The Informational Value of Lobbying in the Tariff Exclusion Process*

Yan Yan
Yale University

November 2025

Latest Version Available Here

Abstract

This paper studies the informational role of lobbying in the U.S.–China trade war. Under Section 301 of the Trade Act of 1974, the United States imposed tariffs on more than $450 billion of Chinese imports and allowed firms to request product-specific exclusions. I develop a signaling model in which firms lobby to transmit information about the welfare consequences of exclusions to a politically motivated regulator who sets both the tariff and the approval rule. Calibrated to 50,000 exclusion requests linked to corporate-group lobbying disclosures, the model shows that informational lobbying moderated the Section 301 tariff and improved the targeting of tariff relief. As a result, total welfare losses fell by roughly 18 percent relative to a world without lobbying, even though the policymaker maximized a political rather than social objective. The framework provides a tractable approach for evaluating how bureaucratic discretion and private information interact in modern trade and industrial policy.

*I am deeply grateful to Giovanni Maggi, Ana Cecilia Fieler, and Guillermo Noguera for their guidance, encouragement, and many conversations that shaped this project. I thank Jingting Fan, Teresa Fort, Kalina Manova, Alejandro Sabal, and Anton Yang for helpful comments and conversations. I thank participants at the Yale International and Spatial Economics Lunch and the International Trade Workshop for valuable feedback. All errors are my own.

1 Introduction

In 2018, the United States imposed a sweeping set of tariffs on imports from China under Section 301 of the Trade Act of 1974, citing unfair trade practices and intellectual-property violations. These measures eventually covered more than $450 billion in imports across four successive product lists. As documented by Bown and Kolb (2021), the tariffs disproportionately targeted intermediate inputs rather than final consumer goods, with rates ranging from 7.5 to 25 percent depending on the item. As part of the implementation, the Office of the U.S. Trade Representative (USTR) established a formal process through which firms could request product-specific exclusions from the new duties. Between 2018 and 2020, more than 50,000 petitions were submitted, each requiring the applicant to provide detailed information on sourcing patterns, domestic alternatives, and the expected economic harm from the tariff. The Government Accountability Office later documented the extraordinary informational burden this system placed on a small bureaucracy operating under tight deadlines and limited staff capacity. Reviewers were required to assess tens of thousands of product-level claims, often relying entirely on information self-reported by petitioners.

This setting, in which a politically directed policy must be implemented by a bureaucracy operating with limited information, raises a broader question of how lobbying affects policy through the transmission of information. The objective of this paper is to quantify the informational role of lobbying in the Section 301 tariff-exclusion process and to evaluate how this informational channel interacts with the political motives embedded in the government's objective, as well as the resulting implications for welfare. In the Section 301 context, these political motives are best interpreted as reduced-form weights that capture the distributional and strategic considerations surrounding the U.S.–China trade action, such as the importance assigned to domestic upstream producers, downstream users of Chinese inputs, tariff revenue, and exposure to Chinese suppliers. I show that lobbying improved the targeting of tariff relief, moderated the tariff set ex ante, and reduced the total welfare loss from the trade war by roughly 18 percent, even though the policymaker maximized a political rather than a welfare objective.

The Section 301 exclusion process offers a rare empirical window into the mechanics of informational lobbying. The introduction of large and heterogeneous tariffs triggered substantial lobbying activity. As documented by Chor et al. (2025), approval outcomes were systematically related to economic fundamentals such as reliance on Chinese suppliers. Building on this evidence, incorporating firm-level lobbying disclosures shows that lobbying did not guarantee success: many lobbying firms saw petitions denied, while some non-lobbying firms still received approvals. Taken together, these empirical patterns indicate that approval

decisions reflected both publicly observable economic fundamentals and additional information supplied by firms. This perspective is consistent with the idea that lobbying in this setting operated as a costly communication channel through which firms conveyed credible private information about the economic consequences of tariffs, and that the USTR used this information, imperfectly but meaningfully, when deciding which products to exempt.

The institutional details reinforce this interpretation. The exclusion process unfolded within a rule-governed bureaucracy rather than a legislative or campaign-finance setting, and decisions were made within established administrative procedures. Firms interacted with the agency primarily by supplying technical and product-specific information, which case reviewers relied on heavily given tight time and staffing constraints. The USTR coordinated with multiple agencies (including U.S. Customs and the International Trade Commission) through a quasi-adjudicatory review process that emphasized documentation and administrability. In this environment, statutory criteria, administrative discretion, and limited information combined to place a premium on credible communication, making the setting naturally suited to an informational role for lobbying.

Motivated by these patterns, the paper develops a signaling model in which lobbying is interpreted as the costly transmission of private information about the economic consequences of tariff relief under a government objective that places reduced-form political weights on different components of welfare. These political weights summarize broader distributional and strategic considerations surrounding the Section 301 action, rather than being shaped by the firm–bureaucracy lobbying studied here. Firms privately observe the product-specific economic impact of an exclusion, information they know more precisely than the regulator, and decide whether to incur a fixed cost to lobby and thereby generate a noisy signal that the government interprets through Bayesian updating. The policy process unfolds in two stages: the government first chooses the Section 301 tariff while anticipating how firms will behave in the subsequent exclusion stage, and then, given realized lobbying decisions and signals, decides which petitions to approve. The signaling game is embedded in a multi-sector trade environment and evaluated under a politically weighted government objective that includes downstream surplus, upstream rents, tariff revenue, and the strategic importance of Chinese suppliers.

A key feature of the framework is that political considerations and informational transmission coexist. The government's objective contains reduced-form political weights that determine how it values domestic upstream producers, downstream firms, tariff revenue, and rents earned by Chinese suppliers. These weights reflect broader distributional and strategic priorities associated with the Section 301 action and are not shaped by the firm–bureaucracy

lobbying studied here. Within this politically weighted environment, firms lobby to convey information about the part of the objective relevant to them. Importantly, the political weights do not change in the model; the role of lobbying is to make the government's politically constrained decisions better informed.

Three empirical facts motivate this structure. First, petitions submitted by firms that engaged in lobbying were more likely to be approved, even after controlling for industry, time, and sector characteristics. Second, consistent with Chor et al. (2025), approval outcomes varied systematically with economic indicators of domestic harm, such as reliance on Chinese inputs, indicating that the agency incorporated economically relevant information when evaluating petitions. Third, a substantial share of lobbying groups received no approvals across all their petitions, reflecting that lobbying was a costly activity with uncertain outcomes. These features, including costly participation, decisions that respond to observable fundamentals, and residual uncertainty, map directly into the signaling environment developed below.

I embed this signaling framework in a multi-sector quantitative environment calibrated to the universe of Section 301 exclusion petitions, combined with firm-level lobbying disclosures and sector-level trade data. Within each sector, products exhibit heterogeneous sourcing patterns based on pre–trade-war information, which generate variation in the product-specific economic impact of tariff relief. The model follows a three-stage sequence. First, the government chooses the Section 301 tariff while anticipating the informational environment. Second, firms privately observe the economic impact of an exclusion, summarized in the model as a type, and decide whether to incur the cost of lobbying in order to generate a signal. Third, the government evaluates petitions using the observed lobbying decisions and, when lobbying occurs, the noisy signals conveyed through those petitions, subject to idiosyncratic political shocks that capture residual randomness arising from factors such as reviewer assignment or transient shifts in political priorities. The calibration uses sector-level tariff outcomes, lobbying participation rates, and the approval rates of lobbying and non-lobbying petitions to identify the cost of lobbying, the informativeness of the signal, and the magnitude of these political shocks.

The calibrated parameters paint a clear picture. Two components of the government's reduced-form political weights are pinned down analytically: the MFN and Column 2 tariffs identify the weights on domestic upstream rents and downstream surplus. The remaining political weight, which governs how the government values Chinese upstream rents, is calibrated at the sector level by requiring the model's tariff-setting condition to match the observed Section 301 tariff. Lobbying signals are estimated to be highly informative relative

to underlying sectoral heterogeneity, implying that the government learns effectively from lobbying when it occurs. The estimated idiosyncratic approval shocks are modest, which means outcomes were predictable on average but still retained the residual randomness inherent in the review process. The main friction lies in the fixed cost of participation, which deters many firms from lobbying. As a result, the information the government receives is reliable but limited in coverage, since only a subset of firms chooses to participate.

Quantitatively, the model shows that informational lobbying reduced the welfare cost of the tariffs by about 18 percent relative to a world without lobbying. The mechanism is twofold. First, lobbying lowers the average Section 301 tariff from 0.243 to 0.207. This moderation arises because the structural model captures the private cost of participation: the government internalizes that a higher tariff would induce more firms to incur lobbying costs, and lowers the rate to mitigate this response. Second, once the tariff is chosen, lobbying conveys precise signals that raise the approval rate from 4.8 to 7.9 percent, enabling more targeted relief for the highest-harm products. Crucially, the finding that these mechanisms improve aggregate welfare is a quantitative result rather than a theoretical necessity. In the calibrated environment, the information conveyed through lobbying allows the government to direct exclusions toward products for which tariff relief generates substantial efficiency gains, while the anticipatory tariff adjustment reduces the distortions created by the tariff itself. Together, these two forces generate a welfare improvement relative to a world without lobbying, even though the government maximizes a political rather than a welfare objective.

This paper contributes to four strands of literature. First, it builds on the literature on informational lobbying, which studies how lobbying conveys private information to policymakers (e.g., Potters and Van Winden, 1992; De Figueiredo and Silverman, 2008; Ludema et al., 2010). Second, it complements the large body of quid-pro-quo models in which lobbying purchases political influence (e.g., Grossman and Helpman, 1994; Goldberg and Maggi, 1999; Gawande et al., 2012), by focusing instead on settings where influence is already embedded in the policymaker's objective and lobbying primarily supplies information. Third, this paper is most closely related to the structural analysis of the Section 301 exclusion process by Chor et al. (2025). They document that approval outcomes correlate with economic fundamentals, consistent with information transmission, and quantify the welfare value of the exclusion mechanism itself, showing that it acts as a safety valve that allows the government to set higher tariffs. I extend their analysis by explicitly modeling the firm's endogenous decision to lobby, the precision of the resulting signals, and the political shocks governing approvals. This allows me to answer a distinct counterfactual question: rather than evaluating the exclusion process as a whole, I quantify the welfare effect of the lobbying channel, showing that

the anticipation of costly signaling leads the government to moderate, rather than raise, the ex-ante tariff. Fourth, it connects to the literature on noisy signaling (e.g., Matthews and Mirman, 1983; Kartik, 2009), using these tools to characterize how information transmission operates within a politically weighted decision environment.

More broadly, the framework provides a template for analyzing administrative decision-making under incomplete information. Many modern industrial and trade policies such as Buy America waivers, export controls, investment screening, and environmental permitting require regulators to implement politically directed rules while relying on information provided by regulated parties. The Section 301 case shows that such systems can harness private information to improve implementation, even when the underlying objective is political rather than welfare maximizing.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and exclusion process in detail. Section 3 presents empirical patterns motivating the signaling framework. Section 4 develops the theoretical model of informational lobbying. Section 5 calibrates the model to sector-level data and quantifies the informational value of lobbying. Section 6 presents counterfactual analyses that decompose the welfare effects of informational lobbying. Section 7 concludes.

