Environmental Disclosures in Global Supply Chains

Christian C. Opp Xingtan Zhang*

Abstract

Economies pursuing environmental goals such as mitigating global warming face the challenge that the relevant externalities largely originate abroad. Governments may respond by taxing imports based on foreign emissions, relying on disclosures that also underpin ESG ratings. We model strategic emissions disclosure regimes for foreign suppliers facing a domestic firm with market power, and characterize optimal tariffs. These differ from prices in carbon credit markets: conditioning on information asymmetries, they are skewed toward the externalities of each rating's greener types. Intervention can make disclosures more or less precise; under severe exposure, ratings become very coarse, excessively pooling brown producers.

Keywords: Carbon credit markets; ESG ratings; Global warming; Asymmetric information; Optimal regulation

*We thank Ehsan Azarmsa, Kostas Dellis (discussant), Vincent Glode, Deeksha Gupta (discussant), Lars Peter Hansen, Nicolas Inostroza, Jan Schneemeier (discussant), Anton Tsoy, Brian Waters, and seminar participants at CKGSB, INSEAD, University of Colorado Boulder, University of Rochester, European Finance Association, Financial Management & Accounting Research Conference, 2025 FTG Fall, Society of Economic Dynamics Meeting, and UT Dallas Fall Finance Conference for their comments. Opp is with University of Rochester and NBER and can be reached by email: opp@rochester.edu. Zhang is with Cheung Kong Graduate School of Business and can be reached by email: xingtanzhang@ckgsb.edu.cn.

1 Introduction

Some of the most consequential externalities affecting economies originate abroad. A leading example is global warming: economies pursuing climate goals (UNFCCC, 2015) face the challenge that emissions largely originate beyond their borders (see Peters et al., 2011; Climate Action Tracker, 2023, among others). At the same time, tariffs are again at the center of economic policy as instruments aimed at influencing foreign behavior. A new class of tariffs targets not just where goods are made, but also how: for instance, the European Union's Carbon Border Adjustment Mechanism (CBAM) charges importers of steel, aluminium, and other carbon-intensive inputs for the emissions embedded in their production abroad. Policies of this type confront a fundamental informational problem: the externalities they target arise beyond the taxing government's jurisdiction and measurement reach, and the parties best informed about them — the foreign producers themselves — are the source of disclosures.

This paper studies how a domestic government should design such policies when it must rely on what foreign suppliers disclose. We examine the design of the foreign disclosure regime — the rating categories — both by a foreign government maximizing the surplus of its firms and by an individual foreign firm. In our baseline model, the foreign government chooses how precisely its producers' externalities are disclosed: it selects the disclosure regime under which each privately informed supplier discloses to a domestic firm with market power. The domestic government sets tariffs that condition on these disclosures.

We obtain several key findings. First, optimal tariffs differ from prices in carbon credit markets: rather than charging a common price per disclosed ton, their charges are skewed toward compensating for the externalities of each rating's greener types, while also prohibiting trade for the brownest disclosures. Second, tariffs shape what gets disclosed, with households' exposure to the externalities determining the direction: when exposure is moderate, tariffs, perhaps paradoxically, can render foreign disclosures more precise; when exposure is severe, disclosures collapse to very coarse ratings. Third, delegating disclosure design to domestic firms does not restore informative ratings: when exposure is severe, both domestic and foreign firms prefer excessively coarse ratings that preserve private supply-chain surplus, even though the trade these ratings sustain destroys global surplus.

The setting we study is quantitatively important: Scope-3 emissions, those occurring in companies' value chains, constitute the majority of total emissions across most sectors (Huang et al., 2009; Bolton and Kacperczyk, 2021). Governments accordingly regulate through domestic firms' supply chains (see, e.g., Copeland and Taylor, 2003), as under CBAM, the EU's Corporate Sustainability Reporting Directive, California's SB 253, and the EU Deforestation Regulation,[1] and

[1] See Regulation (EU) 2023/956 (CBAM); Directive (EU) 2022/2464 (CSRD); Regulation (EU) 2023/1115

environmentally minded investors rely on ESG scores that are meaningful only if they reflect supply chain practices (see Azarmsa and Shapiro, 2025, on the market for such ratings). The mechanism we study applies whenever reducing emissions is privately costly and foreign producers retain information about their production technologies.

Our model captures two key features of global sourcing: foreign suppliers' private information about their production practices, and the substantial market power that large firms in developed economies wield over international suppliers (see, e.g., Antràs and Staiger, 2012). The externalities generated by foreign suppliers harm domestic households and are not sufficiently internalized by firms. Externality-generating production is privately cheap but socially costly. Information about it therefore plays a dual role: a disclosure of production practices informs the domestic firm, the bargaining partner, about how cheaply its supplier can produce, thereby shaping the terms of trade, and it simultaneously informs the domestic country about the externality its households will bear.[2]

This duality, in turn, generates a central force in our analysis. Because browner production is privately cheaper, a tariff that lowers the price the domestic firm offers does not eliminate the brownest suppliers first: instead it rations the greener, higher-cost suppliers. A government attempting to screen out high-emissions trade through tariffs therefore faces adverse selection. Moreover, the foreign side anticipates how the domestic government will use disclosed information when choosing its disclosure regime.

Absent government intervention, disclosures play a purely private role: by reducing the information asymmetry between supply chain partners, they alleviate the privately inefficient rationing of trade, raising the supply chain partners' joint surplus. The foreign side then finds it optimal to commit to a coarse disclosure regime that is informative enough to prevent inefficient trade breakdowns, yet coarse enough to preserve information rents. The additional trade that disclosures enable, however, can destroy global surplus when externalities are accounted for.

Once the domestic government intervenes to protect its households, the foreign disclosure regime itself responds: tariffs and the disclosed information they rely on interact. A central commitment problem underlying our analysis is that the domestic government cannot commit to punishing coarse disclosures with punitive tariffs it would not want to impose once such ratings are disclosed. In other words, such threats are not credible. We find that the impact on equilibrium informativeness crucially depends on the strength of domestic households' exposure to the externalities, that is, the degree to which the domestic economy is harmed by brown production abroad.

(EUDR); California SB 253 (2023). We return to these initiatives below to compare them with the optimal policies our analysis identifies.

[2]Beyond the environmental applications we focus on, the same informational structure arises whenever domestic parties bear disutility from foreign production practices themselves and border measures condition on those practices. Examples include import restrictions tied to forced labor, as under the long-standing U.S. import ban of Section 307 of the Tariff Act of 1930 and recently introduced U.S. tariffs conditioning on forced-labor practices abroad, as well as the supply-chain due-diligence obligations under the EU's Corporate Sustainability Due Diligence Directive.

When exposure is moderate, optimal tariffs shrink the private gains from trade to a limited extent and cause the foreign side to choose finer disclosures that protect the remaining private surplus from trade breakdowns. As such, tariffs can improve the informativeness of foreign disclosures.

When the domestic economy is strongly exposed to foreign-produced externalities, the adverse selection problem faced by the domestic government is particularly severe, since browner production is more harmful to its households. Then, conditional on a foreign disclosure (i.e., a rating), the domestic government's best response is one of two extremes: either to admit all trade under the disclosed rating or to block it entirely. Anticipating this drastic behavior, the foreign side has an incentive to create disclosures that pool the largest possible set of relatively green producers under one broad rating category with which the domestic side is still willing to trade. The resulting rating is excessively broad in the sense that the trade it sustains destroys global surplus.

This result is, moreover, not an artifact of the foreign side controlling disclosure precision. We show that a domestic firm designing the ratings itself would choose the same excessive pooling, that is, both the domestic and the foreign firm benefit from this disclosure regime. Our analysis characterizes these results under a range of distributional assumptions, showing that finely differentiated ratings for greener producers can coexist with overly broad rating categories that destroy global surplus.

The tariffs that implement the domestic country's optimal policy take the residual information asymmetry— the uncertainty about production practices that remains conditional on a disclosure — as given: conditional on a disclosure, the tariffs affect the terms of trade, which determine both how much trade occurs and how green the trading pool is. Optimal conditional tariffs set each disclosure category's charge according to the externalities of the marginal — relatively green — producers, as greener production is costlier and greener types are the first to be rationed.

These prescriptions contrast with prevailing policy designs. The EU Emissions Trading System features one market price for carbon allowances across sectors, and CBAM extends this uniform per-ton price to imports by charging the allowance price on reported embedded emissions (or on default values when producer-specific data are not supplied). Once foreign producers retain private information and domestic importers hold market power, a uniform price is no longer optimal: the effective charge per disclosed ton should vary across disclosure categories and sectors, skewed toward the greener end of each disclosed range rather than anchored to reported or default values.

The principle behind this skew is that tariffs change outcomes by altering which suppliers participate in trade: the optimal charge is determined by the suppliers whose participation changes at the margin — within any rating category, its relatively green, higher-cost producers — rather than by average reported emissions. The contrast extends to the choice of instrument itself. In our environment, the optimal quantity of high-emissions imports is type-contingent and unknown to the regulator: a disclosure-contingent price induces privately informed producers to self-select in ways

that an unconditional quantity restriction on a supply relationship cannot replicate, an asymmetric-information counterpart to the classical price-versus-quantity comparison of Weitzman (1974).

The theoretical novelty of our analysis lies in introducing externalities into strategic disclosure — and in the resulting feedback between regulation and information. Voluntary disclosure to a counterparty with market power serves the trading parties: it trades off information rents against the surplus lost when trade breaks down. In our setting, the disclosed production practice does not only save the supplier costs; it also harms parties outside the trade — domestic households — and this externality is what motivates regulation in the first place. Disclosures thereby acquire a second role: the same information that shapes the bilateral terms of trade informs the government's tariffs, and the regulator's anticipated use of it changes how much information the foreign side produces. We characterize when intervention renders disclosures finer, when it induces coarse pooling, and why regulation that is optimal given the available information nevertheless sustains globally inefficient trade.

Related literature. Our paper contributes most directly to the theoretical literature on ESG and environmental disclosure. Lyon and Maxwell (2011) model greenwashing: the selective disclosure of favorable environmental information, disciplined by the threat of audits. More recent work examines how ESG disclosure interacts with financial markets and regulation: Chen and Schneemeier (2022) study the feedback between responsible investors' trading and firms' green decisions, Aghamolla and An (2023) compare mandatory and voluntary ESG disclosure regimes and their real effects, Xue (2023) analyzes how market forces shape the efficiency consequences of ESG disclosure, Gupta and Starmans (2024) study the dynamic design of green disclosure requirements, and Azarmsa and Shapiro (2025) model the supply of ESG ratings by competing information intermediaries. Our focus differs along two dimensions that make our analysis complementary to these papers. First, disclosure takes place in an international supply chain: the domestic regulator has no authority over the disclosing party and must instead operate through tariffs on domestic firms, which hold buyer-side market power. Second, the granularity of disclosures — the coarseness of equilibrium ratings — is itself the object of interest, chosen by the rated side and responding to the anticipated regulatory use of the information.

More broadly, our paper builds on the general theoretical literature examining voluntary disclosure, which abstracts from externalities affecting third parties and the associated scope for government interventions. Grossman (1981b) and Milgrom (1981) find that when disclosures are ex post verifiable, it is optimal for an agent to fully reveal his private information: since the disclosing agents in these settings do not face counterparties with market power, only beliefs about the mean payoff matter, and any withheld information is interpreted unfavorably, so that disclosure "unrav-

els" (see also Grossman and Hart (1980); Milgrom and Roberts (1986)).[3] As in Glode et al. (2018), the disclosing party in our setting instead faces a counterparty with market power, which screens based on the full conditional distribution of types that a disclosure induces; partial rather than full disclosure is then optimal. The presence of externalities distinguishes our environment: contrary to the efficiency results in Glode et al. (2018), disclosures generally do not ensure trade that maximizes total surplus, even when the domestic government employs an optimal mechanism.[4] Our focus on market power also relates our paper to Gal-Or (1985), where oligopolists can commit ex ante to sharing noisy signals about uncertain demand, although no information sharing arises in equilibrium there.

Our analysis also relates to the theoretical literature on environmental policy instruments, beginning with Weitzman (1974)'s classic comparison of price and quantity instruments. In our cross-border context, the price instrument takes the form of a tariff on transactions with foreign suppliers. Golosov et al. (2014) derive optimal carbon taxes in a dynamic general equilibrium model, and Döttling and Rola-Janicka (2025) provide a rationale for quantity instruments based on financial constraints that is distinct from, and complementary to, the informational forces favoring price instruments in our setting. Classic assessments and surveys of climate policy include Nordhaus (1991), Poterba (1991), and Pindyck (2013).

Finally, a complementary *empirical* literature supports the premises of our setting. On the regulation of carbon costs, Elliott et al. (2010) quantify emissions leakage and the mitigating role of border adjustments, Fowlie et al. (2016) estimate leakage through imports under market-based emissions regulation, and Fabra and Reguant (2014) document that emissions costs are almost fully passed through to prices — pass-through likewise being the margin through which tariffs shift the composition of trading supplier types in our model. On disclosure behavior, Kim and Lyon (2011) document strategic greenhouse-gas disclosure under voluntary regimes, Darendeli et al. (2022) show that corporate buyers use suppliers' CSR information in contracting decisions, and Bonetti et al. (2024) find that customers' climate-information demands predict suppliers' disclosures and subsequent decarbonization. Christensen et al. (2021) review the economics of mandatory sustainability reporting (see also Wagenhofer, 2024; Hodder and Schipper, 2025). This evidence supports

[3] Verrecchia (1983) modifies a setting akin to Grossman (1981b) and Milgrom (1981) by adding disclosure costs, whereas Fishman and Hagerty (1990) assume that a subset of private information cannot be disclosed. In both cases, maximal disclosure may not be optimal for the informed party. Admati and Pfleiderer (2000), however, show that a firm may pick a socially optimal disclosure plan despite disclosure costs if that firm is a monopolist that captures all gains to trade. Matthews and Postlewaite (1985), Okuno-Fujiwara et al. (1990), Fishman and Hagerty (2003), Shin (2003), Acharya et al. (2011), and Guttman et al. (2014) show, in different environments, that full disclosure becomes suboptimal once there is uncertainty about the existence of private information or its content.

[4] Ali et al. (2023) echo the results of Glode et al. (2018) in the context of voluntary disclosures in interactions with a monopolist. Monopoly pricing is also studied by Bergemann et al. (2015) who analyze how exogenously providing monopolists with additional information for price discrimination affects total surplus and its allocation. Inostroza and Tsoy (2025) jointly study security design and information design; in their setting, trade always occurs in equilibrium.

two key premises of our model: regulators rely on information produced by foreign firms, and their authority over those firms is limited.

2 Model

Two firms from distinct countries interact in a global supply chain relationship. A domestic firm (D) sources an input from a foreign firm (F) and provides the final product to domestic households (H). The foreign firm's production of the input involves technologies that generate environmental externalities e. The foreign firm is privately informed about the value e, which captures the notion that in practice, a producer is best informed about the specific technologies and practices it is employing. In contrast, outside parties have imprecise prior information about the amount of negative externalities created by the foreign company's manufacturing activities. Outsiders' beliefs are characterized by the cumulative distribution function (CDF) $G(e)$. The function $G(e)$ is continuous and differentiable and the probability density function (PDF), denoted by $g(e)$, takes strictly positive values on the support $[\underline{e}, \overline{e}]$.

The foreign firm can be a cost-efficient producer, which motivates the domestic firm to potentially source its inputs from this firm. However, the foreign firm's production costs are not only affected by its efficiency, but also negatively related to the brownness of its production process: Pollutive production is cheaper from a private perspective. For instance, a brown production process may be cheaper as it economizes on abatement efforts. Formally, the foreign firm's production costs are given by $(c-e)$, where c represents production costs unrelated to externalities, and e represents the cost savings from externality-generating production: one unit of e is the amount of environmental harm whose creation saves the foreign firm one unit of production cost.

Two key features of the global supply chain we model are (1) the foreign firm's private information about its environmental standards, and (2) the domestic firm's market power. Specifically, the domestic firm chooses a price p that it will offer to the foreign firm in a take-it-or-leave-it offer.[5] If the offer is rejected by the foreign firm, it is not used to source inputs and all parties obtain a payoff that is normalized to 0.[6] In contrast, if the foreign firm accepts the offer, it produces and

[5]The take-it-or-leave-it offer captures the bilateral nature of cross-border supply relationships and the buyer-side market power that large multinational firms exercise within them. Cross-border inputs are typically specialized and traded within long-term relationships rather than on spot markets (see, e.g., Antràs and Helpman, 2004; Nunn, 2007), and switching costs arising from qualification, certification, and relationship-specific investments sustain such sourcing arrangements as relational contracts (Macchiavello, 2022). What our analysis requires is that the domestic firm holds price-setting power vis-à-vis a privately informed supplier; the take-it-or-leave-it offer is a standard tractable representation of this friction.

[6]This normalized payoff could reflect in reduced form the possibility that the domestic firm instead relies on an outside option, such as a domestic supplier with known environmental standards.

delivers the input to the domestic firm, yielding it a payoff:

$$u_F = p - \underbrace{(c-e)}_{\text{costs}}, \tag{1}$$

where the first term represents the proceeds from the transaction and the terms in brackets capture the production costs accounting for the cost savings associated with brown technologies that generate high environmental externalities.

Absent government interventions, the domestic firm is not directly impacted by the negative externalities e. This assumption captures the more general notion that the domestic firm does not fully internalize the externalities associated with its sourcing of inputs. In contrast, the domestic household sector is negatively affected by the environmental externalities. In particular, given that production occurs and generates externalities e, domestic households suffer a disutility of $-\lambda e$, where $\lambda > 0$ captures households' exposure to environmental damages. Formally, while e measures the foreign firm's cost savings, λe measures the externality costs incurred by domestic households. The parameter λ is thus the marginal damage households bear per unit of the foreign firm's cost savings. We will refer to λ as households' "exposure" to foreign production practices, reflecting how severely households are affected by these practices.

Production is beneficial to the domestic sector. We let v represent the total value the domestic sector obtains. Among them, δ is the fraction that the domestic firm internalizes, and $(1-\delta)$ is the domestic households' share. Overall, conditional on production and trade occurring, the domestic firm's payoff is given by:

$$u_D = \delta v - p, \tag{2}$$

and the domestic households' payoff is given by:

$$u_H = (1-\delta)v - \lambda e. \tag{3}$$

The domestic surplus conditional on trade occurring is then given by:

$$u_D + u_H = v - \lambda e - p, \tag{4}$$

and the global surplus conditional on trade occurring takes the form:

$$u_D + u_H + u_F = v - c + (1-\lambda)e. \tag{5}$$

Environmental disclosures. Central to our analysis is the possibility of environmental disclosures

by the foreign firm. Because the brownness of production both lowers the foreign firm's private costs and harms domestic households, information about e inherits a corresponding dual role: a disclosure informs the domestic firm about its trading partner's costs $(c-e)$ — and thus about the prices the foreign firm is willing to accept — while informing the domestic government about the disutility λe that trade imposes on households. The first margin shapes bilateral bargaining; the second shapes optimal tariffs. We suppose that the disclosure policy is chosen by the foreign government. The foreign regulator maximizes the foreign surplus knowing that the cross-section of foreign firms is described by the distribution $G(e)$. So the disclosure policy is chosen based on the distribution of e, without knowing the specific realization e. This increases the tractability of our baseline analysis, as it eliminates the signaling concerns. In Section 4.3, we also consider the case where a disclosure plan is instead chosen by the foreign firm after it has learnt its type e (so called "interim" disclosure), and we show that our results are robust to this type of interim disclosure under reasonable equilibrium refinements.

Regarding the disclosure technology, we consider ex post verifiable disclosures. This assumption is motivated by the fact that it is in principle *technologically feasible* that a foreign firm or country would subject itself to environmental certification processes administered by international rating agencies. In the context of CO2 emissions, it is for example possible to allow foreign agencies to monitor such emissions. However, this technological feasibility still leaves the question whether a foreign firm or foreign country finds it privately optimal to commit to such verification. Under ex post verifiability, choosing *not* to commit to providing incremental information is always an option. Ex post verifiability is also a common assumption in the disclosure literature (see Verrecchia (2001); Milgrom (2008a); Beyer et al. (2010) for related surveys). We formally define ex post verifiability as follows.

