118	312	31,819	0.013	0.987	0.986	0.467	0.473
Input up	185	324	398	31,562	0.017	0.978	0.977	0.328	0.339
Input down	721	1,055	432	30,264	0.045	0.954	0.950	0.469	0.492
Input price up	7,185	804	2,193	22,159	0.268	0.907	0.847	0.765	0.827
Input price down	284	125	308	31,702	0.015	0.987	0.986	0.561	0.567
Inventory product up	612	764	333	30,775	0.036	0.966	0.964	0.510	0.527
Inventory product down	154	150	320	31,865	0.012	0.986	0.985	0.389	0.396
Inventory input up	619	125	1,041	30,691	0.037	0.964	0.961	0.499	0.515
Inventory input down	21	51	184	32,228	0.004	0.993	0.993	0.149	0.152
R&D up	183	172	266	31,848	0.012	0.987	0.986	0.448	0.455
<i>R&D down</i>	12	4	20	32,468	0.001	0.999	0.999	0.500	0.500
Lobbying up	799	339	279	31,087	0.034	0.981	0.980	0.711	0.721
<i>Lobbying down</i>	0	0	0	32,510	0.000	1.000	1.000	–	–
Profit margin up	1,516	434	2,396	28,103	0.090	0.913	0.896	0.475	0.517
Profit margin down	8,139	1,425	4,154	18,702	0.337	0.828	0.689	0.618	0.745
Delay	1,699	512	1,901	28,356	0.089	0.926	0.911	0.546	0.585
Supply chain adj.	8,916	2,498	1,437	19,588	0.336	0.879	0.781	0.728	0.819
Financing adj.	131	111	358	31,877	0.011	0.986	0.985	0.352	0.358
Domestic investment up	367	137	845	31,059	0.026	0.970	0.968	0.415	0.428
Domestic investment down	91	46	395	31,879	0.010	0.986	0.986	0.287	0.292
Domestic labor up	19	13	484	31,933	0.008	0.985	0.984	0.069	0.071
Domestic labor down	60	27	234	32,143	0.006	0.992	0.992	0.312	0.315

Notes: A = Gemma, B = Qwen. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by tariffs (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.16: Inter-model agreement on tariff effects: Gemma vs. GLM

Flag	Counts				π	Raw	Agreement		
	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	18,717	2,776	1,525	5,421	0.734	0.849	0.752	0.614	0.897
Effect future	12,647	3,171	1,689	10,932	0.530	0.829	0.659	0.658	0.839
Negative impact	16,652	3,876	1,100	6,811	0.673	0.825	0.687	0.607	0.870
Positive impact	4,367	2,515	716	20,840	0.210	0.886	0.830	0.660	0.730
Investment up	1,790	1,673	391	24,585	0.099	0.927	0.912	0.596	0.634
Investment down	113	192	114	28,020	0.009	0.989	0.989	0.419	0.425
Labor up	27	65	23	28,324	0.002	0.997	0.997	0.379	0.380
Labor down	55	65	57	28,262	0.004	0.996	0.996	0.472	0.474
Sales up	194	1,749	20	26,476	0.038	0.938	0.933	0.169	0.180
Sales down	1,451	2,753	171	24,064	0.102	0.897	0.874	0.453	0.498
Sales price up	3,100	2,078	269	22,992	0.150	0.917	0.889	0.679	0.725
Sales price down	114	188	45	28,092	0.008	0.992	0.992	0.491	0.495
Input up	202	306	185	27,746	0.016	0.983	0.982	0.443	0.451
Input down	139	1,663	37	26,600	0.035	0.940	0.936	0.131	0.141
Input price up	6,373	1,379	1,307	19,380	0.271	0.906	0.844	0.761	0.826
Input price down	172	197	74	27,996	0.011	0.990	0.990	0.555	0.559
Inventory product up	357	1,016	75	26,991	0.032	0.962	0.959	0.381	0.396
Inventory product down	93	204	113	28,029	0.009	0.989	0.989	0.364	0.370
Inventory input up	485	260	580	27,114	0.032	0.970	0.969	0.521	0.536
Inventory input down	14	55	33	28,337	0.002	0.997	0.997	0.240	0.241
R&D up	146	231	75	27,987	0.011	0.989	0.989	0.483	0.488
<i>R&D down</i>	3	12	13	28,411	0.001	0.999	0.999	0.193	0.194
Lobbying up	520	519	110	27,290	0.029	0.978	0.977	0.612	0.623
<i>Lobbying down</i>	0	0	0	28,439	0.000	1.000	1.000	–	–
Profit margin up	863	991	346	26,239	0.054	0.953	0.948	0.540	0.563
Profit margin down	6,629	2,644	1,432	17,734	0.305	0.857	0.751	0.663	0.765
Delay	1,175	924	318	26,022	0.063	0.956	0.950	0.632	0.654
Supply chain adj.	9,596	1,730	1,933	15,180	0.402	0.871	0.752	0.732	0.840
Financing adj.	64	170	31	28,174	0.006	0.993	0.993	0.386	0.389
Domestic investment up	259	251	267	27,662	0.018	0.982	0.981	0.491	0.500
Domestic investment down	44	79	38	28,278	0.004	0.996	0.996	0.427	0.429
<i>Domestic labor up</i>	11	22	16	28,390	0.001	0.999	0.999	0.366	0.367
Domestic labor down	39	43	28	28,329	0.003	0.998	0.997	0.522	0.523