2 Institutional Background

2.1 The Section 301 Exclusion Process

Section 301 of the Trade Act of 1974 authorizes the U.S. Trade Representative (USTR) to impose trade measures when foreign practices are deemed "unreasonable or discriminatory" and burden U.S. commerce. In August 2017, the USTR launched an investigation into China's technology-transfer, intellectual-property, and innovation policies. Following its findings in March 2018, the United States imposed four successive tariff lists, widely known as the Section 301 tariffs, covering more than $460 billion of annual imports from China. Each list was implemented through notice-and-comment procedures under the Administrative Procedure Act and published in the *Federal Register*. As coverage expanded from roughly $50 billion to about $460 billion, the measures ultimately encompassed almost 90 percent of U.S. imports from China, with tariff rates ranging between 7.5 and 25 percent.[1]

Recognizing that these tariffs could unintentionally harm U.S. producers and consumers dependent on Chinese inputs, the USTR established a formal exclusion process in July 2018.

[1]See Chor et al. (2025) for a detailed chronology of the four lists and their coverage.

This mechanism allowed importers, trade associations, and other stakeholders to request temporary, product-specific exemptions from the additional duties. Each notice of action opened a three-month filing window during which applicants were required to demonstrate that (i) the product was available only from China, (ii) the tariff would cause severe economic harm to U.S. interests, (iii) the product was not strategically linked to China's *Made in China 2025* program, (iv) granting the exclusion would not undermine the objectives of the investigation, and (v) that the product could be defined precisely enough for Customs to administer the exemption. These criteria guided reviewers but were not applied mechanically; no single factor was dispositive.[2] The third criterion, whether a product was associated with China's *Made in China 2025* industrial policy, reflected sector-specific strategic concerns and suggests that the government's sensitivity toward Chinese producers varied across industries.

Between 2018 and 2020, the USTR received roughly 53,000 exclusion petitions across the four tariff tranches. The agency initially relied on a small team of attorneys within its Office of General Counsel to review the first lists and later expanded capacity through detailees from other agencies, contract attorneys, and technical case experts as the volume of requests surged. The review followed a four-stage structure: (1) public submission and comment through an online portal, (2) internal substantive review by assigned analysts, (3) an inter-agency "administrability" assessment involving U.S. Customs, the International Trade Commission, and representatives from 22 federal agencies, and (4) publication of final decisions in the *Federal Register*. Approved exclusions were applied retroactively to the date of the original tariff and were product-based rather than firm-specific, meaning any importer of the qualifying good could claim relief. There was no formal appeal procedure, so petitioners had a single opportunity to make their case.

Administrative strain was a defining feature of the process. Average processing times reached 143 days for List 3 alone. Nearly 87 percent of petitions were denied, most for failing to demonstrate severe economic harm or to show lack of alternative supply. The GAO attributed these patterns to the USTR's limited staffing, absence of standardized guidance, and reliance on self-reported data from firms.[3] Reviewers therefore depended heavily on the credibility and clarity of information voluntarily supplied by petitioners. This delegation of information provision effectively turned the exclusion process into an informal communication channel between private actors and a resource-constrained bureaucracy.

By early 2020, the share of U.S. imports from China subject to Section 301 tariffs exceeded

[2]Government Accountability Office (2021).

[3]Government Accountability Office (2021), pp. 2, 10–14.

50 percent, while products excluded from tariffs accounted for about 10 percent before the program wound down. The combination of massive petition volume, multi-agency review, and limited administrative capacity made the exclusion process one of the largest quasi-judicial exercises in modern U.S. trade policy, offering a natural setting for studying how lobbying and information transmission shape bureaucratic decision-making under political constraints.

The combination of open participation and limited administrative capacity created fertile ground for lobbying activity. Unlike legislative lobbying, which targets elected officials and emphasizes access, persuasion, and influence, exclusion-related lobbying targeted bureaucrats tasked with quasi-adjudicatory decisions. Firms could not alter statutory criteria but could shape the *information* available to reviewers, for example, by providing technical evidence on input substitution or production bottlenecks.

In parallel with the USTR's formal submission portal, firms engaged in lobbying activities governed by the Lobbying Disclosure Act (LDA). Disclosures compiled by the Center for Responsive Politics and Kim (2018) show a sharp rise in filings mentioning "Section 301 tariffs" or relevant issues beginning in 2018, documenting several hundred firms and trade associations that hired lobbyists to represent them or self-represented before the USTR. This registered lobbying provides a credible, legally-monitored record of firms' engagement.

3 Empirical Motivation

This section documents the empirical patterns that motivate the theoretical framework. The goal is twofold: first, to show that approval decisions were systematically linked to both lobbying participation and indicators of economic harm (Section 3.2); and second, to show that approval outcomes exhibited uncertainty, with patterns consistent with an informational channel in which lobbying conveys useful but imperfect information (Section 3.3). Together, these findings motivate a model in which lobbying operates as a costly mechanism for transmitting information to the regulator.

3.1 Overview of Section 301 Exclusion Petitions

Table 1 provides a summary of the petition universe. Across all four tranches, roughly nine percent of applicant firms engaged in lobbying that explicitly referenced Section 301 in their disclosure filings. Approval rates declined sharply over time, from more than 30 percent in the first two tranches to below 7 percent in the final lists, consistent with the rise in workload and tighter political oversight as coverage expanded.

Table 1: Overview of Section 301 Exclusion Petitions

Tranche	Total Applications	Unique Companies	Lobbied (%)	Approval Rate (%)	Approval (Lobbied, %)
Overall	52,745	4,575	9.17	12.90	22.38
Tranche 1	10,813	1,200	12.59	32.81	33.11
Tranche 2	2,869	453	13.38	37.43	62.76
Tranche 3	30,283	2,513	7.45	4.94	14.13
Tranche 4	8,780	1,249	9.53	6.55	10.27

Notes: Each row reports the total number of exclusion petitions and approval outcomes by tariff tranche.

These aggregates already suggest two salient features. First, lobbying participation was selective: only a small fraction of petitioners invested in formal lobbying, suggesting that the cost of participation was substantial relative to the associated increase in approval probability. Second, while lobbying correlates positively with approval rates, the relationship is far from mechanical: many lobbying firms still saw all their petitions denied. This pattern is consistent with a setting in which lobbying is costly and informative and in which approval outcomes remain uncertain rather than guaranteed.

The remainder of this section explores the underlying structure of these patterns.

3.2 Determinants of Exclusion Approval

Data sources. The empirical analysis links exclusion petitions from the USTR public docket to firm-level lobbying disclosures from LobbyView, consolidating subsidiaries under common corporate-group identifiers from the Bureau van Dijk *Orbis* database. This aggregation is essential, as firms often channel lobbying expenditures through a parent entity while submitting exclusion petitions through multiple operational affiliates. Industry characteristics and trade elasticities are drawn from Soderbery (2015), while tariff schedules come from the U.S. International Trade Commission (USITC) and import values from the U.S. Census Bureau. To construct product-level measures of domestic reliance, I combine data from the 2017 BEA Input–Output tables and the 2017 Economic Census Product Statistics. I start from sector-level domestic production totals in the BEA tables and disaggregate them to the product level using shipment-value weights derived from the Economic Census. This approach allocates each sector's domestic production across products within the same *North American Industry Classification System* (NAICS) industry, yielding granular domestic-to-import ratios for the *Harmonized Tariff Schedule* (HTS)–NAICS concordance used in the analysis. The resulting dataset covers the full population of Section 301 petitions and associated lobbying activities at the product level, mapped to input–output sectors used later in the calibration.

Specification. The baseline regression takes the exclusion-approval indicator as the dependent variable. The key explanatory variable is a dummy for whether the firm group (`guo25`) engaged in lobbying that explicitly targeted Section 301 in the same tranche. Control variables include the standardized import demand elasticity (σ), the standardized inverse export supply elasticity (ω), the China import share at the HTS 10-digit level, the domestic share constructed from the BEA Input–Output and Economic Census data described above, and the number of exclusion applications filed under the same HTS product code. The specification includes tranche, year, and industry fixed effects. I estimate the model using both an ordinary least squares linear probability model (LPM) and a logit specification, as reported in Table 2.

Table 2: Determinants of Exclusion Approval

	(1) OLS (LPM)	(2) Logit
Lobbied	0.0500**	0.376**
	(0.0238)	(0.191)
Demand Elasticity (σ)	-0.0156***	-0.156***
	(0.0021)	(0.0347)
Inverse Supply Elasticity (ω)	-0.0029***	-0.148
	(0.0009)	(0.112)
China Import Share	0.309*	3.343**
	(0.185)	(1.603)
Domestic Share	0.265*	2.288*
	(0.145)	(1.194)
Tranche & Year FEs	Yes	Yes
Industry FE	Yes	Yes
Observations	52,355	52,355

Notes: Dependent variable = indicator for exclusion approval. All specifications include tranche, year, and industry fixed effects. Each specification also controls for the number of exclusion applications per HTS product code (coefficients not reported). Standard errors are two-way clustered by corporate group (`guo25`) and six-digit HTS code. Elasticities (σ and ω) are standardized to mean 0 and unit variance.

Results and interpretation. Lobbying is a strong positive predictor of approval: petitions associated with lobbying firms are about 5.0 percentage points more likely to be granted in the OLS model, and the implied marginal effect from the logit specification is approximately 3.4 percentage points (average marginal effect). Given an unconditional approval rate of 13 percent, this represents a sizable difference even after accounting for industry and time fixed effects.

Beyond lobbying, approval probabilities vary systematically with sectoral fundamentals in ways consistent with the USTR's stated goal of mitigating domestic harm. The economic interpretation of the coefficients is straightforward.

- A higher **import demand elasticity** ($\boldsymbol{\sigma}$) means that buyers can more easily substitute away from affected suppliers. Greater substitutability allows firms to avoid tariff burdens, resulting in lower domestic harm and, consequently, a lower likelihood of approval.
- A higher **inverse export supply elasticity** ($\boldsymbol{\omega}$) implies that foreign suppliers have more inelastic supply, causing them to absorb a larger share of the tariff through lower export prices. This reduction in domestic hardship *should* correspond to a lower probability of approval. We find a significant negative relationship in the OLS model (Col. 1), though the point estimate in the logit specification (Col. 2), while negative, is not statistically precise.
- Finally, products with higher **China import shares** or higher **domestic shares** are significantly more likely to be approved, reflecting the government's responsiveness to sectors more directly exposed to tariff-induced disruption.

Together, these patterns indicate that the agency's decisions aligned with economic harm considerations. The regression results therefore support two complementary findings: lobbying independently predicts approval outcomes, and the pattern of coefficients on σ, ω, and exposure variables is consistent with the government acting on economically relevant information. At the same time, lobbying did not guarantee success, and non-lobbying firms also secured approvals, indicating residual variation beyond observable sectoral characteristics.

3.3 Uncertainty and Imperfect Screening

The data indicate that the relationship between lobbying and approval outcomes is noisy and imperfect. Table 3 summarizes success rates at the *corporate-group* level. A substantial share of lobbying groups had every petition denied, despite in some cases having filed a large number of requests. Conversely, aggregate data show that roughly 9 percent of petitions submitted by non-lobbying firms were approved.

These patterns motivate one core feature of the theoretical framework. In the data, both lobbying and non-lobbying firms receive a mix of approvals and denials, indicating that approval was not mechanically determined by participation status. To capture this, the model incorporates *idiosyncratic political shocks* that introduce residual uncertainty into the

approval decision.

In addition, while lobbying is associated with higher approval rates, it does not transmit information perfectly. Submissions vary in clarity, and agencies do not always interpret complex materials with full precision. To reflect this imperfect screening, the model allows for *signal noise*. This feature makes the informational channel more realistic and gives the framework enough flexibility to accommodate the magnitude of the approval-rate gap between lobbying and non-lobbying petitions.