Definition 1. *A signal whose realization s belongs to a set S is called "ex post verifiable" if it can be represented by a function $D:[\underline{e},\overline{e}]\to S$ such that $D^{-1}(s)\equiv\{e:D(e)=s\}\in\mathscr{B}([\underline{e},\overline{e}])\ \forall s\in S$, where $\mathscr{B}([\underline{e},\overline{e}])$ denotes the Borel algebra on $[\underline{e},\overline{e}]$.*

This definition ensures that the trade partner's conditional beliefs about e given any disclosure are well defined. In the literature, it is common to impose monotonicity restrictions with regard to the disclosure functions agents choose, since doing so significantly increases analytical tractability and reflects the ordered structure of the rating and reporting systems used in practice (see Glode, Opp, and Zhang, 2018; Orlov, Zryumov, and Skrzypacz, 2023).[7] We follow this approach as stated in the following assumption.

[7]Monotonicity is also imposed in the context of the security design literature, where it has similar justifications. See, e.g., Innes (1990), Nachman and Noe (1994), and DeMarzo and Duffie (1999).

Assumption 1. *Disclosure plans $D(e)$ are restricted to be monotone in e.*

In particular, a disclosure plan $D(e)$ can be represented as follows:

$$D(e) = \begin{cases} \underline{e}, & \text{for } e \in [\underline{e}, e_1), \\ e_1, & \text{for } e \in [e_1, e_2), \\ e_2, & \text{for } e \in [e_2, e_3), \\ \ldots & \\ e_n, & \text{for } e \in [e_n, \overline{e}], \end{cases} \tag{6}$$

where the partition cutoffs $e_1, ..., e_n$, with $\underline{e} < e_1 < e_2 < ... < e_n < \overline{e}$, separate the domain of types $[\underline{e}, \overline{e}]$ into disjoint subsets. Disclosure plans of this type resemble commonly used discrete ratings systems in practice, which typically pool ranges of types under a given ratings category. It is, however, worth noting that full revelation is also feasible under ex post verifiability, in which case the number of possible ratings goes to infinity.

Government interventions. Since the domestic firm does not internalize the negative externalities affecting domestic households, there is a potential role for interventions by the domestic government. In the presence of a tariff τ that is applied to the transaction between the domestic and the foreign firm, the domestic firm's payoff conditional on trade occurring is given by:

$$u_D = \delta v - p - \tau. \tag{7}$$

Tariff revenues accrue to the domestic government and are rebated lump-sum to domestic households. They are therefore transfers within the domestic economy and do not enter total domestic surplus, which conditional on trade equals

$$u_D + u_H + \tau = v - \lambda e - p. \tag{8}$$

Timing and equilibrium. The timing of the game is as follows:

1. The foreign government chooses a disclosure policy represented by a function $D : [\underline{e}, \overline{e}] \to S$.
2. The domestic government chooses a tariff schedule that is a function of the possible disclosures S.
3. The foreign firm privately observes e.

4. A disclosure signal $D(e) \in S$ is made public according to the plan $D(e)$.

5. The domestic firm makes a take-it-or-leave-it offer to the foreign firm.

6. The foreign firm accepts or rejects the offer.

7. International production occurs conditional on an acceptance and payoffs are realized.

Figure 1 illustrates this timeline. It is worth highlighting that the domestic government determines its tax policy *after* the disclosure policy has been chosen. The same equilibrium analysis would apply if the domestic government moved first, but could not commit to a taxation policy, that is, if it optimally revised its taxation policy once the foreign disclosure regime is actually set. The underlying economic friction is thus a form of limited commitment in negotiations: once a reporting regime is in place, the domestic government chooses the policy that is optimal given the information actually available, and it cannot credibly commit to arbitrary welfare-destroying punishments, such as prohibitive tariffs following any unwanted disclosure regime, whose sole purpose is to influence the earlier information-design choice. Under unrestricted off-path commitments of that kind, the domestic government could dictate the disclosure regime outright and the design problem would disappear. The economics we are interested in concerns credible policies rather than punitive off-path threats.

An equilibrium consists of the foreign government's disclosure policy, the domestic government's tariff scheme, the domestic firm's take-it-or-leave-it offer, and the foreign firm's acceptance decision. We analyze Perfect Bayesian Equilibria (PBE) of this game.[8]

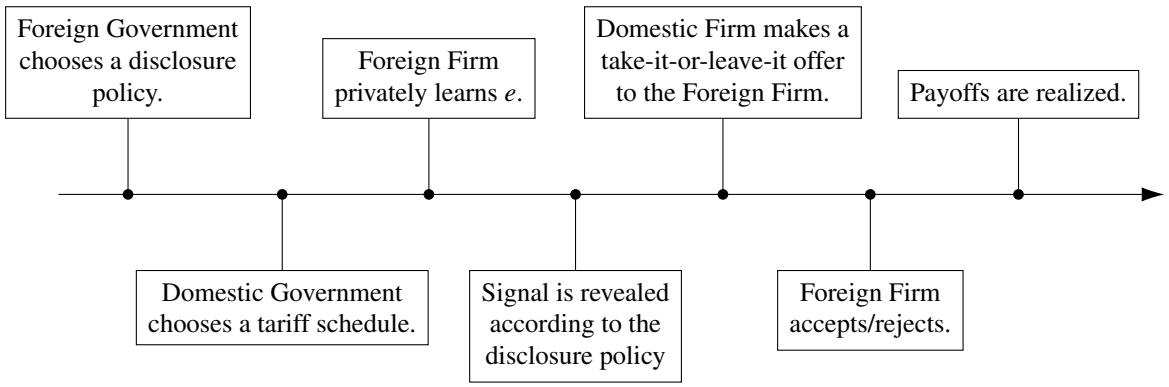


Figure 1: Timeline. The figure illustrates the timeline of the game.

[8]Two tie-breaking conventions are maintained throughout: an indifferent foreign firm accepts the domestic firm's offer; and a domestic government indifferent among marginal types conditional on a disclosure selects the lowest one.

Observability of ex post realizations absent disclosures. One might wonder to what extent it might be feasible for a domestic government to tax parties creating high emissions based on the ex post realizations of emissions (rather than using their voluntary disclosures). A typical feature of externalities such as carbon emissions is that they are not measurable at the level of an individual firm unless measurements are taken locally (at the foreign firm's plant), at the time of their creation.[9] Otherwise, it is not possible to trace emissions back to their source. Formally, this is consistent with an environment where each supply chain relation is atomistic in size, yet households are exposed to aggregate emissions which represent the integral across all individual supply chain interactions. In this case, observing global emissions does not reveal an individual foreign firm's production externalities and thus cannot be used to determine penalties for their domestic supply chain partners, consistent with the setting we study.

A full-information benchmark. Before turning to strategic disclosure, it is useful to discuss what happens when e is publicly observed. With its supplier's costs known, the domestic firm extracts the entire bilateral surplus by offering the price $p = c - e$, so that domestic surplus conditional on trade coincides with global surplus, $v - c + (1 - \lambda)e$. The domestic government's optimal policy then implements trade if and only if domestic surplus (which equals global surplus) is non-negative, for instance, via a type-contingent Pigouvian tariff. Under full information there is thus no conflict between the domestic government's objective and global efficiency: because the domestic side appropriates the entire surplus, maximizing domestic surplus maximizes global surplus. The inefficiencies identified in the analysis that follows are accordingly attributable to the foreign side's private information and the endogenous coarseness of its disclosures, as well as the domestic firm's exercise of market power under this information asymmetry.

3 Model Analysis

We proceed as follows. Section 3.1 characterizes equilibrium disclosures and trade absent government intervention — the laissez-faire benchmark. Section 3.2 turns to optimal intervention: we characterize the domestic planner solution conditional on disclosures, the foreign government's equilibrium disclosure response, and how both vary with households' exposure to externalities. Section 3.3 shows how the planner solution is implemented through disclosure-conditional tariff schedules. Section 3.4 provides a complete characterization for uniformly distributed externalities and compares the resulting equilibrium tariff schedule with carbon market prices.

[9] See Barwick et al. (2024) and Cong et al. (2024) among others.

3.1 Laissez-Faire Equilibrium

We first study equilibrium outcomes absent interventions by the domestic government. By backward induction, we start by analyzing the trading game with and without disclosures and then turn to the foreign government's optimal choice of the disclosure plan.

Trade absent disclosure. To start, consider the case absent disclosures by the foreign government. The domestic firm's optimal pricing decision is characterized by the following program:

$$p^* = \arg\max_p \{(\delta v - p) \cdot (1 - G(c - p))\}, \tag{9}$$

where the first term represents the domestic firm's payoff conditional on trade occurring, and where the second term is the probability of trade; all foreign firm types satisfying $e \geq (c - p)$ accept a price offer p. This problem is equivalent to the domestic firm choosing the marginal foreign firm type $x = c - p$ that is just indifferent between accepting and rejecting the offer. The solution to the following problem yields the corresponding optimal marginal type:

$$e^* = \arg\max_{x \in [\underline{e}, \overline{e}]} \{(\delta v - c + x) \cdot (1 - G(x))\}. \tag{10}$$

If $e^* = \underline{e}$, trade occurs with probability one, that is, irrespective of the realization of e. In this case, not disclosing information is optimal for the foreign firm, since the equilibrium offer by the domestic firm is already the highest possible price. To focus our analysis on scenarios where disclosures may optimally occur in equilibrium absent government interventions, we make the following assumption going forward.

Assumption 2. *Absent disclosure, the probability of trade is less than one, $e^* > \underline{e}$.*

When Assumption 2 is satisfied, trade breakdowns occur with positive probability absent disclosures due to the presence of asymmetric information and the domestic firm's exercise of market power. In particular, trade breaks down for *low* levels of externalities $e < e^*$, since greener production is associated with relatively high production costs for the foreign firm.

In case an interior solution obtains, $e^* \in (\underline{e}, \overline{e})$, the following first-order condition holds:

$$\Pi'(e^*) = 1 - G(e^*) - (\delta v - c + e^*) \cdot g(e^*) = 0. \tag{11}$$

We restrict attention to environments in which $2g(e) + g'(e)(\delta v - c + e) > 0$ for all e, so that the marginal profit function $\Pi'(\cdot)$ is strictly decreasing and the first-order condition (11) characterizes the unique optimum, unless there is a corner solution. This condition holds, for example, under the uniform distribution and under the truncated normal parameterizations (see Appendix B.1).

Moreover, since conditioning on an interval $[a,b)$ only replaces the constant 1 by $G(b)$ in the marginal profit function, the same condition ensures that the marginal profit conditional on any disclosure is strictly decreasing as well.

Trade with disclosure. Conditional on the foreign firm disclosing a signal $\tilde{s} \in S$ revealing that $e \in [a,b)$ and implying the associated conditional CDF:

$$G_{\tilde{s}}(e) \equiv \frac{G(e)-G(a)}{G(b)-G(a)}, \tag{12}$$

the domestic firm solves the program:

$$e^*_{\tilde{s}} = \arg\max_{x\in[a,b)} \{(\delta v - c + x)\cdot(1-G_{\tilde{s}}(x))\}. \tag{13}$$

Conditional on an interior solution, this yields the first-order condition:

$$\Pi'_{\tilde{s}}(e^*_{\tilde{s}}) = 1 - G_{\tilde{s}}(e^*_{\tilde{s}}) - (\delta v - c + e^*_{\tilde{s}})\cdot g_{\tilde{s}}(e^*_{\tilde{s}}) = 0. \tag{14}$$

The following lemma characterizes a key property of disclosure plans that are privately optimal for foreign firms.

Lemma 1. *Suppose Assumptions 1 and 2 hold. Absent government interventions, the foreign government optimally designs a partial disclosure plan that ensures that trade occurs if and only if it generates positive joint surplus for the domestic and foreign firm, that is, for*

$$u_D + u_F \geq 0 \;\;\Leftrightarrow\;\; e \geq c - \delta v. \tag{15}$$

Proof. See Appendix. □

The logic is a splitting argument that recurs throughout our analysis.[10] If trade broke down with types whose bilateral surplus is positive, the foreign government could split the corresponding disclosure interval at the marginal type: prices and trading outcomes above the split would be unchanged, while a positive measure of previously excluded types would then trade at a price that leaves them information rents. Such a split strictly raises the foreign firm's ex ante payoff, so a plan inducing bilaterally inefficient breakdowns cannot be optimal.

[10] A related result is derived in Glode, Opp, and Zhang (2018).

Optimal disclosure plans. Going beyond this general property, we now characterize optimal disclosure plans. To start, we define

$$\hat{e}^o \equiv \max(c - \delta v, \underline{e}), \tag{16}$$

that is, the lowest type of e such that there are positive bilateral gains from trade between the domestic and foreign firm (which excludes the damages to households). Given the result of Lemma 1, we can assume without loss of generality that all types in $[\underline{e}, \hat{e}^o)$ are pooled together in one disclosure interval, since trade will not occur with those types in equilibrium.

Next, we define functions that formalize the maximum range of types that can be pooled under one signal without inducing inefficient screening by the domestic firm. Whether trade occurs with all types in a disclosure interval $[a,b)$ depends only on the interval's own endpoints: conditional on the disclosure, the domestic firm's pricing problem involves only the distribution of types within $[a,b)$, and full trade requires the interval to be narrow enough that the firm optimally prices at the lower endpoint a. Accordingly, we define the left efficiency bound $lb(b)$ as the smallest lower endpoint a for which the interval $[a,b)$ sustains full trade, and conversely the right efficiency bound $rb(a)$ as the largest upper endpoint b for which it does. Explicit expressions are given in the appendix, equations (A3)–(A4).

In the appendix (Lemma 3), we show that given any two neighboring partition cutoffs e_{i-1} and e_{i+1}, the interior cutoff that maximizes the foreign firm's expected payoff takes one of only two candidate values: it either equals $rb(e_{i-1})$ or $lb(e_{i+1})$; that is, the cutoff makes either the disclosure interval below it or the one above it as wide as full trade permits. The function $\Phi(u,w)$, defined explicitly in the appendix (equation (A5)), measures the foreign firm's expected net benefit of the second candidate, $lb(w)$, relative to the first, $rb(u)$, given neighboring cutoffs u and w. The sign of Φ is not determined in general, but under a large range of specifications for the distribution $G(\cdot)$ it is weakly positive everywhere on its domain Ω (also defined in the appendix), including the uniform distributions we consider below and the truncated Normal distributions in the appendix, rendering the characterization of an optimal disclosure plan highly tractable. Correspondingly, we specify the following technical condition.

Assumption 3. $\Phi(u,w) \geq 0, \forall (u,w) \in \Omega$*; that is, given any two neighboring partition cutoffs u and w, the foreign firm's expected payoff is weakly higher when the interior cutoff between them is placed at* $lb(w)$ *than when it is placed at* $rb(u)$.

We show analytically in the appendix (Lemma 4) that if externalities are uniformly distributed, Assumption 3 is satisfied. Further, we have verified the assumption numerically for the truncated normal parameterizations used in the appendix (Section B.1).

We can now proceed to characterizing optimal disclosure plans absent government interventions.

Proposition 1. *Suppose Assumptions 1–3 hold. Absent government interventions, the partition cutoffs of an optimal disclosure plan function are given by a descending sequence starting with* $e_{n+1} = \overline{e}$ *and where:*

$$e_i = lb(e_{i+1}), \qquad for \quad i = n, (n-1), ..., 3, 2. \tag{17}$$

When $c - \delta v < \underline{e}$*, this sequence* $\{e_n, e_{n-1}, e_{n-2}, ..., e_2\}$ *is finite. Otherwise, this sequence is infinite (that is,* $n \to \infty$*) and converges to* $\hat{e}^o$ *from above.*

While we will discuss examples with concrete distributions below, it is worth highlighting a general property of disclosure plans based on Lemma 1 and Proposition 1: The foreign firm benefits from maximizing the size of the disclosure intervals subject to the constraint that the disclosure does not induce trade breakdowns that are bilaterally inefficient. As such, for values of e where the bilateral gains from trade $(\delta v - c + e)$ become small, the disclosure intervals and the associated information asymmetries also need to shrink to prevent such trade breakdowns. Naturally, the smaller the disclosure interval, the less information asymmetry remains between the supply chain partners. If bilateral gains from trade approach zero (for $e \downarrow \hat{e}^o$), the disclosure intervals in fact become locally arbitrarily small in the vicinity of $\hat{e}^o$. We illustrate Proposition 1 using an example of uniform distributions.

Example 1. Suppose $e \in U[\underline{e}, \overline{e}]$. Assumption 3 holds by Lemma 4 in the appendix. Applying Proposition 1, absent government interventions, the partition cutoffs of an optimal disclosure plan function are given by a descending sequence starting with $e_{n+1} = \overline{e}$ and where:

$$e_i = lb(e_{i+1}) = \frac{e_{i+1} - (\delta v - c)}{2}, \qquad \text{for} \quad i = n, (n-1), ..., 3, 2. \tag{18}$$

Here, n is chosen so that $lb(e_2) \leq \underline{e}$, equivalently $\frac{e_2 - (\delta v - c)}{2} \leq \underline{e}$. That is, the cutoffs of the optimal disclosure plan are obtained by sequentially applying the left-bound function, starting from the right end of the support of the distribution.

3.2 Government Interventions

In this subsection, we consider government interventions. While government interventions can take various forms, we focus on the optimal interventions from the perspective of the domestic

country. A disclosure-conditional tariff schedule $\tau(\tilde{s})$ constitutes a sufficient policy instrument in this environment: by varying $\tau(\tilde{s})$, the domestic regulator can induce the domestic firm to make any feasible price offer $p(\tilde{s})$. The optimal intervention therefore reduces to designing the tariff schedule that maximizes domestic surplus subject to the residual information asymmetry. In the appendix we consider the case where the government imposes an unconditional transaction tax τ that does not vary with either the foreign firm's disclosures or the domestic firm's price offers (naturally τ is a subsidy when it takes negative values).

Before characterizing equilibrium disclosures, we state the government's objective conditional on a given disclosure. Suppose a disclosure reveals $e \in [a,b)$ and the government's tariff induces the marginal type $x \in [a,b]$, so that trade occurs at the price $c-x$ if and only if $e \geq x$. Since domestic surplus conditional on trade equals $v - \lambda e - p$, the domestic country's expected surplus conditional on the disclosure is proportional to

$$d(x;b) \;\equiv\; (G(b)-G(x))\,[v-c+x] \;-\; \lambda \int_x^b e\, dG(e), \tag{19}$$

where $x = b$ corresponds to no trade and $d(b;b) = 0$. The objective does not depend on the category's lower endpoint a. This expression exhibits the central tension of this section. Raising the tariff lowers the price the domestic firm offers and thereby raises the marginal type x: the producers that stop trading first are the relatively green ones, whose costs $c-e$ are highest, while browner producers continue to trade. A marginal tariff increase therefore removes the trades with the *smallest* externalities and leaves a browner trading pool, an adverse-selection effect. The results below characterize when the government's fine-tuning of tariffs remains optimal despite this effect, and when it fails entirely.

We make the following assumption.

Assumption 4. *$v - c + (1-\lambda)\underline{e} > 0$ and $v - c + (1-\lambda)\overline{e} < 0$.*

This assumption implies that $\lambda > 1$, since $\overline{e} > \underline{e}$. Assumption 4 states that global surplus is positive at the minimum level of externalities $e = \underline{e}$ and negative at the maximum level of externalities $e = \overline{e}$. As a result, there is a unique interior threshold level $\hat{e} \in (\underline{e}, \overline{e})$ at which the global surplus conditional on trade occurring is exactly zero. This threshold level is given by:

$$\hat{e} = \frac{v-c}{\lambda - 1}. \tag{20}$$

Proposition 2. *Suppose Assumptions 1 and 4 are satisfied and that the hazard rate $\frac{g(e)}{1-G(e)}$ is an increasing function. Then, if the domestic government implements the optimal regulatory policy,*

there exists a threshold level $\underline{\lambda} \geq 1$, such that if the domestic sector's exposure satisfies $\lambda < \underline{\lambda}$, the foreign government's optimal disclosure policy implies that total global surplus is maximized, that is, trade occurs if and only if

$$u_H + u_D + u_F > 0. \tag{21}$$

As shown in Section 3.1, absent government interventions, voluntary disclosure plans generally ensure that trade occurs if and only if it generates positive private bilateral surplus for the two firms in the supply chain. However, one may be naturally interested in total global surplus, which additionally accounts for the disutility households experience due to production externalities. Proposition 2 shows that when households are not highly exposed to externalities, the optimal disclosure plans also ensure that total global surplus is maximized, provided that the domestic government intervenes optimally.