Notes: A = Gemma, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by tariffs (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.17: **Inter-model agreement on tariff effects: Llama vs. Qwen**

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	15,913	2,924	1,223	4,092	0.745	0.828	0.723	0.551	0.885
Effect future	9,367	3,008	2,565	9,166	0.504	0.769	0.538	0.538	0.771
Negative impact	13,623	3,346	2,081	5,060	0.678	0.775	0.600	0.486	0.834
Positive impact	3,277	1,345	3,132	16,359	0.229	0.814	0.713	0.478	0.594
Investment up	1,326	1,516	1,672	19,627	0.121	0.868	0.832	0.379	0.454
Investment down	70	335	412	23,338	0.018	0.969	0.968	0.142	0.158
Labor up	75	143	712	23,205	0.021	0.965	0.963	0.137	0.149
Labor down	81	192	248	23,624	0.012	0.982	0.981	0.260	0.269
Sales up	541	442	1,560	21,556	0.064	0.917	0.906	0.313	0.351
Sales down	1,606	996	2,006	19,513	0.129	0.876	0.840	0.448	0.517
Sales price up	3,608	2,207	1,091	17,137	0.219	0.863	0.792	0.600	0.686
Sales price down	73	59	365	23,609	0.012	0.982	0.982	0.250	0.256
Input up	54	135	432	23,518	0.014	0.977	0.976	0.150	0.160
Input down	130	566	899	22,540	0.036	0.939	0.935	0.120	0.151
Input price up	6,560	3,936	1,662	11,884	0.389	0.767	0.556	0.515	0.701
Input price down	74	54	356	23,610	0.012	0.983	0.983	0.259	0.265
Inventory product up	284	455	597	22,812	0.034	0.956	0.953	0.328	0.351
Inventory product down	124	376	293	23,357	0.019	0.972	0.971	0.256	0.270
Inventory input up	490	271	1,011	22,364	0.047	0.947	0.942	0.408	0.433
Inventory input down	24	242	147	23,730	0.009	0.984	0.984	0.102	0.110
R&D up	112	770	241	23,015	0.026	0.958	0.956	0.164	0.181
R&D down	2	28	26	24,105	0.001	0.998	0.998	0.068	0.069
Lobbying up	205	238	670	23,048	0.027	0.962	0.960	0.294	0.311
<i>Lobbying down</i>	0	0	0	24,165	0.000	1.000	1.000	–	–
Profit margin up	1,005	525	2,087	20,504	0.096	0.892	0.869	0.382	0.435
Profit margin down	6,313	2,091	4,149	11,542	0.391	0.741	0.505	0.461	0.669
Delay	780	539	2,321	20,491	0.092	0.881	0.858	0.299	0.353
Supply chain adj.	7,869	5,449	1,009	9,785	0.460	0.732	0.468	0.479	0.709
Financing adj.	73	456	357	23,251	0.020	0.966	0.965	0.135	0.152
Domestic investment up	339	863	687	22,202	0.046	0.936	0.929	0.271	0.304
Domestic investment down	37	174	353	23,518	0.012	0.978	0.978	0.113	0.123
Domestic labor up	44	76	379	23,624	0.011	0.981	0.981	0.156	0.162
Domestic labor down	60	126	193	23,754	0.009	0.987	0.987	0.267	0.273