Table 3: Lobbying Corporate Groups with Zero Approvals

Metric	Zero-Approval Groups	All Lobbying Groups
Number of Groups	89	203
Total Petitions	1,310	4,839
Mean Petitions per Group	14.7	23.8
Median	4	5
Maximum	201	655
Minimum	1	1

Notes: Each observation represents a corporate group (`guo25`) aggregating all affiliated firms' petitions. Zero-approval groups submitted at least one petition but received no approvals.

3.4 Summary

Three empirical regularities emerge from the Section 301 data:

1. **Lobbying predicts approval.** Petitions associated with lobbying firms are substantially more likely to be granted, even after controlling for industry, time, and sectoral characteristics (Section 3.2).

2. **Approvals track economic harm.** The probability of approval varies systematically with economic fundamentals such as substitutability and import exposure, suggesting that the regulator incorporated economically relevant information when evaluating petitions (Section 3.2).

3. **Approval is uncertain for all firms.** Some lobbying groups received no approvals despite extensive participation, and a nontrivial share of non-lobbying firms secured approvals (Section 3.3). These features indicate that approval was not mechanically tied to participation and that success remained uncertain.

Taken together, these patterns suggest a consistent interpretation. Consistent with Chor et al. (2025), approval outcomes correlated with indicators of economic harm, which points to a decision process that responded to information about potential domestic damage. The

patterns documented here complement their evidence by showing that this informational transmission was imperfect. While lobbying is strongly associated with a higher probability of success, the data also show substantial residual uncertainty, including lobbying denials and non-lobbying approvals, which indicates that the process was noisy rather than deterministic. These empirical regularities motivate the theoretical framework developed next, which relaxes the full-information structure in prior work and models lobbying as a costly and noisy signaling mechanism.

4 Theoretical Framework

4.1 Overview

This section develops the model of informational lobbying that governs the product–level exclusion process. The exclusion stage begins after tariffs are chosen, when firms may lobby to convey private information about the economic impact of tariff relief. The analysis lays out the timing of the interaction and the information available to each player, and then characterizes how lobbying, noisy signals, and idiosyncratic approval shocks shape incentives and approval decisions. Section 4.8 derives the object that firms privately observe from a structural trade environment, embedding the signaling stage within the broader model of tariff determination.

Tariffs are chosen at the *sector* level, consistent with the structure in Chor et al. (2025). Within each sector, individual products are more granular than the tariff category and differ in their contributions to the government's objective. The government observes only the distribution of product–level fundamentals, not their realizations.

The overall environment is represented as a sequential policy problem in which tariff setting (Stage 1) precedes a Bayesian exclusion game (Stages 2 and 3). In Stage 1, the government chooses the sectoral tariff while anticipating the equilibrium of the exclusion stage. In Stage 2, after tariffs are set, firms observe product–specific fundamentals and decide whether to incur a fixed cost to submit information through lobbying; lobbying generates a noisy signal. In Stage 3, the government combines this signal with its prior and an idiosyncratic political shock to form posterior beliefs and decide whether to approve the exclusion.

Equilibrium is defined as a Perfect Bayesian Equilibrium of the exclusion stage. Under the single–crossing structure derived below, firms adopt cutoff lobbying strategies and the government's approval rule is based on posterior beliefs. The subsections that follow develop the model's timing and information structure, characterize equilibrium behavior, and then

connect the signaling environment to its structural microfoundations.

4.2 Model Setup and Timing

The exclusion process is modeled as a three–stage game between the government and the firm associated with a product g drawn from a continuum of products within a sector s.

Sectors and Products. The model distinguishes between the level at which tariffs are set (the sector) and the level at which exclusions are requested (the product). The government's objective is additively separable across products, as shown in Section 4.8, so each product–level problem can be analyzed independently.

Stage 1: Tariff Choice. The government chooses a sector–specific tariff τ_s to apply to all products in sector s. This decision maximizes its expected politically weighted objective for the sector, anticipating the equilibrium of the exclusion stage. The government knows the distribution of economic characteristics across products, but not the realization for any specific product.

Stage 2: Information, Payoffs, and Lobbying. After the tariff is set, the firm associated with product g privately observes its type

$$\theta_{sg}(\tau_s),$$

which represents the contribution of granting an exclusion for that product to the government's politically weighted objective at tariff level τ_s. Section 4.8 derives $\theta_{sg}(\tau_s)$ and the mapping $B(\theta)$ from a structural trade environment. In the signaling analysis below, $B(\theta)$ is treated as a reduced-form object whose monotonicity follows under mild restrictions on sectoral primitives.

The firm chooses whether to lobby,

$$d_{sg} \in \{0, 1\},$$

at a fixed cost c. If the firm lobbies, the policymaker observes a noisy signal:

$$\tilde{\theta} = \theta_{sg}(\tau_s) + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

If the firm does not lobby, the government receives no informative signal.

The firm's expected payoff from choosing d_{sg}, taking the government's approval behavior as given, is

$$U(d_{sg}) = \Pr(A_{sg} = 1 \mid d_{sg})\, B(\theta_{sg}(\tau_s)) - c \cdot d_{sg}.$$

Stage 3: Approval Decision. The government observes d_{sg} and, when $d_{sg} = 1$, the signal $\tilde{\theta}$. It also receives an idiosyncratic political shock

$$\eta \sim \mathcal{N}(0, \sigma_\eta^2),$$

which is realized and observed by the government before the approval decision. Given its posterior belief about $\theta_{sg}(\tau_s)$, the government approves the exclusion when the expected change in its objective is non–negative:

$$A_{sg} = 1 \quad \Longleftrightarrow \quad \mathbb{E}[\theta_{sg}(\tau_s) \mid d_{sg}, \tilde{\theta}] + \eta \geq 0.$$

Model Objects and Microfoundation. The signaling stage depends on the firm's type $\theta_{sg}(\tau_s)$, the benefit $B(\theta_{sg})$, the cost c, the signal noise variance σ_ε^2, and the political shock variance σ_η^2. Section 4.8 derives $\theta_{sg}(\tau_s)$ from a structural trade environment and links $B(\theta)$ to underlying sectoral primitives, while c, σ_ε^2, and σ_η^2 are treated as exogenous primitives.

4.3 Information and Beliefs

Each good g in sector s is characterized by a privately observed type

$$\theta_{sg}(\tau_s),$$

which represents the change in the government's objective from granting an exclusion at tariff level τ_s. Once τ_s is set, this mapping pins down the realization of θ_{sg} for that product. Because the signaling subgame analyzes the interaction for a fixed sector s and product g, we omit the subscripts and write the realized type simply as θ. Throughout Sections 4.3–4.6, θ therefore stands for the object $\theta_{sg}(\tau_s)$. The distribution of types within sector s has density $f_s(\theta)$ with support $[\underline{\theta}_s, \overline{\theta}_s]$. Let $\Theta_s = [\underline{\theta}_s, \overline{\theta}_s]$ denote this type space.

A firm that lobbies pays a cost $c > 0$ and generates a noisy signal:

$$\tilde{\theta} = \theta + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

This signal conveys information without revealing the type perfectly.

The policymaker forms posterior beliefs about θ using Bayes' rule. Let $m_s(\tilde{\theta})$ denote the marginal density over observed signals induced by the equilibrium strategy. Under normal noise, the posterior mean satisfies Tweedie's formula:[4]

$$\mathbb{E}[\theta \mid \tilde{\theta}] = \tilde{\theta} + \sigma_\varepsilon^2 \frac{\partial}{\partial \tilde{\theta}} \ln m_s(\tilde{\theta}).$$

Because signals are normally distributed around the true type, they satisfy the Monotone Likelihood Ratio Property (MLRP): higher signal realizations imply stochastically higher types. This monotonicity is consistent with the informational structure in Matthews and Mirman (1983). The marginal density $m_s(\tilde{\theta})$ incorporates any endogenous selection implied by equilibrium behavior, including the fact that only certain types may lobby in equilibrium.

If the firm does not lobby, the policymaker receives no signal beyond non–participation. Let

$$\mu_s^0 = \mathbb{E}[\theta \mid d = 0]$$

denote the expected type of a non–lobbying firm.

The government's approval decision incorporates both the posterior belief about θ and the realized political shock $\eta \sim \mathcal{N}(0, \sigma_\eta^2)$.

4.4 Strategies and Payoffs

Strategies. A firm of type $\theta \in \Theta_s$ chooses a lobbying action $d(\theta) \in \{0, 1\}$. If $d = 1$, the firm generates a signal

$$\tilde{\theta} = \theta + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2),$$

while if $d = 0$, no signal is observed.

The policymaker observes the action d and, when $d = 1$, the signal $\tilde{\theta}$. Given an information set $\mathcal{I} \in \{d = 0, (d = 1, \tilde{\theta})\}$ and the realization of the political shock η, the policymaker chooses an approval outcome

$$A(\mathcal{I}, \eta) \in \{0, 1\}.$$

Government's Decision. Let $\mu(\cdot \mid \mathcal{I})$ denote the policymaker's belief about θ given the information set $\mathcal{I}$. The government approves the exclusion whenever

$$\mathbb{E}_\mu[\theta \mid \mathcal{I}] + \eta \geq 0.$$

[4] See Robbins (1956) and Efron (2011).

The approval condition implies the following ex ante approval probability at information set $\mathcal{I}$:

$$\mathcal{P}(\mathcal{I}) = \Phi\left(\frac{\mathbb{E}_\mu[\theta \mid \mathcal{I}]}{\sigma_\eta}\right),$$

where $\Phi(\cdot)$ is the standard normal CDF.

Firm's Decision. The firm chooses d to maximize its expected net payoff, taking as given the government's decision rule and belief system $\mu(\cdot \mid \cdot)$.

If the firm lobbies ($d = 1$), type θ generates a signal $\tilde{\theta} = \theta + \varepsilon$, which the government uses to form $\mathcal{P}(d = 1, \tilde{\theta})$. The expected payoff from lobbying is

$$U(d = 1 \mid \theta) = B(\theta) \cdot \mathbb{E}_\varepsilon\Big[\mathcal{P}(d = 1, \tilde{\theta}) \,\Big|\, \theta\Big] - c,$$

where the expectation is taken over ε.

If the firm does not lobby ($d = 0$), the government conditions only on the event $d = 0$, so the expected payoff from not lobbying is

$$U(d = 0 \mid \theta) = B(\theta) \cdot \mathcal{P}(d = 0).$$

A firm of type θ chooses $d = 1$ if and only if $U(d = 1 \mid \theta) \geq U(d = 0 \mid \theta)$.

4.5 Equilibrium Concept

This subsection characterizes the equilibrium of Stages 2 and 3, taking the sectoral tariff τ_s chosen in Stage 1 as given. The tariff influences the exclusion stage only through its effect on $\theta(\tau_s)$. In the full model, the optimal tariff maximizes a sectoral objective that integrates over equilibrium outcomes in the exclusion stage.

Strategy Spaces. A firm of type $\theta \in \Theta_s$ chooses whether to lobby, $d(\theta) \in \{0, 1\}$. If the firm lobbies ($d = 1$), the policymaker observes $\tilde{\theta} = \theta + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$; $\tilde{\theta}$ has full support on $\mathbb{R}$. If the firm does not lobby ($d = 0$), no signal is observed. Given an information set $(d, \tilde{\theta}) \in \{(0, \cdot), (1, \tilde{\theta})\}$ and the realization of the political shock η, the policymaker chooses an approval outcome $A(d, \tilde{\theta}, \eta) \in \{0, 1\}$.