The proof of Proposition 2 (in the appendix) establishes that the threshold $\underline{\lambda}$ admits the explicit form

$$\underline{\lambda} = \inf_{(b,x):\underline{e}\leq x<b\leq\bar{e}} \left\{ 1 - \frac{d}{dx}\left[\frac{G(b)-G(x)}{g(x)}\right] \right\}, \tag{22}$$

which satisfies $\underline{\lambda} \geq 1$ under the increasing-hazard-rate assumption. The threshold strictly exceeds 1, so that the moderate-exposure range $(1, \underline{\lambda})$ is nonempty, whenever the derivative of the ratio $(G(b)-G(x))/g(x)$ is bounded away from zero; in the uniform case, that derivative equals -1, so $\underline{\lambda} = 2$.[11]

However, if domestic households' exposure to externalities is sufficiently large, this result no longer holds. The exposure parameter λ modulates the extent to which the domestic sector's view on the desirability of high-externality transactions differs from that of the foreign trade partner: for higher values of λ, the domestic sector's benefits from trade under high-externality production are decreased, whereas the foreign firm's benefits remain unchanged. For sufficiently high values of λ any marginal increase in border tariffs is associated with such a potent negative selection effect (as only the higher-externality suppliers continue to trade) that it never improves domestic surplus. As a result, the domestic planner solution is an extreme bang-bang solution: the domestic sector either optimally trades with the whole pool of foreign firm types conditional on a given disclosure or with none of them. Put differently, technical conditions that were satisfied absent government interventions and ensured that marginal optimality conditions held are now violated. Facing such

[11] The strict inequality in Proposition 2 is also not restrictive at the boundary: in the uniform case with $\lambda = 2$ exactly, the disclosure plan characterized in Proposition 7 below has $e_1 = \hat{e}$, so trade remains globally efficient, while excessive pooling arises for every $\lambda > 2$. The efficient and excessive-pooling regimes thus partition the parameter space in the uniform case, with no intermediate region.

an extreme domestic policy, the foreign government optimally designs a disclosure policy that leads to excessive pooling, in the sense that trade occurs when it destroys total global surplus.

Proposition 3. *Suppose Assumptions 1 and 4 hold and that $e \cdot g(e)$ is strictly increasing with a continuous, strictly positive derivative, which is for example satisfied when e is uniformly distributed. Then, if the domestic government implements the optimal regulatory policy, there exists a threshold level $\hat{\lambda}$, such that when $\lambda \geq \hat{\lambda}$, the foreign government's optimal disclosure plan induces excessive trade from a global perspective, that is, trade occurs for a range of externalities e such that*

$$u_H + u_D + u_F < 0. \tag{23}$$

Proof. See Appendix. □

When the domestic households are sufficiently strongly exposed to environmental damages, which is the case for high enough values of the parameter λ, the foreign government chooses a disclosure policy that pools types above the threshold level $\hat{e}$ with those below that level. As a result, when the disclosure indicates that externalities are in that region, trade occurs even though it may destroy global surplus. We will discuss this issue further in Section 3.4, where we examine disclosure plans under specific distributions.

While the two propositions characterize the polar regimes, one feature of the severe-exposure regime — the bang-bang form of the realized policy — turns out to be general.

Proposition 4 (No rationing within categories)**.** *Suppose Assumptions 1 and 4 hold and the domestic government implements the planner solution conditional on each disclosure. Then, in any equilibrium, trade occurs with probability one or zero conditional on each disclosure category.*

Proof. See Appendix. □

The logic mirrors Lemma 1: if the planner rationed trade within a category, the foreign government could split off the rationed segment into its own category; the planner would then trade with part of that segment — its types are the greenest in the original pool — and the foreign firm would collect additional rents. In equilibrium, all screening therefore operates through the coarseness of the disclosure plan itself, and the realized policy is bang-bang for *any* distribution: the exposure level governs how granular the equilibrium plan is and where trade stops, not the form of the realized policy. Interior tariffs of the form characterized in Section 3.3 discipline equilibrium behavior off the equilibrium path.

The intermediate exposure regime. The two forces behind Propositions 2 and 3 — the compositional gains from screening out high-emissions types, and the adverse selection that renders marginal tariff increases counterproductive — are present for any distribution of externalities; what varies across distributions is where in the support each force dominates. A virtue of the uniform specification is that this tension resolves cleanly: the planner's objective is globally concave for $\lambda < 2$ and globally convex for $\lambda > 2$, so the two regimes meet at $\lambda = 2$ and the full equilibrium is available in closed form (Section 3.4).[12] For general distributions, the same economics apply locally, and at intermediate exposures the two regimes coexist across disclosure categories. Section 4.1 characterizes this coexistence for a class of densities nesting the uniform case: the equilibrium plan features finely screened trade among greener types together with an excessively pooled top rating category in which trade still occurs, the logics of Propositions 2 and 3 operating side by side.

3.3 Tariff Implementation of the Domestic Planner Solution

In the previous section, we solved the domestic planner's solution by assuming that the government uses a tariff scheme that ensures that the domestic firm acts as if it maximizes total domestic surplus. We now examine how this can be implemented by taxing the domestic firm's supply chain interaction with the foreign firm. We assume that government charges the lowest-absolute-value tax whenever government is indifferent among several different tariffs. This ensures that no tariff is charged if doing so achieves the optimum.

Suppose the foreign firm does not possess private information about e, but the planner sets the tariff before the realization of e is observed, conditioning on a disclosure revealing $e \in [a,b)$ alone. The domestic households' expected payoff is then $(1-\delta)v - \lambda \cdot \mathbb{E}[e \,|\, e \in [a,b)]$. By charging the following Pigouvian tax

$$\tilde{\tau} = \lambda \cdot \mathbb{E}[e \,|\, e \in [a,b)] - (1-\delta)v, \tag{24}$$

the domestic firm internalizes the households' payoff when quoting a price to the foreign firm, and the domestic planner solution is implemented.

Turning to the case where the foreign firm has private information, we examine the domestic government's optimal tariff scheme as a function of the foreign firm's disclosure. Suppose the foreign firm discloses a signal $\tilde{s}$ revealing that $e \in [a,b)$ and that the optimal marginal type $e^{**}_{\tilde{s}}$ for the domestic planner is an interior solution $e^{**}_{\tilde{s}} \in (a,b)$, which satisfies the planner's first-order

[12] The planner's objective (19) has second derivative $d''(x;b) = -(2-\lambda)g(x) - g'(x)[v - c + (1-\lambda)x]$; for non-constant g, this can change sign within a disclosure interval, in which case the planner's response is regular on part of the support and bang-bang elsewhere.

condition:

$$1 - G_{\tilde{s}}(e^{**}_{\tilde{s}}) = (v - c + (1-\lambda)\cdot e^{**}_{\tilde{s}})\cdot g_{\tilde{s}}(e^{**}_{\tilde{s}}). \tag{25}$$

The following proposition characterizes how the optimal tax depends on this marginal type $e^{**}_{\tilde{s}}$.

Proposition 5. *Suppose that the hazard rate $\frac{g(e)}{1-G(e)}$ is an increasing function. Suppose the foreign firm discloses a signal $\tilde{s}$ revealing that $e \in [a,b)$, and the domestic planner solution induces trade for $e \geq e^{**}_{\tilde{s}}$, where $e^{**}_{\tilde{s}} \in (a,b)$. Then the government can implement this solution by imposing a tariff*

$$\tau = \lambda e^{**}_{\tilde{s}} - (1-\delta)v. \tag{26}$$

Comparing the optimal tax under asymmetric information (26) to the optimal tax under symmetric information (24) reveals that both tariffs have a similar structure. However, whereas the tax under symmetric information leads the domestic firm to internalize the *expected* externality if trade were to occur (potentially off equilibrium), the optimal tax under asymmetric information causes the domestic firm to internalize the externality caused by the *marginal* type $e^{**}_{\tilde{s}}$ with which trade occurs under the domestic planner solution. This marginal type depends on the equilibrium disclosure choices of the foreign government and the residual information asymmetry characterized by the conditional distribution $G_{\tilde{s}}(e)$.

Two comparative statics follow directly from the structure of the optimal tariff; formal statements and proofs are in the appendix. First, the equilibrium under optimal regulation is invariant to how the value v is split between the domestic firm and households: the induced marginal type, the disclosure plan, and all equilibrium quantities are independent of δ, with the tariff adjusting one-for-one (Corollary 3). We therefore set $\delta = 1$ in the worked examples below without loss of generality. Second, if the domestic firm itself internalizes a share α of the household externality — for example, through pressure from environmentally responsible consumers or investors — the planner solution is unchanged and the optimal tariff falls one-for-one with the internalized share (Corollary 4).

Why prices, not quantities. The optimal instrument we have characterized takes the form of a conditional tariff schedule $\tau(\tilde{s})$, with the foreign firm's accept/reject decision endogenously determining whether trade occurs. This is a *price* instrument: the regulator commits to a charge per accepted transaction and lets the privately informed foreign firm decide whether to participate. A direct *quantity* instrument, such as cap-and-trade applied to a given supply chain relation, is strictly dominated in this environment. The reason is that within any disclosure category $\tilde{s}$, the optimal quantity is the indicator $\mathbb{1}\{e \geq e^{**}_{\tilde{s}}\}$: it depends on the realization of e, which the regulator does not observe, while the threshold $e^{**}_{\tilde{s}}$ itself varies with the conditional distribution of types in the

category. Fixing a quantity *ex ante* would either fail to admit some transactions that should occur or admit some that should not, and either error reduces welfare. The optimal mechanism therefore sets prices and elicits quantity, rather than the converse, providing an asymmetric-information counterpart to the classical price-versus-quantity comparison in Weitzman (1974). Cap-and-trade can implement the optimum only in the limiting case of a continuum of independent and identically distributed foreign producers within a category, where a law-of-large-numbers argument pins down the optimal aggregate quantity deterministically; this limit is not the empirically relevant case for the concentrated cross-border supply relationships our analysis emphasizes.

3.4 Uniformly Distributed Externalities: A Complete Characterization

In this section, we provide a full characterization of optimal government policies and optimal disclosure plans when externalities are uniformly distributed, $e \sim U[\underline{e}, \overline{e}]$. As a first step, the following proposition characterizes the government's optimal planner solution conditional on any disclosure by the foreign firm. Moreover, it details how the planner solution can be implemented via tariffs. In stating the proposition, we adopt the convention that the case-defining conditions are expressed in terms of the planner's interior-optimum break-even formula (involving $v-c$), while the tariff levels in each case are expressed in terms of the firm's break-even formula (involving $\delta v - c$). The two formulas differ because the planner's objective is $u_D + u_H = v - p - \lambda e$ whereas the firm's is $\delta v - p - \tau$; the optimal tariff τ exactly bridges the two so that the firm's induced marginal type matches the planner's.

Proposition 6 (Tax Implementation)**.** *Suppose externalities are uniformly distributed, $e \sim U[\underline{e}, \overline{e}]$, and $\lambda > 1$.*

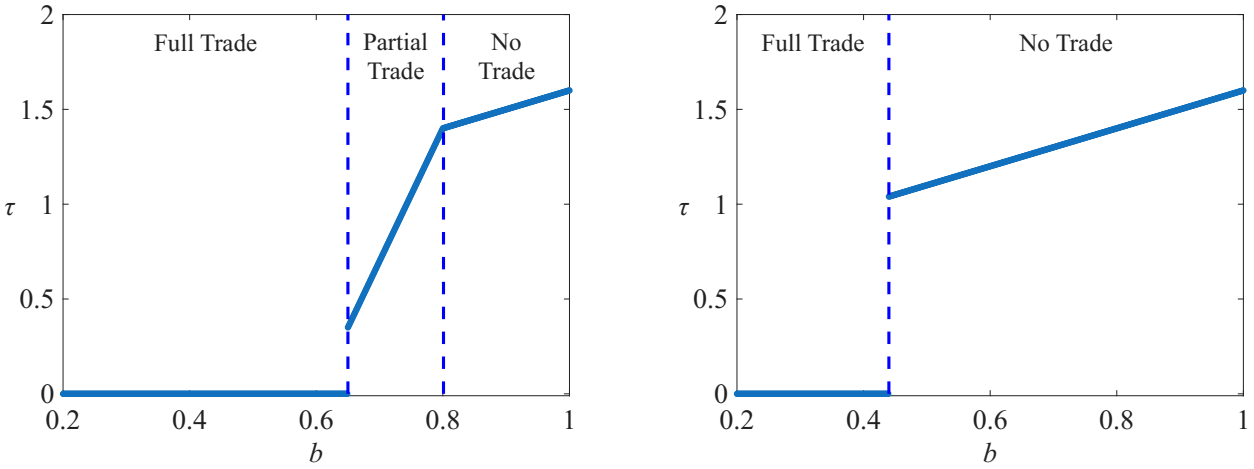
1. *Suppose $\lambda < 2$ and the foreign firm's disclosure reveals that $e \in [a,b)$. The domestic planner solution and its implementation via tariffs have the following properties:*

$$\tau \begin{cases} \leq 2a - b + \delta v - c & \text{if } a \geq \frac{b-(v-c)}{2-\lambda}, \\ = \frac{\lambda(b-\delta v+c)-2(1-\delta)v}{2-\lambda} & \text{if } a < \frac{b-(v-c)}{2-\lambda} < b, \\ \geq b + \delta v - c & \text{if } \frac{b-(v-c)}{2-\lambda} \geq b. \end{cases}$$

2. *Suppose $\lambda \geq 2$ and the foreign firm's disclosure reveals that $e \in [a,b)$. The domestic planner solution and its implementation via tariffs have the following properties:*

$$\tau \begin{cases} \leq 2a - b + \delta v - c & \text{if } v - \lambda \frac{a+b}{2} - (c-a) \geq 0, \\ \geq b + \delta v - c & \text{otherwise.} \end{cases}$$

Proof. See Appendix. □



(a) $\lambda = 1.75$ **(b)** $\lambda = 2.5$

Figure 2: The figure considers a parameterization where $e \sim U[0,1]$, $\delta = 1$, and $(v-c) = 0.6$. It illustrates the tariff implementing the domestic planner solution for a generic disclosure revealing that $e \in [0.2, b)$ as a function of the upper bound b of the disclosure interval. The left-hand side panel considers the case where $\lambda = 1.75$, and the right-hand side panel the case where $\lambda = 2.5$.

Proposition 6 shows that outcomes differ dramatically between the cases of low and high environmental exposures, as represented by $\lambda < 2$ and $\lambda > 2$. At the boundary $\lambda = 2$, the planner's objective is linear in the marginal type, so the bang-bang characterization of the second case applies. The domestic planner's payoff is given by the probability of trade times the payoff conditional on trade, where the latter is

$$v - c - \lambda \, \mathbb{E}[e | e \geq e^{**}] + e^{**}, \tag{27}$$

where e^{**} is the cut-off-type firm that trades. Since a higher induced cut-off type leads to a lower trade probability, the probability of trade is a decreasing function of e^{**}. When e follows a uniform distribution with upper bound b, the domestic planner's expected payoff conditional on trade is $v - c - \lambda \frac{e^{**}+b}{2} + e^{**}$. It is evident that when $\lambda < 2$, the payoff conditional on trade is an increasing function of e^{**} while when $\lambda > 2$, it is a decreasing function. In the former case ($\lambda < 2$), there is typically an interior e^{**} that maximizes the domestic planner's payoff. In the latter case ($\lambda > 2$), the domestic planner optimally induces an equilibrium that features either trade with all types or none of them, depending on whether trade with all types leads to a positive expected payoff. Anticipating this policy by the domestic government, the foreign government then optimally chooses a plan with at most two intervals where one of them pools as many green types (low e) as possible subject to trade occurring and the other one containing all remaining browner (high e) types.

Figure 2 illustrates Proposition 6 for a generic disclosure revealing that $e \in [0.2, b)$, plotting the implementing tariff as a function of the upper bound b. Absent intervention, trade would occur with probability one throughout. Under severe exposure ($\lambda = 2.5$, right-hand side panel), the planner solution switches discretely between full trade — implemented with no tariff — and no trade, implemented with a prohibitive tariff that increases in b. Under moderate exposure ($\lambda = 1.75$, left-hand side panel), an intermediate range of b features a strictly positive tariff that induces an interior marginal type: trade then occurs only if externalities *exceed* that type, because brown production's low costs are passed through to the domestic economy via low prices; for yet higher b, the tariff again shuts trade down. In light of Proposition 4, such interior marginal types arise only off the equilibrium path: their role is to discipline the foreign government's disclosure choice.

Anticipating this domestic planner solution, the foreign government chooses its optimal disclosure plan, which is characterized in the following proposition.

Proposition 7 (Optimal Disclosure)**.** *Suppose externalities are uniformly distributed, $e \sim U[\underline{e}, \overline{e}]$, $\lambda > 1$, and the government implements the domestic planner solution.*

1. *When $\lambda < 2$, the privately optimal disclosure plan chosen by the foreign government ensures that trade occurs if and only if $u_H + u_D + u_F > 0$. Its partition cutoffs are given by the increasing sequence*

$$e_{i+1} = (2 - \lambda)\, e_i + (v - c), \qquad e_0 \equiv \underline{e}, \tag{28}$$

   which converges to $\hat{e}$; all types in $[\hat{e}, \overline{e}]$ are pooled under a single signal, conditional on which no trade occurs.

2. *When $\lambda \geq 2$, the privately optimal disclosure plan chosen by the foreign government pools types into at most two intervals, $[\underline{e}, e_1)$ and $[e_1, \overline{e}]$, where*

$$e_1 = \frac{2(v - c) + 2\underline{e}}{\lambda} - \underline{e}. \tag{29}$$

   If the expression in (29) *weakly exceeds $\overline{e}$, then e_1 is capped at $\overline{e}$: the optimal plan is uninformative and pools all types into the single trading interval $[\underline{e}, \overline{e}]$. Trade occurs with probability 1 conditional on the first interval and with probability 0 otherwise. Global surplus, $(u_H + u_D + u_F)$, is strictly negative when externalities take values in the interval $(\hat{e}, e_1)$.*

Proof. See Appendix. □

Proposition 7 makes the dichotomy of Section 3.2 exact: the two regimes are delimited by $\lambda = 2$. When even the full pool $[\underline{e}, \overline{e}]$ yields positive average domestic surplus, expression (29)

exceeds $\bar{e}$: the optimal plan is then uninformative and trade occurs for all types, which remains excessive from a global perspective.

Combining the equilibrium disclosure plan with the tariffs that implement it delivers a sharp characterization of the implementing tariff *schedule* — and of how it differs from prices in carbon credit markets, under which importers would pay a common price q per disclosed ton, i.e., a proportional schedule $\tau = qm$ in disclosed expected emissions m. In what follows, a *trading category* is a rating category conditional on which trade occurs in equilibrium.

Corollary 1 (Optimal tariffs differ from prices in carbon credit markets)**.** *Suppose externalities are uniformly distributed with $\underline{e} \geq 0$ and $1 < \lambda < 2$, and consider the equilibrium characterized in Propositions 6 and 7. Let m_i denote expected emissions conditional on trading category $[e_i, e_{i+1})$. Then:*

(i) The maximal tariff implementing the planner solution on trading category i — the unique tariff at which the domestic firm's first-order condition holds exactly at the category's marginal type e_i — is $\tau_i = \lambda e_i - (1-\delta)v$: it prices the externality of the category's greenest *type.*

(ii) Across trading categories, the schedule is affine in disclosed expected emissions,

$$\tau_i \;=\; \frac{2\lambda}{3-\lambda}\, m_i \;-\; \frac{\lambda\,(v-c)}{3-\lambda} \;-\; (1-\delta)\,v, \tag{30}$$

with a slope that strictly exceeds the marginal damage per ton, λ, while its level satisfies $\tau_i < \lambda m_i - (1-\delta)v$.

(iii) For disclosures above $\hat{e}$, the schedule jumps to prohibitive levels. Consequently, no proportional schedule $\tau = qm$ reproduces the equilibrium tariffs.

Proof. See Appendix. □

Corollary 1 reports, for each trading category, the *maximal* implementing tariff. Conditional on each such category, the planner prefers that every type in the category trades, and the domestic firm is willing to purchase the entire category at any tariff up to τ_i: the implementing tariffs form the interval $\tau \leq \tau_i$, and every tariff in this interval — including, whenever $\tau_i \geq 0$, a zero tariff — induces the same trades at the same prices, so neither country's welfare nor the foreign government's disclosure incentives depend on which one is levied. Reporting the maximal tariff is nonetheless not arbitrary, for two reasons. First, it characterizes the entire interval of implementing tariffs, so the corollary describes the full set of implementing schedules rather than a particular element. Second, the multiplicity of implementing tariffs arises only on the equilibrium path: if the foreign

government deviated to a disclosure plan with a category on which the planner's response excludes some types from trade, the implementing tariff would be unique — the domestic firm's first-order condition must then hold exactly at the planner's marginal type, yielding (26) — and τ_i is precisely this unique tariff evaluated at the equilibrium marginal type e_i. The corollary thus describes the tariff policy where it disciplines the foreign government's disclosure choice. A proportional schedule cannot replicate this discipline: a common price per disclosed ton low enough to preserve trade on the equilibrium categories would not deter the foreign government from choosing a coarser disclosure plan.