Notes: A = Llama, B = Qwen. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by tariffs (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: *a* = both set it, *b* = A only, *c* = B only, *d* = neither, and *n* = *a* + *b* + *c* + *d*. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.18: Inter-model agreement on tariff effects: Llama vs. GLM

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	16,082	2,473	1,136	3,681	0.765	0.846	0.759	0.572	0.899
Effect future	9,639	2,367	2,526	8,840	0.517	0.791	0.582	0.581	0.798
Negative impact	13,717	2,909	1,558	5,188	0.682	0.809	0.663	0.561	0.860
Positive impact	2,856	1,671	1,326	17,518	0.186	0.872	0.816	0.577	0.656
Investment up	1,075	1,807	775	19,715	0.101	0.890	0.865	0.396	0.454
Investment down	58	347	154	22,813	0.013	0.979	0.978	0.178	0.188
Labor up	26	189	15	23,142	0.005	0.991	0.991	0.201	0.203
Labor down	67	212	45	23,048	0.008	0.989	0.989	0.338	0.343
Sales up	137	857	60	22,318	0.025	0.961	0.959	0.219	0.230
Sales down	1,037	1,556	438	20,341	0.087	0.915	0.899	0.467	0.510
Sales price up	2,847	2,917	415	17,193	0.193	0.857	0.793	0.551	0.631
Sales price down	54	77	94	23,147	0.006	0.993	0.993	0.383	0.387
Input up	43	151	317	22,861	0.012	0.980	0.979	0.146	0.155
Input down	48	645	119	22,560	0.018	0.967	0.966	0.101	0.112
Input price up	6,014	4,372	1,076	11,910	0.374	0.767	0.562	0.512	0.688
Input price down	59	68	128	23,117	0.007	0.992	0.992	0.372	0.376
Inventory product up	218	518	200	22,436	0.025	0.969	0.968	0.363	0.378
Inventory product down	92	404	102	22,774	0.015	0.978	0.978	0.258	0.267
Inventory input up	455	308	542	22,067	0.038	0.964	0.961	0.498	0.517
Inventory input down	9	253	36	23,074	0.007	0.988	0.987	0.056	0.059
R&D up	106	799	81	22,386	0.023	0.962	0.961	0.183	0.194
<i>R&D down</i>	0	29	15	23,328	0.001	0.998	0.998	-0.001	0.000
Lobbying up	171	252	385	22,564	0.021	0.973	0.972	0.336	0.349
<i>Lobbying down</i>	0	0	0	23,372	0.000	1.000	1.000	—	—
Profit margin up	600	917	420	21,435	0.054	0.943	0.936	0.444	0.473
Profit margin down	5,500	2,833	1,897	13,142	0.337	0.798	0.634	0.548	0.699
Delay	657	645	717	21,353	0.057	0.942	0.935	0.460	0.491
Supply chain adj.	8,883	4,442	1,200	8,847	0.501	0.759	0.517	0.526	0.759
Financing adj.	43	486	48	22,795	0.013	0.977	0.977	0.133	0.139
Domestic investment up	234	975	227	21,936	0.036	0.949	0.945	0.259	0.280
Domestic investment down	19	192	58	23,103	0.006	0.989	0.989	0.128	0.132
Domestic labor up	14	107	10	23,241	0.003	0.995	0.995	0.192	0.193
Domestic labor down	41	146	28	23,157	0.005	0.993	0.992	0.317	0.320

Notes: A = Llama, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by tariffs (**effect_any** = 1 for each). **effect_any** itself is omitted, being 1 by construction. Each row is one flag: *a* = both set it, *b* = A only, *c* = B only, *d* = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