Beliefs. Beliefs are mappings $\mu(\theta \mid d, \tilde{\theta})$ from information sets into posterior distributions over types. We write $\mathbb{E}_\mu[\theta \mid d, \tilde{\theta}]$ for the posterior mean implied by these beliefs. Beliefs must satisfy Bayes' rule at all on–path information sets. Because the signal density is

positive everywhere conditional on lobbying, Bayes' rule determines beliefs following any observed signal. At off–path information sets (such as observing $d = 1$ when no types lobby, or $d = 0$ when all types lobby), beliefs must satisfy the standard admissibility requirements for monotone signaling games: they must place probability only on types for whom the deviation is feasible and remain consistent with the monotonicity of the lobbying rule. Any belief system satisfying these principles is admissible in a PBE.

Definition (Perfect Bayesian Equilibrium). A Perfect Bayesian Equilibrium (PBE) of the exclusion stage consists of (i) a lobbying strategy $d^*(\theta)$, (ii) a government decision rule $A^*(d, \tilde{\theta}, \eta)$, and (iii) beliefs $\mu^*(\theta \mid d, \tilde{\theta})$, such that:

1. *Sequential rationality for firms.* Given A^* and μ^*, each type θ chooses $d \in \{0, 1\}$ to maximize its expected payoff.

2. *Sequential rationality for the policymaker.* Given beliefs and the realization of the political shock η, the policymaker approves an exclusion whenever

$$\mathbb{E}_{\mu^*}[\theta \mid d, \tilde{\theta}] + \eta \geq 0.$$

3. *Belief consistency.* Beliefs follow Bayes' rule at all on–path information sets and satisfy the admissibility requirements for monotone signaling games described above at off–path information sets.

Monotone Private Benefit. We first formalize the property of the firm's private gain from exclusion that will underlie the equilibrium characterization.

Assumption 1 (Monotone private benefit). The benefit from approval $B(\theta)$ is weakly increasing in the firm's type: if $\theta' > \theta$, then $B(\theta') \geq B(\theta)$. In Section 4.8 we show that this monotonicity property follows from the structural trade environment under mild restrictions on sectoral primitives.

Monotone Approval for Lobbyists. Given the information structure in Section 4.3, higher types generate stochastically higher signals and therefore face higher approval probabilities when they lobby.

Lemma 1 (Monotone approval for lobbying firms). Fix a PBE of the exclusion stage. Let $P_1(\theta)$ denote the approval probability for a firm of type θ that chooses to lobby, given the equilibrium beliefs and approval rule. Then $P_1(\theta)$ is strictly increasing in θ.

Proof. Consider two types $\theta' > \theta$. When a firm of type $x \in \{\theta, \theta'\}$ lobbies, the signal realization is $\tilde{\theta} = x + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. Because the signal density satisfies the Monotone Likelihood Ratio Property, the distribution of $\tilde{\theta}$ conditional on θ' first–order stochastically dominates that conditional on θ.

In any PBE, the government's approval rule is increasing in the posterior mean $\mathbb{E}[\theta \mid \tilde{\theta}, d = 1]$, and the posterior mean is strictly increasing in $\tilde{\theta}$ under MLRP with normal signals. Hence, for every realization of the political shock η, the probability that the approval condition $\mathbb{E}[\theta \mid \tilde{\theta}, d = 1] + \eta \geq 0$ is satisfied is strictly higher when the true type is θ' than when it is θ. It follows that $P_1(\theta') > P_1(\theta)$. □

Single–Crossing and Threshold Equilibrium. We now combine Lemma 1 with Assumption 1 to show that the net gain from lobbying is increasing in type and that equilibrium lobbying behavior takes a cutoff form.

Lemma 2 (Single–crossing in lobbying incentives). Let P_0 denote the approval probability for a firm that does not lobby. In this environment, P_0 is constant across types because the government conditions only on the event $d = 0$ when no signal is submitted. Define the net benefit from lobbying as

$$\Delta U(\theta) \;\equiv\; B(\theta)\big[P_1(\theta) - P_0\big] - c.$$

Under Assumption 1 and Lemma 1, $\Delta U(\theta)$ is strictly increasing in θ on any domain of types where $P_1(\theta) > P_0$.

Proof. By Lemma 1, $P_1(\theta)$ is strictly increasing in θ, while P_0 is constant (for any fixed belief about the non-lobbying pool). By Assumption 1, $B(\theta)$ is weakly increasing. For any $\theta' > \theta$ such that $P_1(\theta') > P_0$, we have

$$\Delta U(\theta') - \Delta U(\theta) = B(\theta')\big[P_1(\theta') - P_0\big] - B(\theta)\big[P_1(\theta) - P_0\big] > 0.$$

Hence $\Delta U(\theta)$ is strictly increasing in θ on any interval where $P_1(\theta) > P_0$. □

Proposition 1 (Threshold equilibrium). In any PBE of the exclusion stage, lobbying strategies take a threshold form: there exists a cutoff type $\theta_s^*(\tau_s) \in \Theta_s$ such that firms with $\theta \geq \theta_s^*(\tau_s)$ lobby and firms with $\theta < \theta_s^*(\tau_s)$ do not. At least one such equilibrium cutoff exists.

Proof. The proof proceeds in two steps.

(i) Necessity of threshold form. Fix any PBE. If a firm of type θ lobbies, then $\Delta U(\theta) \geq 0$, which implies $P_1(\theta) > P_0$ (since $c > 0$). By Lemma 2, for any $\theta' > \theta$, $\Delta U(\theta') > \Delta U(\theta) \geq 0$. Thus, if type θ lobbies, all higher types must also lobby. The lobbying set must therefore be an upper interval $[\theta^*, \overline{\theta}_s]$.

(ii) Existence. Define the function $\Psi(\theta)$ as the net incentive to lobby for a firm of type θ when the government believes the cutoff is exactly θ:

$$\Psi(\theta) \equiv B(\theta)[P_1(\theta;\theta) - P_0(\theta)] - c.$$

Here, $P_1(\theta;\theta)$ and $P_0(\theta)$ are the approval probabilities induced by the belief cutoff θ. An equilibrium cutoff θ^* is a solution to $\Psi(\theta^*) = 0$, or a boundary solution.

Because P_1 and P_0 are continuous in the cutoff (as integrals of smooth densities) and $B(\cdot)$ is continuous, $\Psi(\theta)$ is a continuous function on the compact interval $\Theta_s = [\underline{\theta}_s, \overline{\theta}_s]$.

We evaluate the function at the boundaries:

- If $\Psi(\underline{\theta}_s) \geq 0$, then even when the government believes the pool includes the lowest types (the most pessimistic belief), the lowest type still finds it profitable to lobby. In this case, $\theta^* = \underline{\theta}_s$ (all types lobby).
- If $\Psi(\overline{\theta}_s) \leq 0$, then even when the government believes the pool consists only of the highest types (the most optimistic belief), the highest type does not find it profitable to lobby. In this case, $\theta^* = \overline{\theta}_s$ (no types lobby).
- If neither boundary condition holds, then $\Psi(\underline{\theta}_s) < 0$ and $\Psi(\overline{\theta}_s) > 0$. By the Intermediate Value Theorem, there exists at least one interior $\theta^* \in (\underline{\theta}_s, \overline{\theta}_s)$ such that $\Psi(\theta^*) = 0$.

Thus, an equilibrium cutoff always exists. □

4.6 Optimal Tariff Choice

In the first stage of the policy environment, the government chooses a sectoral tariff τ_s anticipating that firms will subsequently request exclusions and that approvals will be determined according to the Perfect Bayesian Equilibrium of the exclusion stage (Stages 2 and 3). This subsection characterizes the government's Stage 1 problem and the connection between τ_s and the exclusion-stage equilibrium derived above.

Government Objective. For each sector s, let $W_s(\tau_s)$ denote the government's expected objective conditional on setting τ_s and on the equilibrium behavior of firms and the policy-

maker in the exclusion stage. Because the government does not observe product-level types, $W_s(\tau_s)$ integrates over the sectoral distribution of types and over the randomness in signals and political shocks. Formally,

$$W_s(\tau_s) = \Phi_s(\tau_s) + \mathbb{E}_{\theta,\varepsilon,\eta}\big[\,\theta(\tau_s)\,A^*\big(d^*(\theta),\,\theta+\varepsilon,\,\eta\big)\big]\,,$$

where the expectation is taken over $\theta \sim f_s(\cdot)$, signal noise ε, and political shocks η.

The term $\Phi_s(\tau_s)$ collects the components of the sectoral objective that depend directly on the tariff when it is applied uniformly to all goods, such as tariff revenue and other sector-level considerations specified in the microfoundation (Section 4.8). The second term captures the expected marginal contribution of the exclusion process by integrating $\theta(\tau_s)$ over those goods for which the exclusion is granted ($A^* = 1$). Because $\theta(\tau_s)$ varies with the tariff, the equilibrium lobbying cutoff $\theta_s^*(\tau_s)$ is also endogenous to τ_s.

Stage 1 Problem. Anticipating this dependence, the government chooses τ_s to maximize its sectoral objective:

$$\tau_s^* \;\in\; \arg\max_{\tau_s \geq 0} W_s(\tau_s),$$

where the mapping from τ_s to equilibrium outcomes is governed by the PBE characterized in Section 4.5. The optimal tariff balances the direct effect of τ_s on $\Phi_s(\tau_s)$ with the indirect effects operating through equilibrium exclusion behavior, namely the lobbying cutoff $\theta_s^*(\tau_s)$ and the approval rule A^*.

Implementation. A closed-form expression for τ_s^* is generally not available because the cutoff $\theta_s^*(\tau_s)$ does not admit an analytical solution once product-level heterogeneity and noisy signaling are introduced. The quantitative analysis therefore evaluates $W_s(\tau_s)$ numerically and computes the tariff that maximizes the sectoral objective given the equilibrium mapping

$$\tau_s \;\mapsto\; \theta_s^*(\tau_s) \;\mapsto\; A^*(d, \tilde{\theta}, \eta).$$

This yields a model-consistent characterization of optimal tariff policy under the informational frictions present in the exclusion stage.

4.7 Connection to the Government's Objective

The government's evaluation of each product combines domestic upstream rents, downstream surplus, tariff revenue, and weighted foreign upstream rents. Formally, for each product g

in sector s, the government's objective is:

$$\Omega_{sg} \;=\; \lambda_s \, \Pi^{\text{up}}_{sgh} \;+\; \beta_s \, \Pi^{\text{down}}_{sg} \;+\; TR_{sg} \;+\; \sum_{i \in \{c,f\}} \gamma_{s,i} \, \Pi^{\text{up}}_{sgi}. \tag{1}$$

Here, Π^{up}_{sgh} and Π^{up}_{sgi} denote upstream producer rents for domestic and foreign varieties, respectively; Π^{down}_{sg} denotes the downstream firm's operating profit; and TR_{sg} represents tariff revenue. The parameters λ_s, β_s, and $\gamma_{s,i}$ are the political weights the government assigns to each component of the objective. Here h denotes the domestic origin, while c and f denote Chinese and rest-of-world foreign origins, respectively.

In the signaling environment, the relevant reduced-form object is the product-level change in this objective that results from granting an exclusion. The firm's private type θ_{sg} is defined precisely as this change:

$$\theta_{sg}(\tau_s) \equiv \Omega_{sg}(0) - \Omega_{sg}(\tau_s).$$

Because only the firm observes its product-specific sourcing fundamentals, the value of θ_{sg} is known to the firm but not to the government, even though the mapping from primitives to Ω_{sg} is common knowledge. Lobbying decisions therefore provide information about θ_{sg}. Firms with higher θ_{sg} have stronger incentives to lobby, and the government interprets observed lobbying behavior and noisy signals as imperfect indicators of the underlying objective impact. The signaling game thus operates on a reduced-form representation of the government's objective: it takes θ_{sg} as the payoff-relevant economic quantity that shapes both the firm's incentives and the government's posterior beliefs.