The schedule is affine, but not proportional, in disclosed expected emissions, and both of its deviations from a carbon price reflect the marginal-type logic of Section 3.3. Its level lies below the Pigouvian benchmark λm_i because taxing a category's *average* emissions would over-deter its greenest, marginal types and inefficiently ration green trade. Its slope exceeds λ because trading categories narrow as they approach $\hat{e}$, so expected emissions rise more slowly across categories than the marginal type deteriorates. For $\lambda \geq 2$, beyond the corollary's hypothesis, the schedule takes only two levels: one admitting trade with the green pool and one prohibiting trade with the residual pool. Together with the prohibitive region above $\hat{e}$ and the binary schedule under severe exposure, these properties imply that the effective charge per disclosed ton varies systematically across disclosure categories, rising with disclosed brownness, rather than being constant as under a common carbon price. The result also identifies *which* end of a disclosed range should anchor the charge: the optimal tariff keys to the greenest emissions level consistent with each disclosure, the opposite end from the conservative default values applied under CBAM.

The anchoring at the *greenest* type reflects the model's sharp rating boundaries, but the economics behind it is robust. First, tariffs allocate trade at the margin: charging a category's average emissions would over-deter its marginal members, whose gains from trade are smallest. Second, because greener production is privately more expensive, those marginal members are a rating's relatively green producers — and this remains true on average even if some producers are misclassified, a case our model does not formalize. Combining the two, the optimal charge per rating lies below the rating's average emissions, skewed toward its greener types, while the schedule steepens across ratings; only the exact formula, the externality of the boundary type, relies on perfect classification.

The tractability of the uniform specification also allows us to cleanly analytically quantify how costly the excessive pooling under severe exposure is. Measuring welfare as expected global surplus and taking as reference the full-information benchmark of Section 2, the welfare loss admits a closed form.

Corollary 2 (The welfare cost of excessive pooling)**.** *Suppose externalities are uniformly dis-*

tributed, $e \sim U[0,\overline{e}]$*,* $\lambda > 2$ *with* $2(v-c) \leq \lambda\overline{e}$*, and the government implements the domestic planner solution. Let* W^{FB} *denote expected global surplus under the full-information benchmark and* W^{eq} *its value in the equilibrium of Propositions 6 and 7. Then*

$$\frac{W^{FB} - W^{eq}}{W^{FB}} = \left(\frac{\lambda - 2}{\lambda}\right)^2. \tag{31}$$

Proof. See Appendix. □

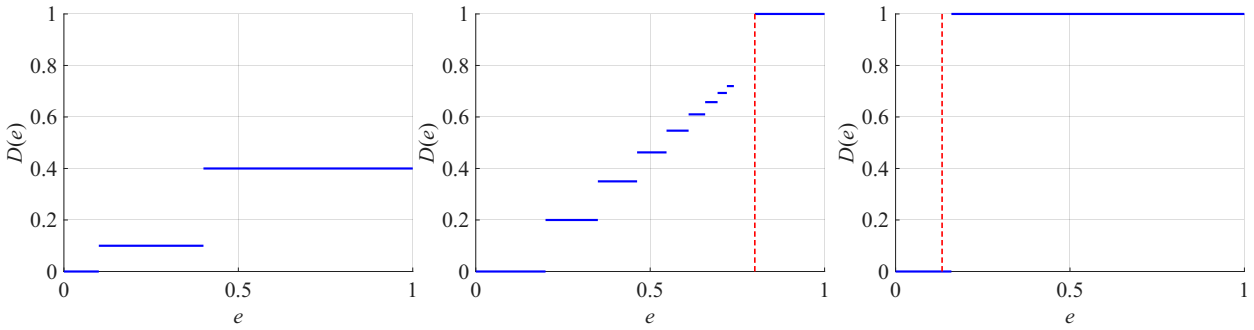
The relative loss depends only on the exposure parameter: it is zero at the regime boundary $\lambda = 2$, rises to one-ninth at $\lambda = 3$ and one-fourth at $\lambda = 4$, and approaches the entire first-best surplus as exposure grows large. That the gains from trade $v - c$ drop out reflects that both welfare levels scale with $(v-c)^2$: the fraction of first-best surplus lost depends only on the exposure parameter, not on the size of the gains from trade.

We now illustrate the results of this section with the standard uniform distribution.

Example 2. Suppose the environmental externalities are uniformly distributed, $e \sim U[0,1]$, that is, $\underline{e} = 0$ and $\overline{e} = 1$, and suppose that $\delta = 1$.

No government interventions. Absent government interventions, the optimal disclosure plan is the one characterized in Example 1; for $(v-c) = 0.2$, it features the three intervals shown in Panel (a) of Figure 3.

Domestic planner solution. Suppose the government imposes an optimal conditional tariff scheme that implements the domestic planner solution as stated in Proposition 6. Further, the foreign government chooses the corresponding optimal disclosure plan as stated in Proposition 7.



(a) No intervention & $(v-c)=0.2$ **(b)** Intervention & $\lambda=1.25$ **(c)** Intervention & $\lambda=2.5$

Figure 3: The figure illustrates the optimal disclosure plans chosen by the foreign government absent intervention (left-hand side panel) and in the presence of the optimal domestic tariff scheme (middle and right-hand side panels). In Panels (b) and (c), the dashed red line marks the efficiency threshold $\hat{e}=(v-c)/(\lambda-1)$: under moderate exposure (Panel (b)), trade stops exactly at $\hat{e}$, whereas under severe exposure (Panel (c)), the trading pool extends beyond $\hat{e}$. Panel (a) does not involve λ, so no threshold is marked.

Figure 3 compares the optimal disclosure plan conditional on no government intervention (left-hand side panel) and in the presence of the optimal domestic tariff scheme (middle and right-hand side panels). When government does not intervene (see Panel (a)), the foreign firm optimally provides a relatively granular disclosure plan that splits the domain of externalities into three intervals. Doing so ensures that trade occurs for all externality levels, despite the presence of lower gains from trade. In contrast, once the government tariffs transactions conditional on such disclosures (see Panels (b) and (c)), the disclosure depends on the level of λ. When $\lambda<2$ (see Panel (b)), the disclosure plan is more granular for low externality levels, while pooling a large range of high-externality types under a single signal. The intuition is as follows. Relative to the no-intervention benchmark, the relevant surpluses shrink once households' disutility from externalities is internalized, and smaller surpluses require more granular disclosures to sustain trade. When $\lambda>2$ (see Panel (c)), the foreign government chooses a less granular plan with only two intervals, whose lower interval pools more types than under the no-intervention benchmark. The dashed red lines make the efficiency benchmark visible: in Panel (b), the trading categories accumulate exactly at $\hat{e}$, so trade stops precisely where global surplus turns negative, whereas in Panel (c) the trading pool extends visibly beyond $\hat{e}$ — the excessive trade established in Proposition 3.

The example illustrates how domestic households' exposures shape both the granularity of equilibrium disclosures and the reach of optimal tariffs.

4 Robustness and Extensions

In this section, we consider variations to our baseline model to examine the robustness of our results. We first move beyond the uniform distribution and characterize equilibrium disclosure at intermediate exposure levels. We then analyze what would happen if the domestic firm, rather than the foreign firm, could choose the disclosure plan. Finally, we consider an alternative timing assumption with "interim" disclosure choices that occur after the foreign firm has learnt its type, which is consistent with the foreign firm choosing its disclosure plan (rather than the foreign government). In the appendix, we additionally verify our results numerically for truncated normal distributions, for which intervention again makes disclosures finer under moderate exposure and coarser under severe exposure (Section B.1).

4.1 The Intermediate Exposure Regime

Propositions 2 and 3 characterize optimal intervention when domestic exposure is either moderate or severe. We now move beyond the uniform distribution and analyze the intermediate range of exposures, asking whether the economics of the two polar regimes carry over and how the transition between them unfolds.

The organizing object is the curvature of the planner's conditional objective. Given a disclosure revealing $e \in [a,b)$, let $d(x;b)$ denote the domestic planner's expected surplus when trade occurs with types $e \geq x$, as in equation (19). Its second derivative in x factorizes as

$$d''(x;b) = -g(x)\left[(2-\lambda) + (\lambda-1)\,\varphi(x)\right], \qquad \varphi(x) \equiv \frac{g'(x)}{g(x)}\,(\hat{e}-x), \tag{32}$$

so that local curvature is governed by a single function of the primitives: the log-slope of the density, weighted by the distance to the efficiency threshold $\hat{e}$. The planner's objective is locally concave — the regular logic of Proposition 2 — where $\varphi(x)$ exceeds the threshold $(\lambda-2)/(\lambda-1)$, and locally convex — the bang-bang logic of Proposition 3 — where it falls below. The uniform distribution is the knife-edge case $\varphi \equiv 0$, for which curvature is uniform in sign and switches globally at $\lambda = 2$.

Recall from Proposition 4 that in any equilibrium, trade occurs with probability one or zero conditional on each disclosure category; it remains to determine which categories trade. To characterize the equilibrium plan, we impose the regularity condition suggested by the factorization (32): φ is nonincreasing on the support. Then, on every disclosure interval, the planner's objective is concave and then convex, and full trade is planner-optimal precisely when two bounds hold — one for each force in the model.

Proposition 8 (Equilibrium disclosure at intermediate exposures)**.** *Suppose Assumptions 1 and 4 hold and φ is nonincreasing. Then:*

(i) Conditional on a category $[a,b)$, full trade is planner-optimal if and only if $b \leq B(a) \equiv \min\{b_1(a), b_2(a)\}$, where the screening bound $b_1(a)$ *solves*

$$G(b_1(a)) = G(a) + (\lambda - 1)\, g(a)\, (\hat{e} - a), \tag{33}$$

and the viability bound $b_2(a)$ *is the largest b for which the category's average domestic surplus $d(a;b)$ remains non-negative.*

(ii) In any equilibrium, the set of trading types is an interval $[\underline{e}, T]$ with $T \geq \hat{e}$, partitioned into consecutive full-trade categories. Either the partition is finite and the terminal category's bound binds, or it consists of infinitely many categories whose cutoffs accumulate at $T = \hat{e}$, in which case trade is globally efficient; for $\lambda > 2$, only the former case arises. Trade is globally excessive if and only if $T > \hat{e}$.

(iii) If, in equilibrium, the screening bound is the binding constraint at every non-terminal step, the equilibrium cutoffs are generated by the recursion $a_{i+1} = B(a_i)$ starting from $a_1 = \underline{e}$.

Proof. See Appendix. □

The two bounds in part (i) are the two polar regimes in local form. The screening bound b_1 states that the category must be narrow enough that the planner does not prefer to ration its greenest types — the marginal logic that generates granular disclosure in Proposition 2, and the analog of the full-trade bounds $rb(\cdot)$ from the laissez-faire analysis, with the planner's surplus replacing the bilateral one. The viability bound b_2 states that the category's *average* domestic surplus must remain non-negative — the bang-bang logic of Proposition 3. The equilibrium plan stretches every category to the tighter of the two bounds, and part (ii) shows that trade always extends at least to the efficient threshold $\hat{e}$: under optimal intervention, equilibrium trade is never insufficient, and the entire welfare loss from strategic disclosure takes the form of excessive trade beyond $\hat{e}$. When the viability bound binds at an interior step, the recursion in part (iii) changes form; we therefore do not claim a fully explicit equilibrium characterization for general distributions. For the linear-density family below, the condition of part (iii) can be verified directly.

A transparent class of distributions satisfying all of our conditions consists of increasing linear densities, which nest the uniform case and make the intermediate regime arbitrarily wide.

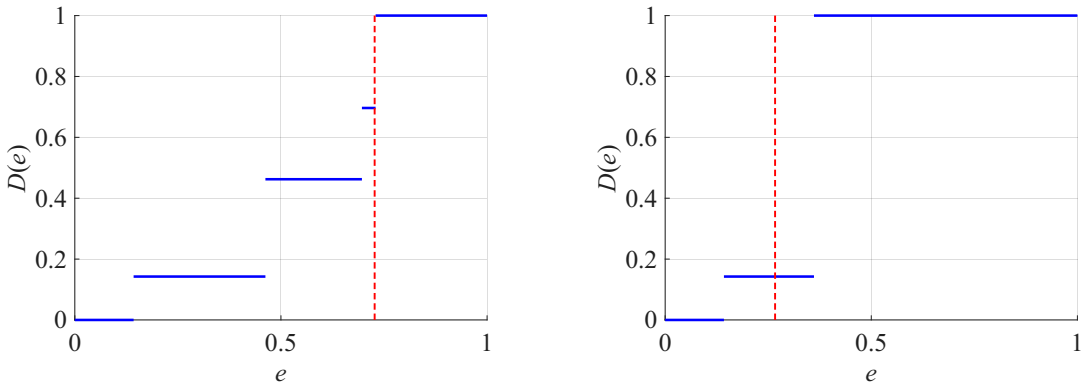
Lemma 2 (Linear densities). *Let $g_\varepsilon(e) = \varepsilon + 2(1-\varepsilon)e$ on $[0,1]$ with $\varepsilon \in (0,1]$. Then φ is strictly decreasing (for $\varepsilon < 1$), the hazard rate is increasing, and $\underline{\lambda} = 2$, while the planner's objective is globally convex on all intervals if and only if*

$$\lambda \;\geq\; \hat{\lambda}(\varepsilon) \;=\; 2 + \frac{2(1-\varepsilon)(v-c)}{\varepsilon}. \tag{34}$$

Equilibrium trade is globally efficient if $\lambda < 2$ and globally excessive if $\lambda > 2$.

Proof. See Appendix. □

At $\varepsilon = 1$ the family collapses to the uniform distribution and $\hat{\lambda} = 2$: the intermediate regime vanishes. As $\varepsilon \downarrow 0$, $\hat{\lambda}(\varepsilon) \to \infty$: the intermediate regime $(2, \hat{\lambda}(\varepsilon))$ becomes arbitrarily wide. Within it, the equilibrium of Proposition 8 exhibits a continuous transition between the two polar regimes. Figure 4 illustrates this transition for $\varepsilon = 0.03$ and $v - c = 0.8$, so that $\hat{\lambda} \approx 53.7$. At $\lambda = 2.1$ (Panel (a)), the equilibrium plan features four trading categories, $[0, 0.143)$, $[0.143, 0.462)$, $[0.462, 0.696)$, and $[0.696, 0.730)$, each limited by the screening bound; trade extends slightly beyond $\hat{e} = 0.727$, stopping at $T = 0.730$. At $\lambda = 4$ (Panel (b)), only two trading categories remain, $[0, 0.143)$ and $[0.143, 0.361)$. The first is again limited by the screening bound, but the terminal one is now limited by the viability bound — as in the severe-exposure regime of the uniform case — and trade extends well beyond $\hat{e} = 0.267$. The lowest cutoff is the same in both panels because, at $\underline{e} = 0$, the screening bound (33) reduces to $G(e_1) = \varepsilon(v-c)$, which does not depend on λ. As exposure rises, disclosure thus becomes coarser, the binding constraint shifts from the screening bound to the viability bound, and the excess of T over $\hat{e}$ widens: the transition from the logic of Proposition 2 to that of Proposition 3 unfolds category by category rather than through a discrete switch.



(a) $\lambda = 2.1$ **(b)** $\lambda = 4$

Figure 4: The figure illustrates the equilibrium disclosure plans in the presence of the optimal domestic tariff scheme, in the format of Figure 3. Externalities follow the linear density $g_\varepsilon(e) = \varepsilon + 2(1-\varepsilon)e$ with $\varepsilon = 0.03$ and $v - c = 0.8$, so that both exposure levels lie in the intermediate regime of Lemma 2. Trade occurs conditional on each of the lower rating categories but does not occur conditional on the topmost category. The dashed vertical line marks the efficiency threshold $\hat{e} = (v-c)/(\lambda - 1)$: in Panel (a), the plan features four trading categories, and trade extends slightly beyond $\hat{e}$; in Panel (b), only two trading categories remain, and trade extends well beyond $\hat{e}$.

4.2 Disclosure Plans Chosen by the Domestic Firm

So far, we have assumed that the disclosure plan is chosen by the foreign government. In this section, we instead study the case when the disclosure plan is chosen by the domestic firm. As before, such a disclosure plan could for example be enforced by an international rating agency that is tasked with measuring and reporting environmental pollution or emissions.

The timing of the game is now as follows:

1. The domestic firm chooses a disclosure plan represented by a function $D : [\underline{e}, \overline{e}] \to S$.
2. The government chooses a tariff schedule that is a function of the possible disclosures S.
3. The foreign firm privately observes e.
4. A signal according to the disclosure plan $D(e)$ is released and becomes publicly known.
5. The domestic firm makes a take-it-or-leave-it offer to the foreign firm.
6. The foreign firm accepts or rejects the offer.
7. International production occurs conditional on an acceptance and payoffs are realized.

Absent interventions by the domestic government, the domestic firm would always choose a perfectly revealing disclosure plan. This is because the domestic firm then can extract all rents by leveraging its market power. However, once the domestic government may intervene with tariffs,

the domestic firm may choose a coarse disclosure plan, in particular for sufficiently high values of λ. The following proposition formalizes this result in the case where externalities are uniformly distributed.

Proposition 9. *Suppose e follows a uniform distribution $U[0,\overline{e}]$ and $\delta = 1$. When the domestic firm chooses the disclosure plan, full disclosure is weakly optimal if and only if $\lambda \leq 2$. For $\lambda > 2$, the optimal plan reveals types below a^* perfectly and pools types in $[a^*, b^*)$ under a single rating, conditional on which trade occurs; no trade occurs for $e \geq b^*$. The pooled rating always extends beyond the efficient threshold, $b^* > \hat{e}$. The cutoffs a^* and b^* are given in closed form in the appendix, equation (A16), together with the characterization of the three parameter regions: an interior plan with $b^* < \overline{e}$; a capped plan with $b^* = \overline{e}$, in which trade occurs for all types; and, for $\lambda \geq 4$ with $2(v-c) \leq \lambda\overline{e}$, the two-rating disclosure plan of Proposition 7.*

The first part of Proposition 9 shows that the domestic firm would not want to choose a disclosure regime that leaves any relevant information asymmetries when the domestic sector has a relatively low exposure to externalities (for $\lambda \leq 2$). When λ is sufficiently small, the domestic firm expects that its government will not impose any tariffs even under full disclosure. In this case, full disclosure is best for the domestic firm as it facilitates extracting all rents from the foreign firm. As λ increases, the domestic firm anticipates that, under full disclosure, its government starts to impose tariffs for high emissions types that render trading with those types privately unprofitable. It is worth noting that in this case, full disclosure is generally only one of many optimal disclosure plans, in the following sense. Suppose, for example, the cutoff of the emission type over which trade does not occur in equilibrium is $\hat{e}$ (because of tariffs being imposed for all firms that are browner than $\hat{e}$). Then a disclosure plan that is perfectly revealing for $e < \hat{e}$ and that pools all types with $e \geq \hat{e}$ is also optimal for the domestic firm. This is because regardless of the disclosure signals for $e > \hat{e}$, the government would impose a tax such that trade does not occur with these types, and the domestic firm's payoff would be identical to that under full disclosure.

Once λ exceeds 2, however, the domestic firm itself begins to coarsen the information it collects. Under full disclosure, its government would eliminate trade with all types above $\hat{e}$; by pooling a range of types around this threshold under a single rating, the domestic firm keeps its government from screening within the pool and thereby preserves trade with the pool as a whole. The optimal plan combines perfect revelation for the greenest types with one pooled trading rating $[a^*, b^*)$ that stretches beyond $\hat{e}$: like the foreign government in Proposition 7, the domestic firm sustains globally excessive trade, now at the top of its own trading range. When gains from trade are large relative to the damage range, the pooled rating absorbs the entire brown tail ($b^* = \overline{e}$): trade never breaks down, and the excess relative to $\hat{e}$ is maximal. As exposure rises further, the revealed range shrinks while the pooled rating widens, and at $\lambda \geq 4$ the plan collapses to the two-rating

plan of Proposition 7, which the foreign firm would choose for any $\lambda > 2$. That is, whenever the domestic firm chooses the binary disclosure plan, the foreign firm would choose the same plan.

In sum, even when the domestic firm can choose the disclosure regime, excessive pooling with coarse ratings emerges in equilibrium, in particular when the domestic sector is sufficiently strongly exposed to externalities.

4.3 Interim Disclosure

In our baseline analysis, we considered the case where the disclosure regime is chosen based on a decision maker that knows the distribution of e but not its specific realization. This assumption is particularly reasonable if one envisions a setting where a foreign government chooses a disclosure regime for its firms ex ante, anticipating the distribution of firms that reside in the foreign country. In contrast, if an individual foreign firm chooses its disclosure plan, it might be reasonable to assume that this firm already has private information when it makes that choice. Correspondingly, we now study the robustness of our results to such an "interim" disclosure game (see also Grossman (1981a) and Milgrom (2008b)).