Table A.19: Inter-model agreement on tariff effects: Qwen vs. GLM

Flag	Counts				π	Raw	Agreement		
	a	b	c	d			AC ₁	κ	Pos.
	both 1	A only	B only	both 0					
Effect current	17,480	1,658	2,021	6,234	0.705	0.866	0.770	0.677	0.905
Effect future	11,296	2,096	2,629	11,325	0.499	0.827	0.654	0.655	0.827
Negative impact	14,513	2,933	2,675	7,220	0.633	0.795	0.617	0.558	0.838
Positive impact	4,077	3,100	751	19,412	0.220	0.859	0.786	0.593	0.679
Investment up	1,599	1,851	420	23,508	0.100	0.917	0.899	0.542	0.585
Investment down	149	377	73	26,794	0.014	0.984	0.983	0.391	0.398
Labor up	41	849	6	26,479	0.017	0.969	0.968	0.085	0.088
Labor down	87	269	27	27,005	0.009	0.989	0.989	0.366	0.370
Sales up	187	2,205	25	24,919	0.048	0.918	0.910	0.131	0.144
Sales down	1,297	2,731	294	23,042	0.103	0.889	0.864	0.413	0.462
Sales price up	3,034	1,815	284	22,138	0.150	0.923	0.897	0.700	0.743
Sales price down	126	338	31	26,853	0.011	0.987	0.986	0.401	0.406
Input up	159	380	219	26,619	0.017	0.978	0.977	0.336	0.347
Input down	142	973	29	26,235	0.023	0.963	0.962	0.212	0.221
Input price up	6,647	2,126	876	17,603	0.299	0.890	0.810	0.738	0.816
Input price down	176	310	59	26,784	0.013	0.986	0.986	0.482	0.488
Inventory product up	296	628	129	26,338	0.025	0.972	0.971	0.427	0.439
Inventory product down	120	340	87	26,849	0.012	0.984	0.984	0.353	0.360
Inventory input up	852	747	181	25,602	0.048	0.966	0.963	0.630	0.647
Inventory input down	32	164	14	27,181	0.004	0.994	0.993	0.262	0.264
R&D up	107	314	93	26,862	0.011	0.985	0.985	0.338	0.345
<i>R&D down</i>	9	21	7	27,372	0.001	0.999	0.999	0.391	0.391
Lobbying up	534	454	97	26,323	0.030	0.980	0.979	0.650	0.660
<i>Lobbying down</i>	0	0	0	27,414	0.000	1.000	1.000	–	–
Profit margin up	943	2,544	219	23,650	0.085	0.899	0.880	0.365	0.406
Profit margin down	6,788	4,630	1,086	14,821	0.353	0.791	0.615	0.550	0.704
Delay	1,187	2,156	277	23,758	0.088	0.911	0.894	0.453	0.494
Supply chain adj.	8,460	1,343	2,569	14,980	0.381	0.857	0.729	0.697	0.812
Financing adj.	68	389	27	26,897	0.010	0.985	0.984	0.242	0.246
Domestic investment up	312	825	177	26,002	0.030	0.963	0.961	0.368	0.384
Domestic investment down	64	386	19	26,849	0.010	0.985	0.985	0.236	0.240
Domestic labor up	21	448	4	26,884	0.009	0.983	0.983	0.083	0.085
Domestic labor down	57	224	13	27,077	0.006	0.991	0.991	0.322	0.325

Notes: A = Qwen, B = GLM. Both models answer the same second-stage questions about the same transcript. The sample is the transcripts where both models say the firm is affected by tariffs (`effect_any` = 1 for each). `effect_any` itself is omitted, being 1 by construction. Each row is one flag: a = both set it, b = A only, c = B only, d = neither, and $n = a + b + c + d$. $p_A = (a + b)/n$ and $p_B = (a + c)/n$ are the two models' positive rates and $\pi = (p_A + p_B)/2$ their mean. Raw agreement is $p_o = (a + d)/n$. Gwet's AC₁ is $(p_o - p_e)/(1 - p_e)$ with $p_e = 2\pi(1 - \pi)$. Cohen's κ is the same expression with $p_e = p_A p_B + (1 - p_A)(1 - p_B)$. Positive agreement (Pos.) is $2a/(2a + b + c)$. AC₁ is in bold. Rows in grey rest on fewer than 50 transcripts where either model set the flag.

D Additional Results


This appendix presents additional results.

D.1 The Supply Chain Response of Russian Firms

In the main text, we focus on the reaction of firms outside of Russia to the sanctions imposed on Russia. We next turn to exploring how Russian firms respond to sanctions in the wake of Russia's invasion of Ukraine in 2022. 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 analyst reports. While these do not directly cover individual firms, they do cover 16 sectors of the Russian economy at a quarterly frequency. Figure A.1 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.¹

¹See [Egorov et al. \(2025\)](#) for a detailed study of the responses of Russian firms to trade sanctions.

Figure A.1: 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.