The next subsection derives the specific functional forms of these rents, profits, and revenue from a structural multi-sector trade environment. This microfoundation explicitly links the components of Ω_{sg} to the underlying primitives, expressing θ_{sg} as a function of sector-level parameters—such as supply elasticities and political weights—and product-specific sourcing fundamentals.

4.8 Microfounding the Objective Component

The signaling game analyzes the interaction between the government and a firm associated with a specific product g. To derive the underlying types $\{\theta_{sg}\}$ and private benefits $B(\theta)$ used in that game, this section embeds the interaction in a multi-sector trade environment. A complete derivation of the production and supply relationships used below appears in Appendix A.

Environment, Utility, and Separability. The economy consists of multiple sectors s. Within each sector, a continuum of differentiated varieties $g \in [0,1]$ are produced. Final consumers have quasi-linear utility over a CES composite of these final goods, together with a freely traded numeraire:

$$U = X_0 \; + \; \left[\int_s \int_{g \in s} \phi_{sg}^{\frac{1}{\rho}} C_{sg}^{\frac{\rho-1}{\rho}} \, dg \, ds \right]^{\frac{\rho}{\rho-1}}, \qquad \rho > 1.$$

Because each final good (s, g) is traded on world markets under perfect competition, its price is fixed from the perspective of the domestic economy. Quasi-linearity implies that total expenditure on this CES composite is constant, and the fixed world prices imply that the quantities demanded of final goods C_{sg} are invariant to a tariff exclusion applied to upstream inputs. Under this utility structure, aggregate welfare is additively separable across products, and because upstream rents, downstream profits, and tariff revenue are also defined at the product level, the government's politically weighted objective inherits this additive separability. Consequently, the government's decision rule for granting an exclusion to any specific product g depends only on that product's contribution to the aggregate objective. This separability allows the exclusion stage to be analyzed as a collection of independent signaling games, one for each product.

Downstream Production. Each product (s, g) is used by a competitive downstream firm that transforms a composite intermediate input X_{sg} into final output according to

$$Y_{sg} = B_s \ln X_{sg},$$

where $B_s > 0$ is a sector-specific technology parameter. The composite input aggregates domestic, Chinese, and rest-of-world varieties through a CES structure. The exact CES aggregator used to construct X_{sg} and the full expression for the associated price index P_{sg} are provided in Appendix A.

With the final-good price normalized to 1 and taking P_{sg} as given, the downstream producer chooses X_{sg} to maximize

$$\Pi_{sg}^{\text{down}} = B_s \ln X_{sg} - P_{sg} X_{sg}.$$

The first-order condition,

$$\frac{B_s}{X_{sg}} = P_{sg},$$

implies

$$P_{sg}X_{sg} = B_s.$$

Thus the downstream firm's total expenditure on intermediate inputs,

$$M_{sg} \equiv P_{sg}X_{sg},$$

is constant and equal to B_s. Tariff exclusions therefore affect downstream producers only through the input price index P_{sg}. Substituting the optimal X_{sg} yields the reduced-form expression

$$\Pi_{sg}^{\text{down}} = -\,B_s \ln P_{sg} + \text{constant},$$

so that P_{sg} is the only tariff-relevant channel through which downstream producers enter the government's objective.

Upstream Supply. Each origin $i \in \{h, c, f\}$ supplies an intermediate-input variety using labor and a specific factor under a standard specific-factor technology, with the full derivation of the supply curve and the constant rent share provided in Appendix A. As shown in Appendix A, this structure implies an inverse supply elasticity ω_s, yielding a constant-elasticity supply curve and a constant specific-factor rent share. Letting p_{sgi} denote the producer price and q_{sgi} the supplied quantity, producer revenue is

$$R_{sgi} \equiv p_{sgi}q_{sgi}.$$

Because downstream expenditure equals $(1+\tau_{sgi})\, p_{sgi}q_{sgi}$ but only $p_{sgi}q_{sgi}$ accrues to upstream producers, the specific factor earns the fraction

$$\Pi_{sgi}^{\text{up}} = \frac{\omega_s}{1+\omega_s} R_{sgi},$$

with the full derivation provided in Appendix A.

Government Preferences. The government evaluates each product using the politically weighted objective function defined in Equation (1). Given the production structure derived above, the contribution of product g in sector s is

$$\Omega_{sg}(\tau_s) = \lambda_s \Pi_{sgh}^{\text{up}} + \beta_s \Pi_{sg}^{\text{down}} + TR_{sg} + \sum_{i \in \{c,f\}} \gamma_{s,i} \Pi_{sgi}^{\text{up}}.$$

Each term depends on the *equilibrium sourcing structure* of the downstream firm, captured by the sourcing shares s_{sgi}. These shares are determined by the product-specific fundamentals

(A_{sgi}, K_{sgi}), the sectoral elasticities (σ_s, ω_s), and the sector-level tariff τ_s.

Mapping to Types. In the signaling model, the reduced-form objects θ_{sg} and $B_s(\theta_{sg})$ represent, respectively, the change in the government's objective from granting an exclusion and the firm's private benefit from that exclusion. Their structural interpretation follows from the fact that both depend on the same equilibrium sourcing shares.

Each firm privately observes its own product-specific fundamentals (A_{sgi}, K_{sgi}) that determine its sourcing shares s_{sgi}. The government observes only the distribution of these fundamentals at the sector level. Because the mapping from fundamentals to sourcing shares and then to Ω_{sg} is common knowledge, a firm's private information about its sourcing structure translates directly into private information about its contribution to the government's objective.

The relevant type is the product-level change in the objective from reducing the tariff from τ_s to 0:

$$\theta_{sg}(\tau_s) \equiv \Omega_{sg}(0) - \Omega_{sg}(\tau_s).$$

Thus θ_{sg} is a function of the sector-level parameters $(\sigma_s, \omega_s, B_s, \lambda_s, \beta_s, \gamma_{s,i}, \tau_s)$ and the product's equilibrium sourcing shares s_{sgi}, which are determined jointly by the firm's private fundamentals (A_{sgi}, K_{sgi}), the sectoral elasticities (σ_s, ω_s), and the tariff τ_s.

Substituting the expressions for upstream rents, downstream profits, and tariff revenue yields

$$\theta_{sg}(\tau_s) = \underbrace{\beta_s B_s \ln\left(\frac{P_{sg}(\tau_s)}{P_{sg}(0)}\right)}_{\text{Gain to Downstream}} - \underbrace{[TR_{sg}(\tau_s) - TR_{sg}(0)]}_{\text{Lost Revenue}} + \Delta\text{Weighted Rents},$$

where the full expression for ΔWeighted Rents appears in Appendix A.

The firm's private benefit is

$$B_s(\theta_{sg}) = \Pi_{sg}^{\text{down}}(0) - \Pi_{sg}^{\text{down}}(\tau_s).$$

Under parameter restrictions on (σ_s, ω_s) (provided in a separate appendix in a future version), $B_s(\theta_{sg})$ is monotone in θ_{sg}, ensuring the single-crossing property required for the signaling equilibrium.

4.9 Equilibrium Implications and Connection to Welfare

The cutoff equilibrium characterized above delivers three equilibrium objects that map directly into the moments used for calibration: (i) the share of firms that lobby, determined by the participation cutoff $\theta_s^*(\tau_s)$; (ii) approval rates conditional on lobbying status, denoted P_1 and P_0; and (iii) the optimal Section 301 tariff implied by the government's politically weighted objective.

These objects depend on two sets of primitives. The informational primitives $(c, \sigma_\varepsilon, \eta)$ govern lobbying incentives and the informativeness of signals, while the sector-specific structural primitives—including technology parameters $(B_s, \sigma_s, \omega_s)$ and political weights $(\lambda_s, \beta_s, \gamma_{si})$—shape the underlying economic impacts θ_{sg} and therefore the model-implied optimal tariff. Together, these primitives determine the equilibrium mapping from economic fundamentals to lobbying behavior, approval patterns, and tariff policy.

In the model, informational lobbying improves targeting conditional on participation, since higher types yield larger gains under the government's objective. At the same time, the presence of lobbying in the exclusion stage feeds back into the government's ex-ante tariff decision: because the government anticipates how firms self-select into lobbying and how signals will be interpreted, the equilibrium tariff τ_s reflects these informational considerations. Moreover, the government's posterior beliefs about non-lobbying firms differ from its prior, since the absence of lobbying is informative in equilibrium.

These forces, combined with the cost of participation, imply that the net effect of lobbying on welfare is theoretically ambiguous. The quantitative analysis evaluates the contribution of the informational channel in light of these competing forces.

5 Quantitative Implementation and Calibration

5.1 Calibration Environment and Parameter Structure

The quantitative implementation combines a small set of economy-wide signaling parameters with sector-specific heterogeneity to match the observed cross-sector variation in tariff levels, lobbying participation, and approval outcomes. Three parameters are common across all sectors: the fixed participation cost of lobbying c, the standard deviation of the noisy signal σ_ε, and the standard deviation of the idiosyncratic political shock η. These parameters jointly determine the informativeness of lobbying, the randomness in approval decisions, and the overall incidence of participation in equilibrium.

Each sector s has two additional parameters. The dispersion parameter α_{0s} governs the heterogeneity of product-level sourcing patterns within the sector, and therefore the dispersion of the welfare-relevant types $\{\theta_{sg}\}$. The political weight γ_{cs} determines how the policymaker values Chinese upstream rents when setting the Section 301 tariff. The weight on non-Chinese foreign upstream rents, γ_{fs}, is set to zero, reflecting that the Section 301 investigation focused exclusively on Chinese suppliers.

A central feature of the environment is that the Section 301 tariff is chosen in a Stage 1 problem that anticipates equilibrium behavior in the exclusion stage. Given the sector's economic fundamentals (σ_s, ω_s) and its political weights, the government selects the tariff that maximizes its politically weighted objective while accounting for how the distribution of types influences firms' lobbying decisions and exclusion outcomes. This forward-looking mapping yields, for each sector, a model-implied Section 301 tariff, the distribution of $\{\theta_{sg}\}$, the lobbying share, and the approval probabilities.

The political weights on domestic upstream rents and downstream surplus, (λ_s, β_s), are pinned down analytically from the sector's MFN (Column 1) and Column 2 tariffs, following Chor et al. (2025). Once these two components are fixed, the only remaining political-economy force shifting the optimal Section 301 tariff is γ_{cs}. The pair $(\alpha_{0s}, \gamma_{cs})$ therefore jointly determines the dispersion and level of the type distribution, and through them the sectoral patterns of lobbying and approval.

Given any candidate values of $(\alpha_{0s}, \gamma_{cs})$ and the common parameters $(c, \sigma_\varepsilon, \eta)$, the model simulates the implied type distribution, solves the signaling-game equilibrium, and computes four sectoral moments: (i) the optimal Section 301 tariff, (ii) the lobbying share, (iii) the approval rate among lobbyists, and (iv) the approval rate among non-lobbyists. These moments discipline the common parameters in the outer loop through their variation across sectors. Once $(c, \sigma_\varepsilon, \eta)$ are fixed, each sector has exactly one parameter, γ_{cs}, that shifts the tariff level and one parameter, α_{0s}, that governs dispersion. Because the three signaling-related moments arise from a single cutoff rule that separates lobbyists from non-lobbyists, sectors generally cannot match all moments exactly. The resulting residual reflects the structural restrictions of the model.