While multiple equilibria emerge in such a game due to signaling concerns, we will show that the equilibrium outcomes characterized in our baseline analysis also obtain in the case of interim disclosure under sensible equilibrium refinements. The timeline of the game we now study is illustrated in Figure 5.

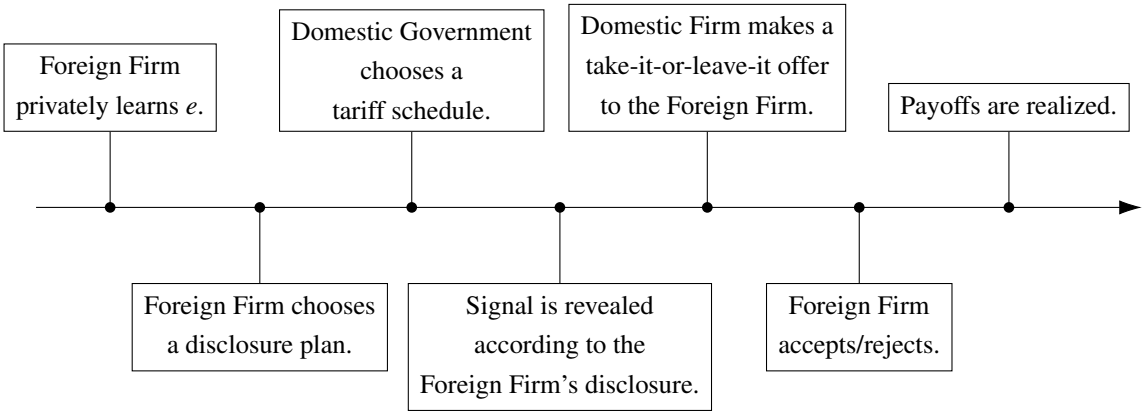


Figure 5: Timeline. The figure illustrates the timeline of the game with interim disclosure.

Since $D(e)$ is now chosen by the foreign firm after it observes e, the disclosure becomes a pure-strategy message in a signaling game. Disclosure operates through the same verifiable, ordered technology as in the baseline analysis: an individual message is an interval of types containing the

sender's type, the interim counterpart of a rating category, and disclosure strategies remain subject to the monotonicity restriction of Assumption 1. Throughout this section we identify a message s with the set of types it discloses: sending s verifiably communicates that $e \in s$. This matches the baseline model, where each signal s corresponds to the set of types $D^{-1}(s)$ that disclose it. Upon observing a signal s, the domestic sector forms beliefs $\mu(s) \in \Delta([\underline{e}, \bar{e}])$ about the foreign firm's type,[13] the domestic firm quotes a price $p(\tau, s)$, and the foreign firm accepts if and only if the price weakly exceeds its production cost $(c - e)$. A perfect Bayesian equilibrium of the interim game consists of a tariff schedule, a disclosure strategy, beliefs, and price offers such that each is sequentially optimal and beliefs obey Bayes' rule on the equilibrium path; the formal definition is in the appendix.

Since beliefs are unrestricted following off-equilibrium deviations, there exist beliefs such that the domestic firm (which has market power) drives the foreign firm's information rents to zero following any off-equilibrium deviation in disclosure. This leads to the existence of multiple perfect Bayesian equilibria with various degrees of information revelation, as opposed to a unique equilibrium with full revelation as in Grossman (1981a) and Milgrom (2008b).[14]

Given this equilibrium multiplicity, we restrict our attention to the sets of equilibria that survive the following two refinements. First, we consider as *foreign-firm-preferred equilibria* the set of equilibria that are not dominated among foreign firm types — in the Pareto sense, with a strict improvement for a positive measure of types — by another equilibrium based on their interim payoffs. Second, we consider an alternative equilibrium refinement known as Grossman-Perry-Farrell, based on the perfect sequential equilibrium of Grossman and Perry (1986) and the neologism-proof equilibrium of Farrell (1993).[15] This refinement is commonly used in models of verifiable disclosure (see, e.g., Bertomeu and Cianciaruso (2018) and the references therein). Instead of comparing equilibria in a Pareto sense as above, this refinement eliminates equilibria with off-equilibrium beliefs that are deemed unreasonable given agents' incentives to deviate from their equilibrium strategies. To streamline the exposition, we defer the formal definition of a Grossman-Perry-Farrell equilibrium to the proof of Proposition 10 in the appendix. We now state our result for the interim disclosure game.

Proposition 10. *An equilibrium disclosure plan of the ex-ante disclosure game can be sustained in*

[13] We use $\Delta([\underline{e}, \bar{e}])$ to denote the set of all possible probability distributions on $[\underline{e}, \bar{e}]$.

[14] For instance, either full disclosure, partial disclosure, or no disclosure can be supported in equilibrium if the seller has the following beliefs: if for any s not in the range of D (that is, whenever s is an off-equilibrium signal), the belief $\mu(s)$ assigns probability 1 to type $\bar{e}(s)$, where $\bar{e}(s) \equiv \sup s$ (recall that a message s is an interval of types). Here, an equilibrium is said to feature full disclosure if $\mu(D(e))$ assigns probability 1 to type e, whereas it is said to feature no disclosure if $D(e) = [\underline{e}, \bar{e}]$ for all $e \in [\underline{e}, \bar{e}]$, and thus $\mu([\underline{e}, \bar{e}])$ is equal to $G(e)$, the prior distribution of e.

[15] We adopt the terminology "Grossman-Perry-Farrell" from Gertner et al. (1988), Lutz (1989), and Bertomeu and Cianciaruso (2018).

both a foreign-firm-preferred equilibrium and a Grossman-Perry-Farrell equilibrium of the interim disclosure game.

Proposition 10 shows that the results of the baseline model apply in the case of interim disclosure under sensible equilibrium refinements. Specifically, the fact that the disclosure plan from our baseline model is a foreign-firm-preferred equilibrium echoes the result that a foreign government would choose this disclosure regime ex ante, anticipating the distribution $G(e)$ of its firms.

5 Conclusion

In this paper, we analyze voluntary environmental disclosures in global supply chains and study a government's optimal reliance on such disclosures in its regulations. Optimal tariff schedules differ from prices in carbon credit markets: rather than charging one common price per disclosed ton of emissions, they are skewed toward the externalities of each rating's greener types, imply effective per-ton charges that rise with disclosed brownness and vary across sectors, and prohibit trade for the brownest disclosures. Penalizing trade with foreign firms that reveal the brownness of their production can make their disclosures more or less precise, depending crucially on the domestic sector's exposure to externalities. In equilibrium, disclosures remain coarse, and under severe exposure they pool excessively, sustaining trade that destroys global surplus.

Our results speak to active regulatory initiatives through one common principle: the regulator's instrument should be calibrated to the residual conditional distribution of externality types within each disclosure category, rather than to the disclosed value alone or to a uniform external price. Applied to the Carbon Border Adjustment Mechanism, this principle implies that charging importers the EU ETS allowance price uniformly per ton of reported (or default) emissions is suboptimal — both across sectors, which differ in their underlying emissions distributions, and within sectors, where the optimal charge is skewed toward each rating's greener supplier types — the opposite end of the disclosed range from CBAM's conservative default values. Our analysis also cautions against expecting mandates to restore granular information: when domestic exposure is severe, coarse and excessively pooled disclosures re-emerge even when domestic parties design the reporting regime.

A Appendix: Proofs

Proof of Lemma 1.

We first record a truncation property of the domestic firm's pricing problem: conditional on any category, the optimal marginal type is unaffected if the category is truncated from below at a point weakly below that optimum. Suppose that conditional on a signal $\tilde{s}$ revealing $e \in [a,b)$, an optimal marginal type is $e^*_{\tilde{s}}$; by equation (13), after scaling by the positive constant $G(b)-G(a)$, it maximizes

$$x \mapsto (\delta v - c + x)\,(G(b) - G(x)) \tag{A1}$$

over $x \in [a,b)$. If the category is truncated from below at t with $a \leq t \leq e^*_{\tilde{s}}$ — revealing $e \in [t,b)$ — the domestic firm's objective becomes the same function scaled by the positive constant $1/(G(b)-G(t))$, maximized over the subset $[t,b)$, which still contains $e^*_{\tilde{s}}$. Hence $e^*_{\tilde{s}}$ remains optimal after the truncation. (Taking $a = \underline{e}$ and $b = \overline{e}$ covers truncations of the unconditional problem (10).)

Now suppose that, contrary to the statement of the lemma, the foreign government chose a disclosure plan that potentially generates a signal conditional on which trade does not occur with types e for which the bilateral surplus is positive, that is, where $e > c - \delta v$. Specifically, the plan includes a possible disclosure $\tilde{s} \in S$ revealing that $e \in [a,b)$ and the optimally chosen marginal type conditional on this disclosure is $e^*_{\tilde{s}} > c - \delta v$. If there are excluded types e satisfying

$$c - \delta v < e < e^*_{\tilde{s}}, \tag{A2}$$

the foreign government can strictly improve the foreign firm's ex ante utility by adjusting the disclosure plan as follows: it can split the original disclosure interval $[a,b)$ into two new disclosure intervals, $[a,e^*_{\tilde{s}})$ and $[e^*_{\tilde{s}},b)$, associated with the new signals $\hat{s}_1$ and $\hat{s}_2$, while keeping the rest of the disclosure plan unchanged. With this adjusted disclosure plan, conditional on a disclosure $\hat{s}_2$ indicating that $e \in [e^*_{\tilde{s}},b)$, the price optimally charged by the domestic firm is still $p = c - e^*_{\tilde{s}}$, implying that all types in the interval $[e^*_{\tilde{s}},b)$ are still trading at the same price under this revised disclosure plan, by the truncation property above.

Moreover, there will be trade with a positive measure of previously excluded types belonging to the interval $[a,e^*_{\tilde{s}})$ at a price exceeding $p = c - e^*_{\tilde{s}}$. Conditional on a signal $\hat{s}_1$ revealing that $e \in [a,e^*_{\tilde{s}})$, the optimal marginal type is a strictly interior element of the interval $[\max\{a, c-\delta v\}, e^*_{\tilde{s}})$. The marginal type will be interior for the following reasons: (1) Offering a price $p = c - e^*_{\tilde{s}}$ would yield the domestic firm a payoff of zero, since the acceptance probability would be zero. (2) Offering a price $p = \delta v$ would yield the domestic firm zero since this is its valuation for the

good. (3) Offering a price in between these two prices (below δv and above $c - e^*_{\hat{s}}$) yields a strictly positive probability of trade (since the PDF $g(\cdot)$ takes strictly positive values everywhere on the domain), while the price is also below δv, and thus will yield the domestic firm a strictly positive expected payoff. At the same time, this (interior) price offer will yield the foreign firm an information rent, since all types above the marginal type (conditional on the disclosure $\hat{s}_1$) will obtain strictly positive surplus. As a result, the foreign firm obtains strictly higher ex ante utility under this revised disclosure plan, implying that the initial supposition is inconsistent with optimality and cannot be an equilibrium. Conversely, trade never occurs with types $e < c - \delta v$: serving such a type requires a price $p > \delta v$, and any offer $p > \delta v$ that is accepted with positive probability yields the domestic firm a strictly negative expected payoff, so the domestic firm never makes such an offer. □

Explicit expressions for the efficiency bounds and the Φ-function. We first record the explicit expressions for the objects defined economically in Section 3.1. The left and right efficiency bound functions $lb, rb : [\hat{e}^o, \overline{e}] \to \mathbb{R}$ are given by

$$lb(e) \equiv \begin{cases} (\Pi')^{-1}(1 - G(e)), & \text{for } \Pi'(\hat{e}^o) > 1 - G(e), \\ \hat{e}^o, & \text{otherwise,} \end{cases} \tag{A3}$$

$$rb(e) \equiv \begin{cases} G^{-1}(1 - \Pi'(e)), & \text{for } \Pi'(e) > 0, \\ \overline{e}, & \text{for } \Pi'(e) \leq 0. \end{cases} \tag{A4}$$

Here, $(\Pi')^{-1}$ is well-defined since $\Pi'(\cdot)$ is strictly decreasing, and the two functions are inverses of each other on the relevant domain: $rb(lb(e)) = e$ and $lb(rb(e)) = e$. The function $\Phi : \Omega \to \mathbb{R}$ is given by

$$\Phi(u, w) \equiv [G(w) - G(rb(u))]\,[rb(u) - u] - [G(w) - G(lb(w))]\,[lb(w) - u], \tag{A5}$$

with domain

$$\Omega \equiv \{(u, w) : u \in (lb(lb(w)), lb(w)] \text{ and } w \in (rb(\hat{e}^o), \overline{e}]\}. \tag{A6}$$

Proof of Proposition 1. We first show the following lemma.

Lemma 3. *Given any two partition cutoffs e_{i-1} and e_{i+1} with $i \geq 2$ that satisfy the constraint $e_{i-1} \geq lb(lb(e_{i+1}))$, the interior cutoff e_i that maximizes the foreign firm's expected payoff over the feasible range $[lb(e_{i+1}), rb(e_{i-1})]$ is either equal to $rb(e_{i-1})$ or equal to $lb(e_{i+1})$.*

Proof. The marginal effect of increasing e_i on the foreign firm's expected payoff equals the partial derivative of the expected price offered by the domestic firm with respect to e_i:

$$\frac{\partial\left(\sum_{j=1}^{n}\left(G(e_{j+1})-G(e_j)\right)(c-e_j)\right)}{\partial e_i} = -(G(e_{i+1})-G(e_i))+g(e_i)(e_i-e_{i-1}). \tag{A7}$$

We first show that for any two partition cutoffs e_{i-1} and e_{i+1} with $i \geq 2$ that satisfy $e_{i-1} \geq lb(lb(e_{i+1}))$, the expression (A7) is strictly increasing in e_i, so that it crosses zero at most once, and from below. Differentiating (A7) with respect to e_i yields

$$2g(e_i)+g'(e_i)(e_i-e_{i-1}). \tag{A8}$$

Recall our maintained restriction that $2g(x)+g'(x)(\delta v-c+x)>0$ for all x. Moreover, all partition cutoffs of the trading region satisfy $e_{i-1} \geq \hat{e}^o \geq c-\delta v$, so that $0 \leq e_i-e_{i-1} \leq \delta v-c+e_i$. If $g'(e_i) \geq 0$, expression (A8) is trivially positive; if $g'(e_i)<0$, then $g'(e_i)(e_i-e_{i-1}) \geq g'(e_i)(\delta v-c+e_i)$, and positivity follows directly from the maintained restriction. Hence the expected price offered by the domestic firm is strictly convex in e_i on the relevant range.

By this strict convexity, the expected price reaches its maximum over $[lb(e_{i+1}), rb(e_{i-1})]$ at one of the two endpoints: $e_i = rb(e_{i-1})$ or $e_i = lb(e_{i+1})$. The range is nonempty under the lemma's hypothesis: since $rb(\cdot)$ is increasing and $rb(lb(x))=x$, the constraint $e_{i-1} \geq lb(lb(e_{i+1}))$ implies $rb(e_{i-1}) \geq lb(e_{i+1})$. □

We now show that the plan described in the proposition is optimal. Its cutoffs are generated by

$$a_1 = \overline{e}, \qquad a_{m+1} = lb(a_m).$$

Step 1: reduction. By Lemma 1 and its splitting argument, any optimal plan induces trade exactly for $e \geq \hat{e}^o$, with trade occurring with probability one conditional on each trading interval; types below $\hat{e}^o$ are pooled in a non-trading interval that does not affect payoffs. It therefore suffices to compare plans whose trading intervals partition $[\hat{e}^o, \overline{e}]$ and each sustain full trade, that is, whose consecutive cutoffs satisfy

$$e_{i+1} \leq rb(e_i), \qquad \text{equivalently} \qquad e_i \geq lb(e_{i+1}).$$

We call such plans feasible in the remainder of the proof.

Step 2: removing a cutoff strictly raises the foreign firm's payoff, whenever the plan remains feasible. Conditional on a trading interval, the domestic firm offers the price c minus the interval's lower endpoint, so each type's payoff equals its distance to that endpoint. If a cutoff y is removed

— where $u < y < x$ denote its neighboring cutoffs — and the merged interval $[u,x)$ sustains full trade, that is, $x \le rb(u)$, then every type in $[y,x)$ strictly gains and no type loses.

Step 3: the sequence (a_m). If $c - \delta v < \underline{e}$, so that $\hat{e}^o = \underline{e}$, there is a first index K with $lb(a_K) \le \underline{e}$, equivalently $a_K \le rb(\underline{e})$: the sequence is strictly decreasing, and if it never crossed $\underline{e}$ it would converge to a fixed point of lb weakly above $\underline{e}$, whereas the only fixed point of lb is $c - \delta v < \underline{e}$. The proposed plan then has the cutoffs $a_K < \cdots < a_1$. Otherwise $\hat{e}^o = c - \delta v \ge \underline{e}$; then $lb(x) > \hat{e}^o$ for every $x > \hat{e}^o$, because

$$\Pi'(\hat{e}^o) = 1 - G(\hat{e}^o) > 1 - G(x)$$

places lb in its interior branch, so the sequence is infinite and converges to $\hat{e}^o$, the unique fixed point of lb. In this case every feasible plan likewise has infinitely many cutoffs: a bottommost trading interval $[\hat{e}^o, y)$ would require $y \le rb(\hat{e}^o)$, whereas

$$rb(\hat{e}^o) = G^{-1}\big(1 - \Pi'(\hat{e}^o)\big) = \hat{e}^o.$$

Step 4: an optimal plan exists. Assign to each type e the lower endpoint $\ell(e)$ of its trading interval. The functions ℓ arising from feasible plans are exactly the nondecreasing functions with $\ell(e) \le e$ that are constant on $[\ell(e), e]$ for every e and satisfy $e \le rb(\ell(e))$; these properties are preserved under pointwise limits, the last by continuity of rb. Since the functions ℓ are monotone and uniformly bounded, any sequence of feasible plans admits a subsequence along which they converge pointwise (Helly's selection theorem), and the foreign firm's expected payoff,

$$\int_{\hat{e}^o}^{\overline{e}} \big(e - \ell(e)\big)\, dG(e),$$

converges along it by dominated convergence. A payoff-maximizing feasible plan therefore exists.

Step 5: every optimal plan can be aligned with the proposed plan from the top down. Fix an optimal plan. By Step 2, none of its cutoffs is removable: a cutoff y with upper neighbor x, and with u denoting the cutoff below y or $\underline{e}$ if there is none, must satisfy

$$x > rb(u), \qquad \text{equivalently} \qquad u < lb(x).$$

Suppose the plan's cutoffs above a_m coincide with $a_1, \ldots, a_m$; for $m = 1$ this holds trivially, since every plan's top interval ends at $\overline{e} = a_1$.

In the finite case with $m = K$, the plan has no cutoffs below a_K: the highest such cutoff, with upper neighbor a_K and with $u \ge \underline{e}$ below it, would satisfy $a_K \le rb(\underline{e}) \le rb(u)$ and hence be removable. The plan thus coincides with the proposed plan, completing the finite case.

If instead m is not terminal, let y be the plan's highest cutoff below a_m, and let u be the cutoff below y or $\underline{e}$ if there is none. Feasibility gives $y \geq lb(a_m) = a_{m+1}$; suppose $y > a_{m+1}$. Non-removability of y gives $u < lb(a_m)$, and feasibility of the interval $[u, y)$ gives

$$u \;\geq\; lb(y) \;\geq\; lb(a_{m+1}),$$

which is the hypothesis of Lemma 3 for the neighbors u and a_m: the payoff, as a function of the position of the cutoff between them, is maximized at one of the two endpoints of its range of feasible positions, $lb(a_m)$ or $rb(u)$ — the latter being the range's upper endpoint since $rb(u) < a_m$. By Assumption 3, the position $lb(a_m)$ is weakly optimal. Since the plan is optimal, moving y to $a_{m+1} = lb(a_m)$ therefore leaves its payoff unchanged; the interval $[u, a_{m+1})$ remains feasible by the display above. The modified plan is again optimal — so again has no removable cutoffs — and its cutoffs down to a_{m+1} are $a_1, \ldots, a_{m+1}$.

Step 6: the infinite case. Iterating Step 5 yields, for every m, an optimal plan whose cutoffs down to a_{m+1} are $a_1, \ldots, a_{m+1}$. Its payoff from types above a_{m+1} coincides with that of the proposed plan, while types below a_{m+1} contribute at least zero and at most

$$G(a_{m+1})\left(a_{m+1} - \hat{e}^o\right) \;\longrightarrow\; 0,$$

since each type's payoff is bounded by the width of its interval. Letting $m \to \infty$, the proposed plan's payoff equals the optimal value: the proposed plan is optimal. □

Lemma 4 (Φ-positivity in the uniform case)**.** *If externalities are uniformly distributed on $[\underline{e}, \overline{e}]$, Assumption 3 is satisfied, with equality only at the boundary of the domain Ω.*

Proof. For uniform G, normalizing $\underline{e} = 0$ and $\overline{e} = 1$ (which is without loss for the sign of Φ) and writing $k \equiv \delta v - c$, we have $lb(e) = (e - k)/2$ and $rb(e) = 2e + k$. Substituting into the definition (A5) and simplifying yields

$$4\Phi(u, w) = (w - 2u - k)(4u + 3k - w).$$

The first factor is positive for $u < lb(w) = (w - k)/2$ and vanishes at $u = lb(w)$; the second factor is positive for $u > lb(lb(w)) = (w - 3k)/4$ and vanishes there. Since the domain Ω restricts u to $(lb(lb(w)), lb(w)]$, both factors are non-negative on Ω, so $\Phi \geq 0$, with equality only at the boundary. □

Proof of Proposition 2. Given a generic disclosure that implies $e \in [a, b)$, we characterize the marginal type induced by the domestic planner solution. Conditional on the disclosure, the plan-

ner's problem is

$$\max_{p} \frac{G(b)-G(c-p)}{G(b)-G(a)} \left[v-p-\lambda\, \mathbb{E}[e \mid e \geq c-p,\, e \in [a,b)] \right].$$

Given the marginal type $x \equiv c-p$, the problem can be rewritten as

$$\max_{x} \frac{G(b)-G(x)}{G(b)-G(a)} \left[v-c+x-\lambda\, \mathbb{E}[e \mid e \geq x,\, e \in [a,b)] \right],$$

and equivalently,

$$\max_{x} (G(b)-G(x))\,[v-c+x] - \lambda \int_{x}^{b} e dG(e).$$

Let

$$d(x) = (G(b)-G(x))\,[v-c+x] - \lambda \int_{x}^{b} e dG(e). \tag{A9}$$

We note that $d(x)$ does not depend on a. Thus, the planner's induced allocation does not rely on the left end of the disclosure interval a. This implies that under the optimal disclosure plan, trade either occurs with probability 1 or 0. To see this, we argue by contradiction. Suppose trade occurs with a probability strictly between 0 and 1 conditional on a category $[a,b)$, so that the planner induces an interior marginal type $x^* \in (a,b)$. Interior optimality requires $d'(x^*) = 0$, that is,

$$G(b)-G(x^*) \;=\; g(x^*)\,[v-c+(1-\lambda)x^*];$$

since $G(b)-G(x^*) > 0$, it follows that $v-c+(1-\lambda)x^* > 0$, that is, $x^* < \hat{e}$. Now consider the finer disclosure plan that splits the category into $[a,x^*)$ and $[x^*,b)$ and keeps everything else the same. Conditional on $[x^*,b)$, the planner faces the objective $d(\cdot)$ restricted to marginal types in $[x^*,b)$; since x^* maximizes d over the larger set $[a,b)$, it remains a maximizer over the subset, so the planner induces the same allocation (selecting x^* under the maintained convention that an indifferent planner chooses the lowest marginal type) and the foreign firm's payoff from these types is unchanged. Conditional on the new category $[a,x^*)$, the planner's payoff from inducing a marginal type x slightly below x^* is

$$d(x;x^*) \;=\; (G(x^*)-G(x))\,[v-c+x] - \lambda \int_{x}^{x^*} e\, dG(e) \;\approx\; (x^*-x)\, g(x^*)\,[v-c+(1-\lambda)x^*] \;>\; 0,$$

whereas inducing no trade yields 0: the planner's optimal response therefore features trade with strictly positive probability, with some marginal type $\tilde{x} < x^*$. All types in $(\tilde{x},x^*)$ — a positive-measure set — then collect strictly positive rents, whereas they collected none under the original

plan, and no type is made worse off: a strict improvement for the foreign government, contradicting the optimality of the original plan.

We are left to show that when trade occurs, the ex post global surplus must be nonnegative. Suppose otherwise: under an optimal disclosure plan, there is a signal $[a,b)$ under which trade occurs with probability 1 but the global surplus at b is negative. Computing the derivative of d with respect to x, we have

$$d'(x) = G(b) - G(x) - g(x)\left[v - c + (1-\lambda)x\right].$$

Since the global surplus at b is negative, $d'(b)$ is weakly positive. Moreover, since trade occurs with probability 1 conditional on this signal, the induced marginal type is $x = a$, and optimality requires $d'(a) \leq 0$; otherwise the planner would gain by inducing a slightly higher marginal type. Let $h(x) = d'(x)/g(x)$, then $h(a) \leq 0$ and $h(b) \geq 0$. By Lagrange's mean value theorem, there exists x_0 such that

$$h'(x_0) = \frac{h(b) - h(a)}{b - a} \geq 0.$$

Note that

$$h'(x) = \left(\frac{G(b) - G(x)}{g(x)}\right)' - (1-\lambda).$$

Plugging into x_0, it follows that

$$\lambda \geq 1 - \left(\frac{(G(b) - G(x_0))}{g(x_0)}\right)'. \tag{A10}$$

We next prove the following lemma.

Lemma 5. *Suppose the hazard rate $\frac{g(e)}{1-G(e)}$ is an increasing function. Let $\tilde{G}$ denote a truncated version of distribution G, and let $\tilde{g}$ denote the associated probability density function. Then $\frac{\tilde{g}(e)}{1-\tilde{G}(e)}$ is also an increasing function.*

Proof. Suppose the truncation is denoted by $[a,b]$, then

$$\tilde{G}(e) = \frac{G(e) - G(a)}{G(b) - G(a)}, \qquad \tilde{g}(e) = \frac{g(e)}{G(b) - G(a)}.$$

Note that

$$\frac{1 - G(e)}{G(b) - G(e)} = 1 + \frac{1 - G(b)}{G(b) - G(e)},$$

which is an increasing function of e. Then,

$$\frac{\tilde{g}(e)}{1 - \tilde{G}(e)} = \frac{g(e)}{G(b) - G(e)} = \frac{g(e)}{1 - G(e)} \cdot \frac{1 - G(e)}{G(b) - G(e)},$$

which is increasing because it is equal to the product of two positive increasing functions. □

Returning to the proof of the proposition, let

$$\underline{\lambda} = \inf\left\{1 - \left(\frac{G(b) - G(x)}{g(x)}\right)'\right\},$$

where the infimum is taken over all pairs (b,x) with $\underline{e} \leq x < b \leq \overline{e}$, as stated in equation (22). Since $(G(b) - G(x))/g(x)$ is decreasing by Lemma 5, the inner derivative is nonpositive and hence $\underline{\lambda} \geq 1$; the threshold strictly exceeds 1 whenever that derivative is bounded away from zero uniformly over the relevant pairs (b,x), in which case the moderate-exposure range $(1, \underline{\lambda})$ is nonempty. In the leading uniform example, $g(\cdot)$ is constant and $(G(b) - G(x))/g(x) = b - x$, with derivative -1, so $\underline{\lambda} = 2$. Now suppose $\lambda < \underline{\lambda}$. It follows that the ex post global surplus must be nonnegative whenever trade occurs, because otherwise (A10) would be violated: the mean-value point x_0 constructed above would deliver

$$\lambda \;\geq\; 1 - \left(\frac{G(b) - G(x_0)}{g(x_0)}\right)' \;\geq\; \underline{\lambda},$$

contradicting $\lambda < \underline{\lambda}$. Conversely, trade cannot break down for types generating strictly positive global surplus. If a positive measure of types below $\hat{e}$ did not trade in equilibrium, the foreign government could split off a small category $[x, x+\varepsilon)$ at the bottom of the excluded region. For small ε, the planner's payoff from full trade on this category is strictly positive,

$$d(x; x+\varepsilon) \;\approx\; \varepsilon\, g(x)(\lambda - 1)(\hat{e} - x) \;>\; 0,$$

so the planner's optimal response on the new category features trade with strictly positive probability, and the foreign firm collects strictly positive rents from the types above the induced marginal type — a profitable deviation, by the same splitting argument as in Lemma 1. Equilibrium trade therefore occurs if and only if global surplus is positive. This concludes the proof of the proposition. □

Proof of Proposition 3.

As in the proof of Proposition 2, conditional on a generic disclosure revealing $e \in [a,b)$, the planner's problem reduces to choosing the marginal type x to maximize

$$d(x) = (G(b) - G(x))\,[v - c + x] - \lambda \int_x^b e\, dG(e),$$

which is linear in λ. We have

$$\frac{\partial^3 d}{\partial x^2 \partial \lambda} = \frac{\partial (xg(x))}{\partial x} > 0.$$

This is to say that for sufficiently large λ, $d(x)$ is convex in x. The hypothesis makes the required threshold finite and uniform. Writing

$$d''(x) \;=\; \lambda \frac{d}{dx}\big[xg(x)\big] \;-\; \big[2g(x) + g'(x)(v - c + x)\big],$$

the derivative of $xg(x)$ is continuous and strictly positive on the compact support and hence attains a strictly positive minimum, while the bracketed term is bounded; consequently there is a finite $\hat{\lambda} > 2$ such that $d''(x) \geq 0$ for every x, every conditioning interval, and every $\lambda \geq \hat{\lambda}$.

Returning to the proof of the proposition, we argue by contradiction. Suppose under the foreign government's optimal disclosure plan, trade is efficient. That is, the disclosure plan separates types above $\hat{e}$ and types below $\hat{e}$. Assume that one of the disclosures reveals $e \in [e', \hat{e})$, where $e' \geq \underline{e}$. We shall consider a small perturbation in the disclosure plan such that F earns a higher ex ante expected payoff, which leads to a contradiction. For ease of exposition, we let $e' = \underline{e}$ so that there are two disclosures: $\{[\underline{e}, \hat{e}), [\hat{e}, \overline{e}]\}$. The argument extends without change to any efficient disclosure plan that subdivides the trading region $[\underline{e}, \hat{e})$ into finitely many sub-intervals: in that case, the perturbation we construct below acts on the rightmost trading sub-interval $[e_{K-1}, \hat{e})$, replacing it with $[e_{K-1}, \hat{e} + \varepsilon)$ while leaving all sub-intervals to its left and the no-trade interval to its right unchanged. The proof then proceeds verbatim with e_{K-1} in place of $\underline{e}$ and shows that F strictly gains because types in $[\hat{e}, \hat{e} + \varepsilon)$ now earn strictly positive payoffs. Consider the following disclosure plan: $\{[\underline{e}, \hat{e} + \varepsilon), [\hat{e} + \varepsilon, \overline{e}]\}$. When $e \in [\hat{e} + \varepsilon, \overline{e}]$, there is no trade due to negative global surplus, so nothing changes compared to the original disclosure plan.

Next, we claim that we can choose ε so that the planner still induces trading for all $e \in [\underline{e}, \hat{e} + \varepsilon)$. Setting $a = \underline{e}, b = \hat{e} + \varepsilon$, the planner's expected payoff function $d(x)$ is convex for large enough λ. So the planner induces either $x = \hat{e} + \varepsilon$ or $x = \underline{e}$. Since inducing $x = \hat{e} + \varepsilon$ implies no trade, it implies that $d(\hat{e} + \varepsilon) = 0$. If we choose ε small enough, $d(\underline{e}) > 0$. This is because, in the original plan, full trade on $[\underline{e}, \hat{e})$ is the planner's optimal response, so

$$d(\underline{e}) \;=\; \max_{x \in [\underline{e}, \hat{e}]} d(x),$$

where d is computed with $b = \hat{e}$. Substituting $e = x + t(\hat{e} - x)$,

$$d(x) \;=\; (\hat{e} - x)^2 \int_0^1 \big[(\lambda - 1) - \lambda t\big]\, g\big(x + t(\hat{e} - x)\big)\, dt,$$

and the integral converges to $g(\hat{e})(\lambda - 2)/2 > 0$ as $x \uparrow \hat{e}$, where $\lambda > 2$ since $\lambda \geq \hat{\lambda} > 2$. Hence

$d(x) > 0$ for x sufficiently close to $\hat{e}$, and therefore $d(\underline{e}) > 0$ strictly. By a continuity argument, when ε is small enough, $d(\underline{e}) > 0$ under the new disclosure plan.

Now, F's payoff is strictly higher under the new disclosure plan than under the original plan, since types in $[\hat{e}, \hat{e}+\varepsilon)$ are earning strictly positive payoffs. Overall, the original disclosure plan cannot be the equilibrium outcome.

It remains to rule out efficient disclosure plans with *infinitely many* trading categories. For such a plan, either some trading category has right endpoint $\hat{e}$ — a rightmost trading category exists after all, and the perturbation above applies to it verbatim — or the trading categories' cutoffs accumulate at $\hat{e}$ from below, so that no rightmost trading category exists. In the latter case, pick a cutoff a_K of the plan close enough to $\hat{e}$ that $d(a_K; \hat{e}) > 0$; this is possible because, by the substitution used above,

$$d(a;\hat{e}) \;=\; (\hat{e}-a)^2 \int_0^1 \big[(\lambda-1)-\lambda t\big]\, g\big(a+t(\hat{e}-a)\big)\, dt,$$

with the integral converging to $g(\hat{e})(\lambda-2)/2 > 0$ as $a \uparrow \hat{e}$. By continuity in the right endpoint, $d(a_K; \hat{e}+\varepsilon) > 0$ for ε small enough. Now consider the deviation that merges the entire tail of trading categories contained in $[a_K, \hat{e})$, together with $[\hat{e}, \hat{e}+\varepsilon)$, into the single category $[a_K, \hat{e}+\varepsilon)$, leaving all other categories unchanged. Since d is convex on the merged category, the planner's response is a corner, and $d(a_K; \hat{e}+\varepsilon) > 0$ makes full trade optimal. Every type in $[a_K, \hat{e})$ now trades at the price $c - a_K$, weakly above the price it received under the original plan (whose category lower endpoints in the tail all lie weakly above a_K), and types in $[\hat{e}, \hat{e}+\varepsilon)$ newly trade and earn strictly positive rents $e - a_K > 0$. The deviation strictly raises the foreign government's ex-ante payoff, so no efficient plan with infinitely many trading categories can be an equilibrium either.

So the equilibrium cannot feature efficient trade. The equilibrium also cannot feature inefficiently low trade: if a positive measure of types below $\hat{e}$ were excluded from trade, the foreign government could split off a small category at the bottom of the excluded region; the planner's payoff from trade on such a category is strictly positive (as in the proof of Proposition 2), so trade would occur there and the foreign firm would collect strictly positive rents — a profitable deviation. As a result, the only possibility is that equilibrium features inefficiently excessive trade. □

Proof of Proposition 4. Conditional on a category $[a, b)$, the government implements the domestic planner solution: the induced marginal type maximizes

$$d(x;b) = (G(b) - G(x))(v - c + x) - \lambda \int_x^b e\, dG(e)$$

over $x \in [a, b]$, where $x = b$ corresponds to no trade and $d(b; b) = 0$; as noted in the proof of Proposition 2, d does not depend on the category's left endpoint. Throughout, we state the tie-

breaking convention directly in terms of the planner's chosen response: a planner indifferent among marginal types selects the lowest one, that is, the response with the greatest trade probability. Since the proposition assumes that the government implements the planner solution conditional on each disclosure, the convention selects among the planner's optimal responses directly, without relying on a selection among implementing tariffs; among the tariffs implementing the selected response, we report the one of smallest absolute value.

Step 1: no trade at marginal types above $\hat{e}$. Writing

$$d(x;b) = \int_x^b (v - c + x - \lambda e)\, dG(e), \qquad v - c + x - \lambda e = [v - c + (1-\lambda)x] + \lambda(x - e),$$

both bracketed terms are non-positive for $x \geq \hat{e}$ and $e \geq x$, and the second is strictly negative for $e > x$. Hence $d(x;b) < 0$ for all $x \geq \hat{e}$ and $b > x$: any planner response featuring trade has marginal type strictly below $\hat{e}$, and no category contained in $[\hat{e}, \overline{e}]$ features trade.

Step 2: splitting. Suppose, toward a contradiction, that in equilibrium some on-path category $[a,b)$ induces trade with probability strictly between zero and one; that is, the planner's response is a marginal type $x^* \in (a,b)$ with $d(x^*;b) \geq 0$, and $x^* < \hat{e}$ by Step 1. Consider the foreign government's deviation that splits this category into $[a,x^*)$ and $[x^*,b)$, leaving all other categories unchanged. On the upper category, the planner maximizes $d(\cdot\,;b)$ over $[x^*,b]$; since x^* maximized $d(\cdot\,;b)$ over the superset $[a,b]$, it remains optimal (selected under our convention if ties arise), and the induced outcome — trade with all types in $[x^*,b)$ at price $c - x^*$ — is unchanged. On the lower category, the planner maximizes $d(\cdot\,;x^*)$ over $[a,x^*]$, where $d(x^*;x^*) = 0$ and

$$\left.\frac{\partial d(x;x^*)}{\partial x}\right|_{x=x^*} = -\,g(x^*)\,[v - c + (1-\lambda)x^*] < 0,$$

the inequality holding because $x^* < \hat{e}$. Hence $d(x;x^*) > 0$ for x slightly below x^*, so the planner's optimal response on $[a,x^*)$ induces trade with strictly positive probability at some marginal type $\tilde{x} < x^*$. The foreign firm's payoff is unchanged for types in $[x^*,b)$, unchanged in all other categories, and strictly higher for the positive measure of types in $(\tilde{x},x^*)$, which now trade at price $c - \tilde{x}$ and earn $e - \tilde{x} > 0$. The deviation strictly raises the foreign government's ex-ante payoff, contradicting equilibrium. □

Proof of Proposition 5.

By hypothesis, the planner solution induces an interior marginal type $e_{\tilde{s}}^{**} \in (a,b)$, so the planner's first-order condition (25) holds; multiplying it by $G(b) - G(a)$ expresses it in unconditional terms:

$$G(b) - G(e_{\tilde{s}}^{**}) = (v - c + (1-\lambda)e_{\tilde{s}}^{**})\, g(e_{\tilde{s}}^{**}).$$

Consider the domestic firm's problem conditional on the disclosure $\tilde{s}$ when there is a tariff $\tau = \lambda e^{**}_{\tilde{s}} - (1-\delta)v$: it chooses the marginal type $x \in [a,b)$ to maximize $(\delta v - c - \tau + x)(1 - G_{\tilde{s}}(x))$, as in program (13) with the tariff added. Multiplying by $G(b) - G(a)$, the associated first-order condition is

$$G(b) - G(e) = (\delta v - c - \tau + e)\, g(e) = (v - c - \lambda e^{**}_{\tilde{s}} + e)\, g(e).$$

Evaluating at $e = e^{**}_{\tilde{s}}$, the right-hand side equals $(v - c + (1-\lambda)e^{**}_{\tilde{s}})\, g(e^{**}_{\tilde{s}})$, so $e = e^{**}_{\tilde{s}}$ satisfies the condition by the planner's first-order condition above. For uniqueness, rewrite the condition as

$$H_b(e) \;\equiv\; \frac{G(b) - G(e)}{g(e)} - e \;=\; \delta v - c - \tau.$$

Since the hazard rate of G is increasing, Lemma 5 implies that the hazard rate of the distribution truncated to $[a,b]$ is increasing as well, so $(G(b) - G(e))/g(e)$ is decreasing; since H_b subtracts from it the strictly increasing term e, H_b is strictly decreasing on $[a,b)$. The first-order condition therefore has at most one solution; moreover, the marginal profit $\Pi'_{\tilde{s}}(e)$ has the same sign as the difference $H_b(e) - (\delta v - c - \tau)$, which crosses zero once from above, so that solution is the domestic firm's optimum. Overall, with the tariff $\tau = \lambda e^{**}_{\tilde{s}} - (1-\delta)v$, the domestic firm optimally charges a price that induces exactly the threshold $e^{**}_{\tilde{s}}$ above which foreign firm types accept. □

Corollary 3 (δ-invariance of the optimal mechanism)**.** *Suppose Assumptions 1 and 4 hold and that the domestic regulator implements the optimal disclosure-conditional tariff schedule. Then:*

(i) The marginal foreign-firm type $e^{**}_{\tilde{s}}$ induced by the planner solution conditional on any disclosure $\tilde{s}$ is independent of δ.

(ii) The foreign government's optimal disclosure plan $D(\cdot)$ is independent of δ.