The common parameters have clear interpretations. The cost c governs the extensive margin of lobbying; a higher cost discourages participation and compresses the share of lobbying firms. The signal noise σ_ε determines how informative lobbying is about the firm's type. The political shock η governs residual randomness in approval outcomes. The next subsection describes how these ingredients enter a nested calibration procedure that aligns the model

with the cross-sectoral patterns in the data.

5.2 Heuristic Identification

Although the mapping from parameters to moments is fully equilibrium-based, it is helpful to describe heuristically how cross-sectoral variation disciplines the three common signaling parameters. The central idea is that the four sectoral moments respond differentially to changes in the cost of lobbying, the informativeness of the signal, and the magnitude of the political shock.

Lobbying participation and the cost of lobbying. Holding the informational environment fixed, the fixed cost c determines the cutoff type in the signaling game and therefore governs the share of firms that choose to lobby. Higher values of c raise the cutoff and reduce participation. Changes in σ_ε and η affect participation only indirectly through expected approval probabilities. As a result, cross-sectoral variation in lobbying shares $\{L_s\}$ primarily disciplines the level of c.

Approval gaps and signal informativeness. The difference in approval rates between lobbyists and non-lobbyists, $A_{1s} - A_{0s}$, is informative about the precision of the lobbying signal. In the model, the signal noise σ_ε directly affects only the approval probability of lobbyists, because non-lobbyists generate no signal and are evaluated using the truncated prior over types below the cutoff. As σ_ε changes, the government's posterior for lobbyists places different weight on realized signals relative to the lobbying prior, and the equilibrium cutoff adjusts, altering the composition and approval rate of the lobbying pool. In contrast, the political shock η primarily shifts approval probabilities for both lobbyists and non-lobbyists in similar ways, so it has a more limited effect on the *difference* between their approval rates. Sectoral variation in $A_{1s} - A_{0s}$, taken together with observed dispersion in outcomes and lobbying shares, therefore provides useful discipline for the signal noise parameter σ_ε.

Approval levels and political noise. The political shock η directly governs the residual randomness in approval decisions for both lobbyists and non-lobbyists. Larger values of η widen the support of the approval decision around the deterministic threshold and therefore raise approval among non-lobbyists in particular, since many non-lobbyists are evaluated close to the cutoff. In contrast, σ_ε primarily influences the sorting of firms into lobbying and affects non-lobbyist approval only indirectly through the equilibrium cutoff. For this reason, the model relies on the level of non-lobbyist approval rates $\{A_{0s}\}$—which are especially sensitive to η—to discipline the magnitude of the political shock.

Taken together, these heuristic patterns show how the three approval- and participation-based moments jointly inform the three common signaling parameters. No single moment isolates one parameter on its own, but the cross-sectoral variation in lobbying participation, approval gaps, and approval levels provides enough joint discipline to pin down the cost of lobbying, the informativeness of the signal, and the magnitude of political noise that characterize the signaling environment.

5.3 Calibration Structure

The calibration exploits four empirical moments for each sector s: (i) the observed Section 301 tariff $\tau_{s,301}$; (ii) the share of firms that lobbied L_s; (iii) the approval rate among lobbyists A_{1s}; and (iv) the approval rate among non-lobbyists A_{0s}. Taken together, these moments summarize how participation and exclusion outcomes vary across sectors, conditional on observed trade structure and political weights determined outside the calibration.

The nested calibration consists of three layers:

1. **Inner layer.** For given $(\alpha_{0s}, \gamma_{cs})$ and $(c, \sigma_\varepsilon, \eta)$, the model constructs the distribution of types $\{\theta_{sg}\}$ and solves the signaling-game equilibrium.

   Equilibrium selection. The fixed-point condition for the lobbying cutoff may, for some parameter values, admit multiple solutions. In these cases, I select the equilibrium with the lowest cutoff. This equilibrium features the largest set of lobbying types and therefore conveys the most information to the policymaker. This selection rule provides a consistent mapping from parameters to moments in the calibration. Given the resulting cutoff and approval rule, the model then produces the implied sectoral moments $(\tau_{s,301}, L_s, A_{1s}, A_{0s})$.

2. **Sector-level profiling.** Holding the common parameters fixed, each sector chooses $(\alpha_{0s}, \gamma_{cs})$ to minimize the squared deviations between the four model-implied moments and their empirical counterparts. The minimized value defines the sector-specific loss $\ell_s(c, \sigma_\varepsilon, \eta)$.

3. **Outer loop.** The common parameters $(c, \sigma_\varepsilon, \eta)$ are chosen to minimize the aggregate loss $\sum_s \ell_s(c, \sigma_\varepsilon, \eta)$.

This structure isolates the informational primitives at the economy level while allowing each sector to flexibly match its tariff and exclusion patterns through its own heterogeneity and political-weight parameters.

5.4 Sector-Level Profiling

Sector-level heterogeneity is captured through the two parameters $(\alpha_{0s}, \gamma_{cs})$, which shape how exclusion decisions respond to the sector's fundamentals.

Dispersion of product-level impacts. Product-level sourcing shares (s_{gh}, s_{gc}, s_{gf}) are drawn from a Dirichlet distribution centered on the sector's observed pre–trade-war sourcing pattern, with total mass α_{0s}. Lower values of α_{0s} generate greater dispersion, producing a wider distribution of welfare impacts $\{\theta_{sg}\}$ and more cross-sectional variation in lobbying incentives. Higher values compress the type distribution around the sector mean.

Political sensitivity to Chinese supply. The parameter γ_{cs} governs the weight placed on Chinese upstream rents and therefore shifts the optimal Section 301 tariff. The weights on domestic upstream rents and downstream surplus, (λ_s, β_s), are determined from MFN and Column 2 tariffs, while the weight on non-Chinese foreign upstream rents is set to zero. Conditional on (σ_s, ω_s), the pair $(\alpha_{0s}, \gamma_{cs})$ therefore jointly determines the sector's implied distribution of types and its Section 301 tariff.

Profiling procedure. For each candidate $(c, \sigma_\varepsilon, \eta)$, the model computes the sector's four implied moments for every trial pair $(\alpha_{0s}, \gamma_{cs})$. The profiling step chooses the pair minimizing the loss

$$(\alpha_{0s}, \gamma_{cs}) = \arg\min_{(\alpha,\gamma_c)} \Big[(\tau_{s,301} - \tau_{s,301}^{\text{model}})^2 + (L_s - L_s^{\text{model}})^2 + (A_{1s} - A_{1s}^{\text{model}})^2 + (A_{0s} - A_{0s}^{\text{model}})^2\Big].$$

Because (λ_s, β_s) are pinned down by observed MFN and Column 2 tariffs, each sector's remaining heterogeneity is summarized by the pair $(\alpha_{0s}, \gamma_{cs})$. These two parameters jointly determine the model-implied tariff and the three signaling-related moments. However, the signaling moments (L_s, A_{1s}, A_{0s}) all arise from a single cutoff equilibrium, so the sector cannot generally match all of them exactly. In practice, the fit for A_{0s} in particular reflects this structural restriction.

5.5 Calibrated Results

Table 4 reports the estimated common parameters. All parameter values are expressed per dollar of total sectoral expenditure, including both domestic production and imports. The calibrated signal noise is small ($\sigma_\varepsilon^* = 9.7 \times 10^{-7}$), implying that lobbying conveys highly precise information about firms' types. The calibrated political shock is modest:

Table 4: Calibrated Common Parameters and Economic Interpretation

Parameter	Value	Economic Meaning
c	0.00020	Lobbying cost per dollar of expenditure; higher $c \Rightarrow$ fewer firms lobby.
σ_ε	9.7×10^{-7}	Std. dev. of signal noise; small $\Rightarrow$ highly informative signals.
η	0.00070	Std. dev. of political shock; moderate relative to within-sector θ dispersion.

$\eta^* = 7.0 \times 10^{-4}$ is about forty percent of the median within-sector standard deviation of θ (1.7×10^{-3}). Approval decisions therefore exhibit some randomness but remain highly predictable on average. The primary friction in the signaling environment is instead the participation cost, $c^* = 2.0 \times 10^{-4}$.

Table 5 summarizes the exogenous trade-structure inputs, which anchor the tariff component of the calibration. Table 6 compares the model-implied moments to their empirical counterparts. The model matches the average Section 301 tariff as well as the extensive and intensive margins of lobbying. The approval rate among non-lobbyists is fitted less tightly. Because each sector has only two parameters to fit four moments, the model cannot in general match all moments exactly. In practice, the approval rate among non-lobbyists, A_{0s}, tends to exhibit the largest residual, reflecting the fact that it depends on the entire lower portion of the type distribution induced by the cutoff rule.

Table 5: Pre-Determined Trade Structure (Exogenous Inputs to Calibration)

Measure	Value
Chinese import share	0.067
Rest-of-world share	0.171
Column 1 (MFN) tariff	0.032
Column 2 tariff	0.380

Table 6: Fitted Moments: Lobbying and Approval Outcomes

Measure	Data Mean	Model Mean
Section 301 tariff (additional)	0.204	0.207
Lobbying share	0.218	0.207
Lobbyist approval rate	0.248	0.231
Non-lobbyist approval rate	0.094	0.031

To assess the informativeness of signals, I compare the calibrated noise level to the dispersion of θ across goods within each sector. Across sectors, the median within-sector standard deviation is 1.7×10^{-3}, implying that the calibrated signal noise is several orders of magnitude smaller than underlying heterogeneity. Even in the least dispersed sector, the noise level

remains below one percent of the natural variation in θ. Lobbying therefore transmits highly precise information about the welfare consequences of tariff relief.

To interpret the magnitude of c^*, consider a representative sector with \$100 million in total expenditure (including domestic production and imports), of which roughly \$6 million originates from China and faces an additional tariff burden of about \$1 million. For interpretive simplicity, I abstract from endogenous responses within the sector, such as tariff pass-through, changes in sourcing shares, and the resulting adjustments in the distribution of welfare impacts $\{\theta_{sg}\}$. These forces matter for the full equilibrium analysis in Section 6, but are second-order for the narrower task of translating the calibrated cost parameter into a dollar value. Under this benchmark, a lobbying cost of $c^* = 2 \times 10^{-4}$ corresponds to an outlay of roughly \$20,000. This level is consistent with the fixed-cost nature of filing, legal assistance, and documentation. Aggregating across sectors, the model implies a total lobbying cost of \$0.037 billion, comparable in magnitude to aggregate expenditures recorded in lobbying disclosures for Section 301 issues.

5.6 Discussion

The calibrated model reproduces the overall magnitude of lobbying and approval behavior under Section 301. A central conclusion that emerges from the calibration is that the informational channel was not constrained by noisy signals. The estimated noise level σ_ε^* is small relative to the natural dispersion of welfare impacts within sectors, indicating that once firms chose to lobby, their submissions conveyed highly informative signals about the consequences of tariff relief. This pattern is consistent with the institutional setting: exclusion petitions were prepared by professional legal counsel, supported by detailed documentation, and evaluated by a technically trained bureaucracy. When information was supplied, it was clear.

The binding friction instead arises from the fixed cost of participation. The calibrated cost c^* implies that only firms with sufficiently large expected gains from an exclusion found it worthwhile to prepare and submit petitions. In equilibrium, lobbying therefore operates as a selective communication channel: signals are highly precise, but only a subset of firms incurs the cost to transmit them. The political shock η^* introduces some randomness into approval decisions, but is modest relative to underlying heterogeneity, implying that the agency's decisions were reasonably predictable once information was available.

Taken together, these findings suggest that the Section 301 exclusion process functioned primarily as a mechanism for aggregating high-quality information under a politically weighted

objective. The informational friction did not lie in the government's ability to interpret signals, but in the cost borne by firms to generate and submit them. This distinction is central for the counterfactual analysis that follows, where the welfare consequences of removing lobbying arise not because lobbying is noisy, but because the government loses access to precise information supplied by participating firms.