(iii) At interior planner solutions, the implementing tariff (26) satisfies $\partial\tau(\tilde{s})/\partial\delta = v > 0$; at corner solutions, the interval of implementing tariffs shifts up one-for-one with δv.

Proof. The planner's first-order condition for the marginal type, equation (25), depends on v, c, λ, and the conditional distribution $G_{\tilde{s}}$, but not on δ. Likewise, the bang-bang criterion at corner solutions depends on

$$v - \lambda\, \mathbb{E}[e \mid e \in [a,b)] - (c - a),$$

which does not involve δ. Hence (i). Under the implementing tariff $\tau(\tilde{s}) = \lambda e^{**}_{\tilde{s}} - (1-\delta)v$, the domestic firm's effective valuation conditional on trade is

$$\delta v - \tau(\tilde{s}) = \delta v - \lambda e^{**}_{\tilde{s}} + (1-\delta)v = v - \lambda e^{**}_{\tilde{s}},$$

which is independent of δ. The price the domestic firm offers and the resulting probability of trade are functions of this effective valuation, so the foreign firm's expected payoff under any disclosure plan is also independent of δ. Therefore the foreign government's optimization over disclosure plans yields a δ-invariant solution, establishing (ii). For (iii), the interior case follows by direct differentiation of equation (26); at corners, every bound of the admissible tariff interval in Proposition 6 contains δv additively, so the interval shifts by $+v$ per unit increase in δ. (A tariff selected from the interior of such an interval by the lowest-absolute-value convention, e.g. $\tau = 0$, may remain locally unchanged as δ varies.) □

For the next result, suppose the domestic firm's objective conditional on trade is $\delta v - p - \tau - \alpha\lambda e$ for an internalization share $\alpha \in [0,1]$, with $\alpha = 0$ recovering our baseline.

Corollary 4 (Partial internalization by the domestic firm)**.** *Suppose the conditions of Proposition 5 hold and the domestic firm internalizes a share $\alpha \in [0,1]$ of the household externality. The domestic planner solution is unchanged, and it is implemented by the tariff*

$$\tau = (1-\alpha)\,\lambda\, e^{**}_{\tilde{s}} - (1-\delta)\,v, \tag{A11}$$

provided the firm's marginal profit under the modified objective still crosses zero once from above (for the uniform distribution, $\alpha\lambda < 2$ suffices). The optimal tariff is thus reduced one-for-one by the share of the externality the firm already internalizes.

Proof. The planner's objective, $v - p - \lambda e$ conditional on trade, does not involve α, so the planner solution $e^{**}_{\tilde{s}}$ is unchanged. Given the tariff above and a disclosure revealing $e \in [a,b)$, the firm chooses the marginal type x to maximize

$$(\delta v - c - \tau + x)\,[G(b) - G(x)] - \alpha\lambda \int_x^b e\, dG(e),$$

with first-order condition

$$G(b) - G(x) = g(x)\,[\delta v - c - \tau + (1-\alpha\lambda)\,x]\,.$$

Substituting the stated tariff, the right-hand side at $x = e^{**}_{\tilde{s}}$ equals

$$g(e^{**}_{\tilde{s}})\,[v - c + (1-\lambda)e^{**}_{\tilde{s}}]\,,$$

which is the planner's first-order condition (25); since the marginal profit crosses zero once from above under the stated condition, this solution is the firm's optimum. □

Proof of Proposition 6. We first consider the domestic planner solution when $\lambda < 2$:

$$p = \arg\max_p (v - \lambda\, \mathbb{E}[e|e \geq c-p, e \in [a,b)] - p)(b-(c-p)).$$

Since $\lambda < 2$, from the first order condition we find that $e^{**} = \frac{b-(v-c)}{2-\lambda}$. We first assume $e^{**} \in (a,b)$, that is, $a < \frac{b-(v-c)}{2-\lambda} < b$. Suppose a tariff is imposed so that D's payoff is $\delta v - p - \tau$ when there is a trade. Following the same argument, the firm's optimal threshold is $\frac{b-(\delta v-\tau-c)}{2}$. Setting

$$\frac{b-(\delta v-\tau-c)}{2} = \frac{b-(v-c)}{2-\lambda},$$

we obtain

$$\tau = \frac{\lambda(b-\delta v+c) - 2(1-\delta)v}{2-\lambda}.$$

We now discuss the cases $e^{**} \leq a$ and $e^{**} \geq b$.

- Suppose $e^{**} \leq a$, or $a \geq \frac{b-(v-c)}{2-\lambda}$. D's optimal price under the domestic planner solution is given by $c-a$, in which case the cut-off type is a. Since when a tariff is imposed the firm's optimal threshold is $\frac{b-(\delta v-\tau-c)}{2}$, government can implement the same outcome by imposing a tariff such that $\frac{b-(\delta v-\tau-c)}{2} = a$, equivalently, $\tau = 2a-b+\delta v-c$.
- Suppose $e^{**} \geq b$, or $\frac{b-(v-c)}{2-\lambda} \geq b$. D's optimal price under the domestic planner solution is given by $c-b$, in which case the threshold is b. Since when a tariff is imposed the firm's optimal threshold is $\frac{b-(\delta v-\tau-c)}{2}$, government can implement the same outcome by imposing a tariff such that $\frac{b-(\delta v-\tau-c)}{2} = b$, equivalently, $\tau = \delta v - c + b$.

Next consider $\lambda > 2$. The domestic planner solution is given by the following program.

$$p = \arg\max_{e^{**}=c-p} (v - \lambda\, \mathbb{E}[e|e \geq e^{**}, e \in [a,b)] - p)(b-e^{**}).$$

The objective is a quadratic function of e^{**} with coefficient $(\lambda-2)/2 > 0$ on $(e^{**})^2$, and is thus convex in e^{**}. So the optimal e^{**} is either a or b. There are two possibilities.

- If $v - \lambda\frac{a+b}{2} - (c-a) \geq 0$, D's optimal price under the domestic planner solution is given by $c-a$, in which case the cut-off type is a. Since when a tariff is imposed the firm's optimal threshold is $\frac{b-(\delta v-\tau-c)}{2}$, government can implement the same outcome by imposing a tariff such that $\frac{b-(\delta v-\tau-c)}{2} = a$, equivalently, $\tau = 2a-b+\delta v-c$.
- Otherwise, D's optimal price under the domestic planner solution is given by $c-b$, in which case the cut-off type is b. Since when a tariff is imposed the firm's optimal threshold is

$\frac{b-(\delta v-\tau-c)}{2}$, government can implement the same outcome by imposing a tariff such that $\frac{b-(\delta v-\tau-c)}{2} = b$, equivalently, $\tau = b+\delta v-c$.

Consider now the boundary case $\lambda = 2$, which part 2 of the statement includes. The coefficient $(\lambda-2)/2$ on $(e^{**})^2$ vanishes, so the planner's objective is affine in e^{**} and attains its maximum at an endpoint of $[a,b]$: the optimal e^{**} is again a or b. The objective's value at $e^{**} = b$ is zero, while its value at $e^{**} = a$ equals

$$(b-a)\left[v-\lambda\,\frac{a+b}{2}-(c-a)\right],$$

so the maximizing endpoint is a exactly when $v-\lambda\frac{a+b}{2}-(c-a)\geq 0$ — the same case condition as for $\lambda > 2$ — and the corner analysis and implementing tariffs above apply verbatim.

Finally, we verify that in each corner case, *every* tariff beyond the stated boundary implements the same corner allocation, which is why the corner cases are stated as inequalities on τ. Conditional on $[a,b)$ and a tariff τ, the domestic firm's objective $(\delta v-c-\tau+x)(b-x)$ is strictly concave in the marginal type x, so its optimum on $[a,b]$ is the unconstrained maximizer $\frac{b-(\delta v-\tau-c)}{2}$ projected onto $[a,b]$. Any $\tau \leq 2a-b+\delta v-c$ places the unconstrained maximizer weakly below a, so the firm prices at $c-a$ and full trade obtains; any $\tau \geq b+\delta v-c$ places it weakly above b, so no type accepts and no trade obtains. The lowest-absolute-value convention adopted in the text then selects, within each interval, the tariff closest to zero: the boundary value when the interval excludes zero, and $\tau = 0$ otherwise. □

Proof of Proposition 7. Following the proof of Proposition 3, when G is a uniform distribution and $e \in [a,b)$, the planner's problem is

$$\max_x (b-x)\left[v-c+x-\lambda\frac{x+b}{2}\right].$$

When $\lambda < 2$, the planner's objective function is concave. From the first order condition, the optimal $x = \frac{b-(v-c)}{2-\lambda}$. We note that if $b > \hat{e}$, then $\frac{b-(v-c)}{2-\lambda} > b$, since $v-c = (\lambda-1)\hat{e}$ and $\lambda > 1$. In this case, the planner's optimal response is $x = b$, yielding a zero payoff. This observation implies that if the right endpoint of an interval is greater than $\hat{e}$, then there is no trade. Thus, the foreign government optimally pools $[\hat{e},\overline{e}]$ under a single signal, which induces no trade. It is left to solve the optimal disclosure plan for $e \in [\underline{e},\hat{e})$. We assume that the disclosure plan is given by

$$[\underline{e},e_1),[e_1,e_2),\cdots$$

We write $e_0 = \underline{e}$ for convenience. We first establish that trade occurs with probability 1 when $e \in [e_i,e_{i+1})$. By Proposition 4, trade occurs with probability one or zero conditional on each category in any equilibrium; and by the converse splitting argument in the proof of Proposition 2

— which applies here since $\lambda < 2 = \underline{\lambda}$ for the uniform distribution — trade cannot break down on a positive measure of types below $\hat{e}$. Full trade must therefore be planner-optimal conditional on each trading category, which, given the concave objective, requires the interior optimum to lie weakly below the category's lower endpoint. This implies that $e_i \geq \frac{e_{i+1}-(v-c)}{2-\lambda}$. F's expected payoff is given by

$$\begin{aligned}&\sum_{i=0}^{\infty}\int_{e_i}^{e_{i+1}}(e-e_i)\,de\\&=\sum_{i=0}^{\infty}\frac{(e_{i+1}-e_i)^2}{2}.\end{aligned}$$

Define the following two functions, the intervention-case counterparts of the bounds $lb(\cdot)$ and $rb(\cdot)$ of Section 3.1, with the planner's conditional surplus replacing the firms' bilateral surplus (we denote them lb^g and rb^g to distinguish them from the laissez-faire bounds): $rb^g(\cdot)$ is the right bound of an interval such that trade occurs with probability 1, and $lb^g(\cdot)$ is the left bound of such an interval,

$$\begin{aligned}rb^g(e)&=(2-\lambda)e+(v-c)\\lb^g(e)&=\frac{e-(v-c)}{2-\lambda}.\end{aligned}$$

Suppose $lb^g(lb^g(e_{i+1})) \leq e_{i-1} < lb^g(e_{i+1})$, so that an interior cutoff is required to sustain trade and an admissible one exists. We ask which value of e_i maximizes F's expected payoff. We note that F's expected payoff is a quadratic function of e_i with a positive coefficient on e_i^2. By this convexity, F's expected payoff is maximized at one of the boundary values, i.e., $rb^g(e_{i-1})$ or $lb^g(e_{i+1})$. Moreover, the quadratic function is minimized at $(e_{i-1}+e_{i+1})/2$. We next show that $rb^g(e_{i-1})$ is always further from $(e_{i-1}+e_{i+1})/2$ than $lb^g(e_{i+1})$:

$$\begin{aligned}&rb^g(e_{i-1})+lb^g(e_{i+1}) > e_{i-1}+e_{i+1}\\\Leftrightarrow&(2-\lambda)e_{i-1}+v-c+\frac{e_{i+1}-(v-c)}{2-\lambda} > e_{i-1}+e_{i+1}\\\Leftrightarrow&-e_{i-1}-\frac{v-c}{2-\lambda}+\frac{e_{i+1}}{2-\lambda} > 0,\end{aligned}$$

which holds (after dividing both sides of the second line by $\lambda - 1 > 0$) given

$$e_{i-1} < lb^g(e_{i+1}) = \frac{e_{i+1}-(v-c)}{2-\lambda}.$$

As a result, to maximize F's expected payoff, $e_i = rb^g(e_{i-1})$. Repeatedly applying this result, we

obtain the optimal disclosure in the proposition.

When $\lambda > 2$, the planner's objective function

$$(b-x)\left[v-c+x-\lambda\,\frac{x+b}{2}\right]$$

is convex in x (linear at $\lambda = 2$, in which case the endpoint characterization below applies directly). Thus, the maximum is achieved at the two extremes when $x = a$ or $x = b$. Since $x = b$ yields a zero profit, the planner's response is full trade at $x = a$ if and only if

$$v-c+a-\lambda\frac{a+b}{2} \geq 0. \tag{A12}$$

The weak inequality reflects the maintained tie-breaking convention that an indifferent planner selects the response with the greater trade probability. Equivalently,

$$b \leq \frac{2(v-c)+(2-\lambda)a}{\lambda}. \tag{A13}$$

We next show that under the optimal disclosure plan, there is at most one signal realization that leads to trade between D and F. We argue by contradiction: suppose the plan contains two trading categories, $[\underline{e},a)$ and $[a,b)$, so that the planner is willing to quote the lower bound on each. We show that the planner is then also willing to quote the lower bound on the merged category $[\underline{e},b)$, and that merging strictly raises F's payoff.

We observe that

$$b \leq \frac{2(v-c)+(2-\lambda)a}{\lambda} \tag{A14}$$

$$\leq \frac{2(v-c)+(2-\lambda)\underline{e}}{\lambda} \tag{A15}$$

where the first inequality comes from (A13) and the second inequality comes from $a > \underline{e}$ together with $\lambda \geq 2$, so that $(2-\lambda)a \leq (2-\lambda)\underline{e}$. Applying (A13) again, it is evident that D is also willing to quote the lower bound when $e \in [\underline{e},b)$. Merging the two categories into $[\underline{e},b)$ then strictly raises F's payoff: types in $[a,b)$ trade at the price $c-\underline{e}$ rather than $c-a$ and earn strictly higher rents, while all other types are unaffected. Hence a plan with two trading categories cannot be optimal.

Overall, there is at most one signal realization that leads to trade. Applying (A13) at $a = \underline{e}$, and taking the largest right endpoint consistent with full trade,

$$e_1 = \frac{2(v-c)+(2-\lambda)\underline{e}}{\lambda},$$

which coincides with the cutoff in (29). If this expression weakly exceeds $\overline{e}$, then (A13) holds at $a = \underline{e}$, $b = \overline{e}$: the full pool $[\underline{e}, \overline{e}]$ sustains trade with probability one, the uninformative plan is optimal, and e_1 is capped at $\overline{e}$, as stated. □

Equilibrium tariff schedule for Corollary 1. Combining the results of Propositions 6 and 7 we obtain the following equilibrium results. We first consider the case when $\lambda < 2$. When the disclosure is $[e_i, e_{i+1})$, the government prefers to induce a price of $c - e_i$ so that trade always occurs. If a tariff τ_i is levied, D's optimal price is given by

$$c - \frac{e_{i+1} - (\delta v - \tau - c)}{2}.$$

Setting it equal to $c - e_i$,

$$\tau_i = (\delta v - c) - (e_{i+1} - 2e_i).$$

Since $e_{i+1} = (2-\lambda)e_i + (v-c)$, we solve

$$e_i = \hat{e} - (2-\lambda)^i(\hat{e} - \underline{e}).$$

With algebraic manipulation, we obtain

$$\begin{aligned}\tau_i &= (\delta v - c) - (e_{i+1} - 2e_i) \\ &= \delta v - c + \hat{e} - \lambda(2-\lambda)^i(\hat{e} - \underline{e}) \\ &= \lambda e_i + (\delta v - c) - (\lambda - 1)\hat{e}.\end{aligned}$$

Upon the signal $[\hat{e}, \overline{e}]$, the government prefers to induce a price of $c - \overline{e}$ so that trade never occurs. The minimal such tariff is given by $(\delta v - c) - (\overline{e} - 2\overline{e}) = \delta v - c + \overline{e}$.

We next consider $\lambda \geq 2$. When the disclosure is $[e_1, \overline{e}]$, the government prefers to induce no trade between D and F. In this case, the government can impose a large enough tariff so that D's optimal pricing induces no trade with F. Let τ^H denote the tariff. Then D's problem is

$$\max_x (\delta v - c - \tau^H + x)\left(1 - \frac{x - e_1}{\overline{e} - e_1}\right)$$

FOC implies

$$(\overline{e} - x) - (\delta v - c - \tau^H + x) = 0$$

Thus,

$$x^* = \frac{\overline{e} - (\delta v - c - \tau^H)}{2}.$$

No trade requires that

$$\frac{\overline{e}-(\delta v-c-\tau^H)}{2}\geq \overline{e}$$

Equivalently,

$$\tau^H \geq \overline{e}+\delta v-c$$

Now consider $e\in[\underline{e},e_1)$. Let τ^L denote the tariff. Then D's problem is

$$\max_x(\delta v-c-\tau^L+x)(e_1-x)$$

Then

$$x^*=\frac{e_1-(\delta v-c-\tau^L)}{2}.$$

Full trade requires that

$$\frac{e_1-(\delta v-c-\tau^L)}{2}\leq \underline{e}.$$

Equivalently,

$$\tau^L \leq 2\underline{e}-e_1+\delta v-c$$

Plugging in e_1,

$$\tau^L \leq 3\underline{e}-\frac{2(v-c)+2\underline{e}}{\lambda}+\delta v-c$$

Overall, if the tariff scheme (τ^L,τ^H) satisfies the above two conditions, then D would optimally quote a price that induces trade when $e\in[\underline{e},e_1)$ and that induces no trade when $e\in[e_1,\overline{e}]$. □

Proof of Corollary 1. *Part (i).* For $\lambda<2$, the equilibrium cutoffs satisfy the binding recursion $e_{i+1}=(2-\lambda)e_i+(v-c)$, at which the planner's corner solution at e_i coincides with the interior optimum. By Proposition 6, the set of implementing tariffs on such a category is the interval $\tau\leq 2e_i-e_{i+1}+\delta v-c=\lambda e_i-(1-\delta)v$, so the maximal implementing tariff is $\tau_i=\lambda e_i-(1-\delta)v$; it is the unique tariff at which the domestic firm's first-order condition holds exactly at e_i (Proposition 5 evaluated at $e_{\tilde{s}}^{**}=e_i$), and it is the schedule plotted in Figure 2. Lower tariffs, including the zero tariff selected by the lowest-absolute-value reporting convention when the interval contains zero, implement the same allocation; the schedule τ_i is the economically meaningful object because it measures the highest charge consistent with the category trading — the shadow price that disciplines the foreign government's disclosure choice.

Part (ii). Uniform G implies $m_i=(e_i+e_{i+1})/2$; substituting the recursion yields

$$m_i=\frac{(3-\lambda)\,e_i+(v-c)}{2},\qquad \text{i.e.,}\qquad e_i=\frac{2m_i-(v-c)}{3-\lambda},$$

and substituting e_i into τ_i delivers the affine form (30). The slope comparison $2\lambda/(3-\lambda)>\lambda$ is

equivalent to $\lambda > 1$, and the level comparison follows from $e_i < m_i$.

Part (iii). No trade above $\hat{e}$ requires $\tau \geq \delta v - c + \overline{e}$, as derived in the proof of Proposition 7's implementation. Since Assumption 4 (evaluated at the *lower* support endpoint $\underline{e}$, with $\underline{e} \geq 0$) implies $v - c > (\lambda - 1)\underline{e} \geq 0$, the intercept in (30) is strictly negative, so no proportional schedule $\tau = qm$ can match the tariffs on two or more trading categories, let alone the prohibitive region. □

Proof of Corollary 2. Expected global surplus from trading with all types in $[0,t]$ is

$$W(t) \;=\; \frac{1}{\overline{e}} \int_0^t [v - c + (1-\lambda)e]\, de \;=\; \frac{1}{\overline{e}} \left[(v-c)t - (\lambda - 1)\frac{t^2}{2} \right],$$

which is independent of δ and of the price and tariff, as these only redistribute surplus. Under the full-information benchmark, trade occurs on $[0, \hat{e}]$ with $\hat{e} = (v-c)/(\lambda - 1)$, so

$$W^{FB} = W(\hat{e}) = \frac{(v-c)^2}{2(\lambda - 1)\overline{e}}.$$

In equilibrium under $\lambda > 2$, by Propositions 6 and 7 with $\underline{e} = 0$, trade occurs exactly on the pooled green rating $[0, e_1]$ with $e_1 = 2(v-c)/\lambda$, which lies in $(\hat{e}, \overline{e}]$ given $2(v-c) \leq \lambda \overline{e}$. Hence

$$W^{eq} = W(e_1) = \frac{1}{\overline{e}} \left[\frac{2(v-c)^2}{\lambda} - (\lambda - 1)\frac{2(v-c)^2}{\lambda^2} \right] = \frac{2(v-c)^2}{\lambda^2 \overline{e}},$$

so that

$$\frac{W^{eq}}{W^{FB}} = \frac{4(\lambda - 1)}{\lambda^2} \qquad \Longrightarrow \qquad \frac{W^{FB} - W^{eq}}{W^{FB}} = \frac{\lambda^2 - 4\lambda + 4}{\lambda^2} = \left(\frac{\lambda - 2}{\lambda} \right)^2. \qquad \square$$

Lemma 6. *Suppose φ is nonincreasing. Then for any b, the function $d(\cdot\,; b)$ is concave and then convex on $[\underline{e}, b]$ (either piece possibly empty), and full trade is planner-optimal on a category $[a, b)$ if and only if $d'(a; b) \leq 0$ and $d(a; b) \geq 0$.*

Proof. By equation (32),

$$\operatorname{sign} d''(x; b) = -\operatorname{sign} \left[(2 - \lambda) + (\lambda - 1)\varphi(x) \right].$$

Since $\lambda > 1$ and φ is nonincreasing, the bracketed term is nonincreasing in x, so d'' changes sign at most once, from negative to positive.