6 Counterfactual Analysis

This section quantifies how informational lobbying affects equilibrium outcomes. In the calibrated economy, the government sets a lower average Section 301 tariff when it anticipates lobbying than when it expects no firm-level information, although this difference need not hold sector by sector. I begin by comparing the aggregate welfare in the calibrated economy to a world where lobbying is absent. I then decompose these effects using alternative counterfactual scenarios to isolate the role of ex-ante tariff choice versus ex-post targeting. All comparisons are expressed relative to the cooperative MFN benchmark.[5]

6.1 Baseline Welfare Comparison

I first compare the calibrated equilibrium ("Equilibrium with Lobbying") to a counterfactual "No Lobbying" equilibrium. In the latter, the government sets the tariff and makes approval decisions based solely on priors, knowing that no firm-level information will be available.

Table 7 reports the aggregate results. The equilibrium with lobbying generates a total welfare cost of $2.12 billion, compared with $2.59 billion when lobbying is absent. Informational lobbying therefore reduces the welfare loss of the trade war by roughly 18 percent. Importantly, because political weights are identical across scenarios, this gain arises entirely from the informational environment and the induced participation decisions, not from changes in political preferences.

Table 7: Aggregate Welfare Comparison: With vs. Without Lobbying (Relative to MFN)

Scenario	Total Welfare Cost ($bn)	Ratio (to No Lobbying)
Equilibrium without Lobbying	2.593	1.00
Equilibrium with Lobbying	2.118	0.82

[5]Throughout this section, positive values denote welfare *losses* relative to MFN.

6.2 Mechanisms and Decomposition

To understand the sources of this welfare gain, I analyze how the components of the government's objective change and distinguish between two distinct mechanisms:

1. **Tariff Choice:** The government may choose a different tariff rate τ_s when it anticipates the ability of firms to lobby.
2. **Implementation:** Conditional on a tariff, lobbying allows the government to screen petitions and approve those with the highest welfare (or political) value.

6.2.1 Counterfactual Scenarios

To isolate these mechanisms, I introduce two intermediate counterfactuals that vary the timing of the lobbying shutdown, holding structural parameters fixed. The relationship between these scenarios is illustrated in Figure 1.

1. **CF1: Surprise Shutdown.** The government sets the tariff τ^* anticipating the baseline equilibrium with lobbying, but the lobbying channel is unexpectedly shut down in the implementation stage. This scenario isolates the value of *targeting*, holding the tariff rule fixed at the lower level.
2. **CF2: Surprise Lobbying.** The government sets the tariff τ^{NL} anticipating no lobbying, but the lobbying channel is unexpectedly opened in the implementation stage. This scenario isolates the value of introducing information when the tariff rule is fixed at the higher level.

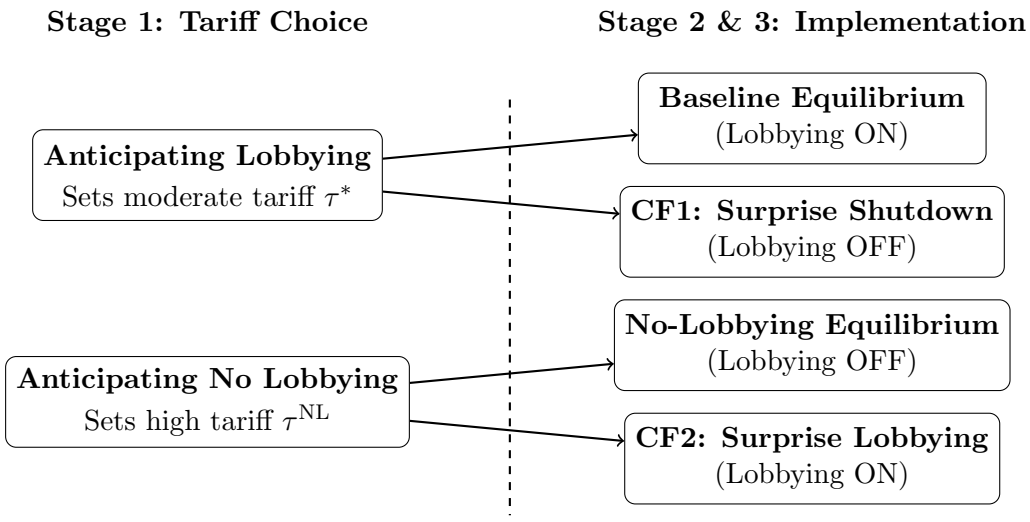

Figure 1: Structure of Counterfactual Scenarios

6.2.2 Results and Interpretation

Table 8 reports the optimal tariffs and approval rates in each scenario. A key finding is that the government sets a significantly lower tariff when it anticipates lobbying (20.7%) compared to when it does not (24.3%). This moderation reflects the fact that the government internalizes the cost of the lobbying process: a higher tariff induces more firms to incur participation costs, which dampens the government's incentive to raise duties.

Table 8: Policy Outcomes Across Scenarios

Scenario	Avg. Tariff (additional) τ_{301}	Avg. Approval Prob. (%)
Equilibrium with Lobbying	0.207	7.88
Equilibrium without Lobbying	0.243	4.83
CF1: Surprise Shutdown	0.207	5.31
CF2: Surprise Lobbying	0.243	7.45

Table 9 decomposes the welfare changes. The results clarify that the aggregate gains are driven primarily by the tariff moderation channel.

Table 9: Welfare Decomposition Across Counterfactuals (Relative to MFN)

	Components of Change ($bn)				**Total**
Scenario	Prod. Surplus	Downstream Profit	Tariff Rev.	Lobby Cost	Welfare Cost ($bn)
Equil. with Lobbying	−0.46	7.99	−5.45	0.04	2.12
Equil. no Lobbying	−0.49	9.24	−6.16	0.00	2.59
CF1: Surprise Shutdown	−0.47	8.37	−5.73	0.00	2.17
CF2: Surprise Lobbying	−0.48	9.04	−6.01	0.05	2.60

6.3 Interpretation

The counterfactuals separate two channels: (a) the impact of information when the *tariff rule* is held fixed, and (b) the impact of the *tariff rule itself* when chosen under different information expectations.

(i) Fixed–policy comparisons. Holding the tariff rule fixed but turning lobbying off (CF1) raises the total welfare loss only slightly, from $2.12 to $2.17 billion. Turning lobbying on after the no-lobbying tariff rule has been chosen (CF2) likewise generates only a small change, from $2.59 to $2.60 billion. These comparisons show that participation, signals, and lobbying costs have largely redistributive effects under a given tariff rule; they are not a first-order source of aggregate welfare differences.

(ii) Policy-rule effect. The dominant quantitative effect comes from the choice of the tariff rule. As shown in Table 8, the government sets a substantially lower tariff when it anticipates lobbying (0.207 versus 0.243). This moderation reflects the mechanism that a higher tariff induces more firms to incur the fixed cost of petitioning, which the policymaker internalizes when choosing the rule.

Information then improves ex-post implementation. With lobbying, the approval rate is 7.9 percent rather than 4.8 percent, reflecting more precise identification of high-harm products. The combination of a more moderate initial tariff and more targeted relief accounts for most of the $470 million difference in aggregate welfare costs between the two equilibria.

(iii) Composition. Differences in downstream profits explain most of the welfare gap. The no-lobbying equilibrium combines the worst of both worlds: a higher, blunt tariff that depresses downstream profits and an inability to screen, which leaves high-harm products unrelieved. The baseline equilibrium mitigates both distortions.

(iv) Summary. The welfare value of informational lobbying does not come from its direct, redistributive effects under a fixed tariff rule, but from the existence of the channel itself. Its availability disciplines the policymaker to choose a more moderate tariff ex ante and enables more efficient, better-targeted exclusions ex post.

7 Conclusion

This paper develops and quantifies a model of *informational lobbying* in which firms lobby not to purchase political access but to transmit costly information to a bureaucracy operating under political constraints. Using the U.S.–China Section 301 tariff exclusion process as a case study, the analysis shows that lobbying can improve the quality of policy implementation even when the policymaker's objective is explicitly political.

The model embeds a noisy signaling game inside a politically weighted objective framework. Firms observe the product-level impact of tariff relief under the government's own objective and decide whether to incur a cost to reveal that information. The government, in turn, interprets lobbying behavior through Bayesian updating and decides on exclusions subject to political shocks. When signals are precise but participation is costly, lobbying endogenously selects firms whose information is most valuable to the policymaker, generating targeted relief without altering the political weights that shape the underlying tariff.

Calibrated to the universe of 50,000 petitions, the model matches key patterns of lobbying

participation and approval rates across sectors. The estimated lobbying cost is modest but sufficient to deter many firms from engaging, while the inferred signal noise is near zero and political shocks are modest. These estimates indicate that most uncertainty in the exclusion process arose from limited participation rather than from arbitrariness.

The quantitative results reveal a conceptually simple but powerful mechanism. Informational lobbying lowers the average Section 301 tariff from 0.243 to 0.207 because the policymaker internalizes that a higher tariff induces more firms to incur the fixed cost of petitioning. Once the tariff is chosen, lobbying provides precise signals that raise the approval rate from 4.8 to 7.9 percent, enabling more targeted relief for the most adversely affected products. Together, tariff moderation and better targeting reduce the aggregate welfare cost from \$2.59 to \$2.12 billion—an 18 percent improvement relative to a world without lobbying.

The counterfactuals clarify that these gains arise from the *existence* of the informational channel rather than from lobbying activity per se. Turning lobbying on or off after the tariff rule has been set affects welfare only mildly: participation costs are second-order, and information is valuable only when the policymaker anticipates it. The main effect comes from how the policymaker chooses the tariff in anticipation of the informational environment. This interaction between discretion and information generalizes to other regulatory settings in which agencies implement politically directed policies using stakeholder-provided information.

Implications and Directions for Future Research

The findings illustrate how bureaucratic discretion can, under certain conditions, harness rather than amplify political distortion. In trade policy, this helps explain why administrative exemption procedures such as product-level rulings, waiver systems, and licensing reviews persist even in highly politicized contexts: they enable policymakers to incorporate decentralized information while retaining political control.

Several extensions follow naturally. First, firms in the model act independently; introducing strategic interaction could illuminate how information aggregates within sectors and how free-riding affects participation. Second, the calibration treats political weights as reduced-form objects; linking them to industry characteristics or electoral geography could connect informational lobbying to the broader determinants of protection. Third, allowing for heterogeneous participation costs across firms would generate richer sorting patterns and could reveal which types of firms are most constrained by the fixed-cost nature of petitioning.

More broadly, the framework offers a template for quantifying the informational content

of lobbying in other administrative environments. Policy regimes such as Buy America waivers, export controls, foreign-investment screening, and environmental permitting share the same institutional logic: policymakers must implement politically directed rules using private information held by firms. Applying similar signaling structures to these settings could deliver a unified view of how information flows shape the effectiveness and legitimacy of modern industrial and trade policy.

References

Bown, C. P. and Kolb, M. (2021). Trump's trade war timeline: An up-to-date guide. Technical report, Peterson Institute for International Economics. Available at `https://www.piie.com/blogs/trade-investment-policy-watch/trumps-trade-war-timeline-date-guide`.

Chor, D., Grant, M., and Li, B. (2025). Exclusions for sale? tariff exclusions in the u.s.–china trade war. Working Paper.

De Figueiredo, J. M. and Silverman, B. S. (2008). Academic earmarks and the returns to lobbying. *Journal of Law and Economics*, 51(4):851–876.

Efron, B. (2011). Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614.

Gawande, K., Krishna, P., and Olarreaga, M. (2012). Lobbying competition over trade policy. *International Economic Review*, 53(1):115–132.

Goldberg, P. K. and Maggi, G. (1999). Protection for sale: An empirical investigation. *American Economic Review*, 89(5):1135–1155.

Government Accountability Office (2021). U.s.–china trade: Opportunities to improve ustr's tariff exclusion process. Technical Report GAO-21-506, U.S. Government Accountability Office, Washington, DC.

Government Accountability Office (2025). Lobbying disclosure: Observations on compliance with requirements. Technical Report GAO-25-107523, U.S. Government Accountability Office, Washington, DC.

Grossman, G. M. and Helpman, E. (1994). Protection for sale. *American Economic Review*, 84(4):833–850.

Kartik, N. (2009). Strategic communication with lying costs. *Review of Economic Studies*, 76(4):1359–1395.

Kim, I. S. (2018). Lobbyview: Firm-level lobbying & congressional bills database. Working Paper, Massachusetts Institute of Technology.

Ludema, R., Mayda, A. M., and Mishra, P. (2010). Protection for free? the political economy of u.s. tariff suspensions. IMF Working Paper 10/211, International Monetary Fund.

Matthews, S. A. and Mirman, L. J. (1983). Equilibrium limit pricing: The effects of private information and stochastic demand. *Econometrica*, 51(4):981–996.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380–391.

Potters, J. and Van Winden, F. (1992). Lobbying and asymmetric information. *Public Choice*, 74(3):269–292.

Robbins, H. (1956). An empirical bayes approach to statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 157–163, Berkeley, CA. University of California Press.

Soderbery, A. (2015). Estimating import supply and demand elasticities: Analysis and implications. *American Economic Journal: Microeconomics*, 7(3):188–222.

Appendix A Microfoundation of the Objective Component

This appendix provides the full structural environment that underlies the expressions used in Section 4.8. It presents complete derivations of upstream supply, sourcing shares, downstream input demand, tariff revenue, and the government's objective $\Omega_{sg}(\tau_s)$. The main text reports the components needed for the signaling model—at times including intermediate steps such as the downstream first-order condition—while the full set of algebraic details appears here for completeness.

Appendix A.1 Final Consumers and Aggregate Demand

Each good (s, g) yields a final output C_{sg} that enters the utility function of a representative consumer together with a freely traded numeraire X_0:

$$U = X_0 \; + \; \left[\int_s \int_{g \in s} \phi_{sg}^{\frac{1}{\rho}} C_{sg}^{\frac{\rho-1}{\rho}} \, dg \, ds \right]^{\frac{\rho}{\rho-1}}, \qquad \rho > 1. \tag{2}$$

The final goods C_{sg} are traded on world markets under perfect competition, so their prices are fixed from the perspective of the domestic economy. Quasi-linearity implies that total expenditure on the CES composite is constant, and with fixed world prices the quantities demanded of final goods C_{sg} are invariant to a tariff exclusion applied to upstream intermediate inputs. This demand structure, together with perfect competition in final-good markets, implies that welfare is additively separable across products. Because upstream rents, downstream profits, and tariff revenue are likewise defined at the product level, the government's objective inherits this product-level separability.

Consequently, tariff exclusions do not affect welfare through consumer demand or final-good prices. All policy-relevant effects operate through the production side: changes in intermediate-input costs faced by downstream firms and the implied changes in upstream rents, downstream surplus, and tariff revenue. This is why the microfoundation of θ_{sg} in the main text focuses exclusively on the production block.

Appendix A.2 Downstream Technology and Input Demand

Each product (s, g) is used by a competitive downstream firm that combines a composite intermediate input X_{sg} into final output. Technology is given by

$$Y_{sg} = B_s \ln X_{sg}, \tag{3}$$

where $B_s > 0$ is a sector-specific parameter.

Composite Input and Price Index. The intermediate input X_{sg} aggregates domestic (h), Chinese (c), and rest-of-world (f) varieties through a CES aggregator with elasticity of substitution $\sigma_s > 1$:

$$X_{sg}^{\frac{\sigma_s-1}{\sigma_s}} = x_{sgh}^{\frac{\sigma_s-1}{\sigma_s}} + x_{sgc}^{\frac{\sigma_s-1}{\sigma_s}} + x_{sgf}^{\frac{\sigma_s-1}{\sigma_s}}, \tag{4}$$

which implies a unit cost (price index)

$$P_{sg} = \left(p_{sgh}^{1-\sigma_s} + \left[(1+\tau_{sgc})p_{sgc}\right]^{1-\sigma_s} + \left[(1+\tau_{sgf})p_{sgf}\right]^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}}. \tag{5}$$

Note that the price index includes the tariffs τ_{sgc} and τ_{sgf} applied to Chinese and rest-of-world varieties, respectively.

Downstream Producer Optimization. Taking the final-good price as 1, the downstream firm chooses X_{sg} to maximize

$$\Pi_{sg}^{\text{down}} = B_s \ln X_{sg} - P_{sg} X_{sg}. \tag{6}$$

The first-order condition is

$$\frac{B_s}{X_{sg}} = P_{sg}.$$

Solving for X_{sg} gives

$$P_{sg} X_{sg} = B_s. \tag{7}$$

Constant Intermediate-Input Expenditure. Define total intermediate-input expenditure as

$$M_{sg} \equiv P_{sg} X_{sg}.$$

Equation (7) implies

$$M_{sg} = B_s,$$

so tariff or price changes affect only the composition of sourcing, not the total scale of intermediate-input use.

Reduced-Form Downstream Profits. Substitute the optimal X_{sg} into (6):

$$\Pi_{sg}^{\text{down}} = B_s \ln\left(\frac{B_s}{P_{sg}}\right) - B_s.$$

Thus up to a constant term,

$$\Pi_{sg}^{\text{down}} = -B_s \ln P_{sg} + \text{constant}, \tag{8}$$

which is the expression used in the main text.

Appendix A.3 Upstream Production and Supply

Each origin $i \in \{h, c, f\}$ supplies an intermediate-input variety specific to product (s, g) using labor L_{sgi} and a specific factor K_{sgi}. Output is produced with a standard specific-factor technology:

$$q_{sgi} = A_{sgi} L_{sgi}^{\frac{1}{1+\omega_s}} K_{sgi}^{\frac{\omega_s}{1+\omega_s}}, \qquad \omega_s > 0, \tag{9}$$

where ω_s is the inverse supply elasticity at the sector level. Producer revenue is $p_{sgi}q_{sgi}$, while downstream expenditure on inputs from origin i equals $(1 + \tau_{sgi})\, p_{sgi}q_{sgi}$. Only the producer-revenue component $p_{sgi}q_{sgi}$ is relevant for the distribution of specific-factor rents. This distinction is important because tariff revenue accrues to the government, not to the upstream producer. Firms are perfectly competitive and take factor prices as given, with the wage normalized to 1 and the specific factor earning a residual return.

Optimal Labor Demand. Given an output price p_{sgi}, the firm solves

$$\max_{L_{sgi}} p_{sgi}q_{sgi} - L_{sgi}.$$

Using (9), optimal labor demand is

$$L_{sgi} = \left(\frac{p_{sgi}A_{sgi}}{1+\omega_s}\right)^{\frac{1+\omega_s}{\omega_s}} K_{sgi}. \tag{10}$$

Supply Function. Substituting (10) into (9) yields yields a constant-elasticity supply curve with elasticity $1/\omega_s$:

$$q_{sgi} = a_{sgi}^{-1/\omega_s} p_{sgi}^{1/\omega_s}, \qquad a_{sgi} \equiv (1+\omega_s) K_{sgi}^{-\omega_s} A_{sgi}^{-(1+\omega_s)}. \tag{11}$$

Thus supply has elasticity $1/\omega_s$.

Appendix A.4 Tariffs, Expenditure, and Sourcing Shares

Let τ_{sgi} denote the ad valorem tariff applied to origin i. The downstream firm pays $(1+\tau_{sgi})p_{sgi}$ per unit of the intermediate. Its total expenditure on inputs from that origin is therefore

$$R_{sgi}^{\text{exp}} = (1+\tau_{sgi})\, p_{sgi} q_{sgi}. \tag{12}$$

For clarity, we use the superscript "exp" here to emphasize that this is *expenditure* by the downstream firm including the tariff.

Downstream firms minimize the cost of acquiring the composite input X_{sg}:

$$X_{sg}^{\frac{\sigma_s-1}{\sigma_s}} = x_{sgh}^{\frac{\sigma_s-1}{\sigma_s}} + x_{sgc}^{\frac{\sigma_s-1}{\sigma_s}} + x_{sgf}^{\frac{\sigma_s-1}{\sigma_s}},$$

which implies a cost $P_{sg}X_{sg}$, where P_{sg} is the CES price index.

The sourcing share of origin i in product (s,g) is then

$$s_{sgi} = \frac{R_{sgi}^{\text{exp}}}{P_{sg}X_{sg}} = (1+\tau_{sgi})\frac{a_{sgi}^{-1/\omega_s}}{P_{sg}X_{sg}} p_{sgi}^{(1+\omega_s)/\omega_s}. \tag{13}$$

Appendix A.5 Upstream Rents

Because the specific factor earns the residual return, upstream operating profits are a constant fraction of producer revenue. Producer revenue is $p_{sgi}q_{sgi}$, while downstream expenditure on origin i is $(1+\tau_{sgi})p_{sgi}q_{sgi}$. Using the standard specific-factor result,

$$\Pi_{sgi}^{\text{up}} = \frac{\omega_s}{1+\omega_s}\, p_{sgi} q_{sgi} = \frac{\omega_s}{1+\omega_s} \frac{R_{sgi}^{\text{exp}}}{1+\tau_{sgi}}. \tag{14}$$

Recall from (12) that R_{sgi}^{exp} denotes downstream expenditure inclusive of the tariff. For the purposes of the government's objective, it is convenient to redefine

$$R_{sgi} \equiv p_{sgi} q_{sgi}$$

as producer revenue net of the tariff. With this convention,

$$\Pi^{\text{up}}_{sgi} = \frac{\omega_s}{1+\omega_s} R_{sgi}, \tag{15}$$

which is the reduced-form expression reported in Section 4.8.

Appendix A.6 Tariff Revenue

Tariff revenue collected from origin i is:

$$TR_{sgi} = \tau_{sgi} p_{sgi} q_{sgi} = \frac{\tau_{sgi}}{1+\tau_{sgi}} R^{\text{exp}}_{sgi} = \tau_{sgi} R_{sgi}, \tag{16}$$

where the last equality again uses $R_{sgi} \equiv p_{sgi} q_{sgi}$. Summing over origins yields TR_{sg}.

Appendix A.7 Government Objective and Exclusion Types

Combining downstream profits, upstream rents, and tariff revenue gives the government's product-level objective in reduced form:

$$\Omega_{sg}(\tau_s) = \lambda_s \Pi^{\text{up}}_{sgh} + \beta_s \Pi^{\text{down}}_{sg} + TR_{sg} + \sum_{i \in \{c,f\}} \gamma_{s,i} \Pi^{\text{up}}_{sgi}. \tag{17}$$

The type $\theta_{sg}(\tau_s)$ used in the signaling game is the change in this objective from granting a tariff exclusion:

$$\theta_{sg}(\tau_s) = \Omega_{sg}(0) - \Omega_{sg}(\tau_s). \tag{18}$$

As noted in the main text, a separate appendix will provide the argument establishing the monotonicity of the firm's private benefit $B(\theta)$ in θ.