For the second claim, suppose $d'(a; b) \leq 0$ and $d(a; b) \geq 0$. On the concave piece, d' is nonincreasing, so $d' \leq 0$ throughout and d is nonincreasing; on the convex piece, d' is nondecreasing,

so d attains its maximum over that piece at one of its endpoints. Since the left endpoint of the convex piece has d below $d(a;b)$ and the right endpoint satisfies $d(b;b) = 0 \leq d(a;b)$, the global maximum over $[a,b]$ equals $d(a;b)$, so full trade at a is optimal. Conversely, if $d'(a;b) > 0$, moving the marginal type slightly above a strictly raises the planner's payoff, and if $d(a;b) < 0$, no trade strictly dominates; in either case full trade is not optimal. □

Proof of Proposition 8. Throughout the proof, "planner-optimal" refers to the *domestic* planner — the government maximizing domestic surplus conditional on each disclosure category via $d(\cdot;b)$ — not to a global planner, whose objective would also include the foreign firm's surplus.

Part (i). By Lemma 6, full trade on $[a,b)$ is planner-optimal iff $d'(a;b) \leq 0$ and $d(a;b) \geq 0$. Using $v - c + (1-\lambda)a = (\lambda - 1)(\hat{e} - a)$, the first condition reads

$$G(b) - G(a) \;\leq\; (\lambda - 1)\, g(a)(\hat{e} - a),$$

i.e., $b \leq b_1(a)$ as defined in (33) (with $b_1(a) = \overline{e}$ if the right-hand side of (33) exceeds $G(\overline{e})$). For the second,

$$\frac{\partial d(a;b)}{\partial b} = (v - c + a - \lambda b)\, g(b),$$

which is positive for $b < (v - c + a)/\lambda$ and negative thereafter, so $d(a;\cdot)$ is single-peaked with $d(a;a) = 0$; hence

$$\{b : d(a;b) \geq 0\} = [a, b_2(a)],$$

with $b_2(a) = \overline{e}$ if $d(a;\overline{e}) \geq 0$. Combining, full trade is planner-optimal iff

$$b \;\leq\; B(a) \;=\; \min\{b_1(a), b_2(a)\}.$$

Note that $B(a) > a$ for every $a < \hat{e}$, since $(\lambda - 1)g(a)(\hat{e} - a) > 0$ and

$$\frac{\partial d(a;b)}{\partial b}\Big|_{b=a} = (\lambda - 1)(\hat{e} - a)\, g(a) > 0.$$

Part (ii). By Proposition 4, every on-path category features trade with probability one or zero, and by Step 1 of its proof, no category contained in $[\hat{e}, \overline{e}]$ trades. First, the set of trading types has no gaps: if some non-trading category $[x,y)$ with $x < \hat{e}$ existed (below or between trading categories), the foreign government could split off $[x,b')$ with $b' < \min\{B(x), y\}$ chosen so that $d'(x;b') < 0$ and $d(x;b') > 0$ hold strictly; full trade on $[x,b')$ is then uniquely planner-optimal, all other categories are unaffected, and the foreign firm collects strictly positive rents from types in (x,b') — a profitable deviation. Hence the trading types form an interval $[\underline{e}, T]$. Second, $T \geq \hat{e}$: if $T < \hat{e}$, appending the category $[T,b')$ with $b' \leq B(T)$, which trades by part (i), is again a profitable

deviation.

Third, we determine when the partition of $[\underline{e}, T]$ can contain infinitely many categories. Since no category contained in $[\hat{e}, \bar{e}]$ trades, every trading category begins strictly below $\hat{e}$, so an infinite collection of cutoffs must accumulate at some $m \leq \hat{e}$. Accumulation at $m < \hat{e}$ is impossible: for a near m, $G(b_1(a)) - G(a) = (\lambda - 1)\, g(a)(\hat{e} - a)$ is bounded below by a strictly positive constant, and $\partial d(a;b)/\partial b|_{b=a} = (\lambda - 1)(\hat{e} - a)g(a)$ is likewise bounded below, so $B(a) - a$ is bounded away from zero near m — whereas accumulation requires category widths to vanish. An infinite partition therefore has cutoffs accumulating at $\hat{e}$ itself, with every category remaining below $\hat{e}$; in that case $T = \hat{e}$ and trade is globally efficient. When $\lambda > 2$, this case cannot arise. By continuity of g and $\lambda - 1 > 1$, for a sufficiently close to $\hat{e}$ we have $(\lambda - 1)g(a) > \sup_{[a,\hat{e}]} g$, so

$$G(b_1(a)) - G(a) \;=\; (\lambda - 1)\, g(a)(\hat{e} - a) \;>\; (\hat{e} - a) \sup_{[a,\hat{e}]} g \;\geq\; G(\hat{e}) - G(a),$$

giving $b_1(a) > \hat{e}$. Moreover, substituting $e = a + t(\hat{e} - a)$,

$$d(a;\hat{e}) \;=\; (\hat{e} - a)^2 \int_0^1 \big[(\lambda - 1) - \lambda t\big]\, g\big(a + t(\hat{e} - a)\big)\, dt,$$

and the integral converges to $g(\hat{e})(\lambda - 2)/2 > 0$ as $a \uparrow \hat{e}$, so $d(a;\hat{e}) > 0$ and hence $b_2(a) > \hat{e}$ near $\hat{e}$ as well. Thus $B(a) > \hat{e}$ for a close to $\hat{e}$. If cutoffs nevertheless accumulated at $\hat{e}$, the foreign government could pick a cutoff a_K close enough to $\hat{e}$ that $B(a_K) > \hat{e}$ and merge the infinite tail of categories contained in $[a_K, \hat{e})$ into the single category $[a_K, b')$ for some $b' \in (\hat{e}, B(a_K)]$: full trade on the merged category is planner-optimal by part (i), every type in $[a_K, \hat{e})$ receives a weakly higher price (its category's lower endpoint falls to a_K), and types in $[\hat{e}, b')$ newly trade with strictly positive rents — a strictly profitable deviation. For $\lambda > 2$ the equilibrium partition is therefore finite.

Fourth, when a terminal trading category exists, its bound binds: if the terminal trading category $[a_K, T)$ had $T < B(a_K)$, extending it to $T + \eta \leq B(a_K)$ would preserve planner-optimality of full trade and strictly raise the foreign firm's payoff (types in $[T, T + \eta)$ earn $e - a_K > 0$). Finally, trade on $(\hat{e}, T]$ generates strictly negative global surplus by the definition of $\hat{e}$, so equilibrium trade is globally excessive if and only if $T > \hat{e}$.

Part (iii). Suppose that, in equilibrium, the screening bound is the binding constraint at every non-terminal step; that is, every non-terminal trading category satisfies $a_{i+1} = B(a_i)$ with $B(a_i) = b_1(a_i) \leq b_2(a_i)$. By part (ii), the trading region begins at $\underline{e}$, so $a_1 = \underline{e}$, and the equilibrium cutoffs are generated by the recursion $a_{i+1} = B(a_i)$, as stated. □

Proof of Lemma 2. Write $\beta = 2(1 - \varepsilon) = g'$.

Monotonicity of φ: $\varphi(x) = \beta(\hat{e}-x)/g(x)$, so

$$\varphi'(x) = -\frac{\beta\,[g(x)+\beta(\hat{e}-x)]}{g(x)^2} = -\frac{\beta\,g(\hat{e})}{g(x)^2} \;<\; 0$$

for $\varepsilon < 1$, since $g(\hat{e}) > 0$.

Hazard rate: g is increasing and $1-G$ decreasing, so $g/(1-G)$ is increasing.

Threshold $\underline{\lambda}$: for any b and $x < b$,

$$1 - \frac{d}{dx}\left[\frac{G(b)-G(x)}{g(x)}\right] = 2 + \frac{(G(b)-G(x))\,\beta}{g(x)^2} \;\geq\; 2,$$

with the bound approached as $x \uparrow b$; hence the infimum in (22) equals 2.

Threshold $\hat{\lambda}$:

$$d''(x) = (\lambda-2)\,g(x) - \beta(\lambda-1)(\hat{e}-x)$$

has derivative $\beta(2\lambda-3) > 0$ in x for $\lambda > 3/2$, so global convexity on all intervals holds iff

$$d''(0) = (\lambda-2)\varepsilon - 2(1-\varepsilon)(v-c) \geq 0,$$

i.e., iff $\lambda \geq \hat{\lambda}(\varepsilon)$ in (34); for $\lambda < 2$, $d''(0) < 0$, so no smaller threshold applies.

Efficiency: for $\lambda < 2 = \underline{\lambda}$, Proposition 2 yields global efficiency. For $\lambda > 2$, we show $T > \hat{e}$. Substituting $e = a + t(\hat{e}-a)$,

$$d(a;\hat{e}) = (\hat{e}-a)^2 \int_0^1 [(\lambda-1) - \lambda t]\, g(a+t(\hat{e}-a))\, dt,$$

and the integral converges to $g(\hat{e})(\lambda-2)/2 > 0$ as $a \uparrow \hat{e}$, so $d(a;\hat{e}) > 0$ and hence $b_2(a) > \hat{e}$ for a sufficiently close to $\hat{e}$. Moreover, once $(\lambda-1)g(a) > g(\hat{e})$ — which holds for a close to $\hat{e}$ since $\lambda - 1 > 1$ —

$$G(b_1(a)) - G(a) \;=\; (\lambda-1)g(a)(\hat{e}-a) \;>\; g(\hat{e})(\hat{e}-a) \;\geq\; G(\hat{e}) - G(a),$$

so $b_1(a) > \hat{e}$ near $\hat{e}$ as well. Hence the equilibrium plan cannot accumulate at or below $\hat{e}$: by Proposition 8(ii), $T \geq \hat{e}$, and any category starting sufficiently close to $\hat{e}$ has $B(a) > \hat{e}$, so $T > \hat{e}$ and trade is globally excessive. □

Closed-form cutoffs for Proposition 9. For $\lambda > 2$, the cutoffs of the domestic firm's optimal

disclosure plan are

$$a^* \;=\; \max\left\{\frac{(v-c)(4-\lambda)}{3\lambda-4},\; \frac{2(v-c)-\lambda\overline{e}}{\lambda-2},\; 0\right\}, \qquad b^* \;=\; \frac{2}{\lambda}\big(v-c+a^*\big)-a^*, \tag{A16}$$

and $b^* > \hat{e}$. When the first term in (A16) attains the maximum, the plan is an interior hybrid with $b^* < \overline{e}$; when the second term attains it, $b^* = \overline{e}$ and trade occurs for all types; and when $a^* = 0$ — which obtains for $\lambda \geq 4$ with $2(v-c) \leq \lambda\overline{e}$ — the plan is the two-rating disclosure plan of Proposition 7, with cutoff given by (29).

Proof of Proposition 9. To solve for D's optimal disclosure plan, we first need the government's optimal tariff conditional on a disclosure that informs $e \in [a,b)$. This problem is solved in Proposition 6. Throughout this proof, $\delta = 1$ by hypothesis, so $\delta v - c = v - c$ and the $(1-\delta)v$ terms in the tariff formulas vanish. Applying the proposition and the lowest-absolute-value convention, the optimal tariff is given by the following.

1. Suppose $\lambda < 2$
   (a) If $a \geq \frac{b-(v-c)}{2-\lambda}$, the optimal tariff is $\tau = 0$.
   (b) If $a < \frac{b-(v-c)}{2-\lambda} < b$, the optimal tariff is $\tau = \frac{\lambda(b-v+c)}{2-\lambda}$.
   (c) If $\frac{b-(v-c)}{2-\lambda} \geq b$, the optimal tariff is $\tau = b+v-c$.
2. Suppose $\lambda \geq 2$
   (a) If $v - \lambda\frac{a+b}{2} - (c-a) \geq 0$, the optimal tariff is $\tau = 0$.
   (b) Otherwise, the optimal tariff is $\tau = b+v-c$.

First, we consider the case when $\lambda < 2$. We first claim that under D's optimal disclosure plan, we should not expect case (b). We argue by contradiction. Suppose $a < \frac{b-(v-c)}{2-\lambda} < b$. Then D can choose a new disclosure plan, which inherits the existing one but divides the interval $[a,b)$ at the point $\frac{b-(v-c)}{2-\lambda}$ into two subintervals. By doing so, D does not lose on the upper subinterval but strictly gains on the lower one, a contradiction. Next, we observe that trade occurs with zero probability under case (c), so D earns a zero profit whenever $\frac{b-(v-c)}{2-\lambda} \geq b$. Thus far, we have shown that D earns a nonzero profit only in case (a). In case (a), $a \geq \frac{b-(v-c)}{2-\lambda}$. Recall that the threshold at which global surplus is zero is given by $\hat{e} = \frac{v-c}{\lambda-1}$. We note that whenever $b > \hat{e}$, we also have $\frac{b-(v-c)}{2-\lambda} \geq b$, which is in case (c). So D does not gain by pooling types above $\hat{e}$. For types below $\hat{e}$, by choosing case (a), government always charges a zero tariff. Thus, for $e \leq \hat{e}$, the finest disclosure will give D the maximal possible surplus by charging a price such that F breaks even. Note that in

case (c), the finest disclosure also gives D a zero profit. Overall, the finest disclosure constitutes an optimal disclosure for D.

Second, consider $\lambda > 2$ (recall $\delta = 1$, so $v - c = v - c$). As in the first case, the domestic firm earns rents only on categories in which trade occurs. On any trading category $[a,b)$, the government's tariff is zero (case (a) above) and the domestic firm's optimal price is $c - a$, so the category contributes $(b-a)(v-c+a)$ to the domestic firm's expected payoff (up to the factor $1/\overline{e}$), whereas perfectly revealing the same interval contributes strictly more:

$$(b-a)(v-c+a) \;<\; \int_a^b (v-c+e)\,de.$$

The optimal plan therefore reveals types perfectly wherever trade is to occur below some point a, and uses at most one pooled category $[a,b)$ to extend trade beyond what revelation can sustain. Trade on the pooled category requires the government's full-trade condition

$$v-c+a-\lambda\,\frac{a+b}{2} \;\geq\; 0, \qquad \text{that is,} \qquad b \;\leq\; b(a) \equiv \frac{2}{\lambda}(v-c+a)-a,$$

and the domestic firm optimally stretches the pool to this bound, since its contribution increases in b. No second pooled trading category above $b(a)$ is viable: $b(x) > x$ if and only if $x < \hat{e}$, and any pooled category satisfying the full-trade condition with $a < \hat{e}$ already extends beyond $\hat{e}$.

The domestic firm thus chooses a and the pool's upper end $b \leq \min\{b(a), \overline{e}\}$ to maximize (up to the factor $1/\overline{e}$)

$$f(a,b) \;=\; (v-c)\,a + \frac{a^2}{2} + (b-a)\,(v-c+a),$$

and since the pool's contribution increases in b, it sets $b = \min\{b(a), \overline{e}\}$. The bound $b(\cdot)$ is strictly decreasing for $\lambda > 2$; let a_{cap} denote the value at which the bound reaches the top of the support, $b(a_{\text{cap}}) = \overline{e}$:

$$a_{\text{cap}} \;\equiv\; \frac{2(v-c)-\lambda\overline{e}}{\lambda-2}.$$

For $a < a_{\text{cap}}$ (a nonempty range only if $a_{\text{cap}} > 0$), the pool is capped at $\overline{e}$ and

$$\frac{\partial}{\partial a} f(a,\overline{e}) \;=\; \overline{e}-a \;>\; 0,$$

so the objective is strictly increasing in a. For $a > a_{\text{cap}}$, the pool is uncapped; substituting $b(a)$ and differentiating,

$$f'(a) \;=\; \frac{4}{\lambda}\,(v-c+a) - (v-c) - 3a, \qquad f''(a) \;=\; \frac{4}{\lambda} - 3 \;<\; 0,$$

so the objective is strictly concave, with unconstrained maximizer equal to the first term in (A16). The objective is therefore single-peaked in a, with maximizer

$$a^* = \max\left\{\frac{(v-c)(4-\lambda)}{3\lambda-4},\, a_{\text{cap}},\, 0\right\}$$

as in (A16) and $b^* = b(a^*)$, which equals $\overline{e}$ when $a^* = a_{\text{cap}}$. Assumption 4 guarantees $a_{\text{cap}} \le \hat{e}$: the inequality

$$\frac{2(v-c)-\lambda\overline{e}}{\lambda-2} \;\le\; \frac{v-c}{\lambda-1}$$

reduces to $v-c \le (\lambda-1)\overline{e}$. Moreover, the interior term is below $\hat{e}$ if and only if $\lambda > 2$; hence $a^* \le \hat{e}$ and all revealed types below a^* trade.

It remains to verify the properties in the statement. If the first term in (A16) attains the maximum, solving $f'(a^*) = 0$ gives

$$a^* = \frac{(v-c)(4-\lambda)}{3\lambda-4}, \qquad b^* = \frac{(v-c)\lambda}{3\lambda-4},$$

with

$$b^* - \hat{e} \;=\; \frac{(v-c)\,(\lambda-2)^2}{(3\lambda-4)(\lambda-1)} \;>\; 0.$$

If $a^* = a_{\text{cap}}$, then $b^* = \overline{e} > \hat{e}$ and trade occurs for all types. If $a^* = 0$ — which requires $\lambda \ge 4$ (so that the first term is nonpositive) and $2(v-c) \le \lambda\overline{e}$ (so that $a_{\text{cap}} \le 0$) — the revealed range vanishes and the plan collapses to the two-rating disclosure plan of Proposition 7, whose cutoff $b(0) = 2(v-c)/\lambda$ equals (29) at $\underline{e} = 0$ and exceeds $\hat{e}$ if and only if $\lambda > 2$. Finally, at $a = \hat{e}$ the pool is empty ($b(\hat{e}) = \hat{e}$) and $f(\hat{e},\hat{e})$ equals the full-disclosure payoff; for $\lambda > 2$ the single peak lies strictly below $\hat{e}$, so the plan with a pooled rating strictly dominates full disclosure, whereas at $\lambda = 2$ the interior term equals $v-c = \hat{e}$, the pool is empty, and full disclosure remains weakly optimal. Combined with the first case, full disclosure is weakly optimal if and only if $\lambda \le 2$. □

Proof of Proposition 10. We first state the equilibrium concept formally.

Definition 2. *A $(\tau(\cdot), D(\cdot), \mu(\cdot), p(\cdot,\cdot))$ profile forms a perfect Bayesian equilibrium of the interim game if:*

1. *For every possible signal s, the tariff schedule $\tau(s)$ maximizes the domestic welfare.*
2. *For every possible signal s and tariff $\tau(s)$, $p(\tau,s)$ solves the domestic firm's problem* $\max_p(\delta v - p - \tau(s)) \cdot \Pr_{\mu(s)}[c-e \le p]$.

3. *For every $e \in [\underline{e}, \bar{e}]$, $D(e)$ solves* $\arg\max_s \{\max[p(\tau(s), s) - (c - e), 0]\}$*, where the maximization is over interval messages s containing e, and the strategy $D(\cdot)$ satisfies the monotonicity restriction of Assumption 1.*

4. *For every s in the range of D (i.e., every interval s that can be disclosed in equilibrium), the domestic sector's belief $\mu(s)$ is obtained by applying Bayes' rule given the particular signal s.*

We now show how to construct a strategy profile that supports the optimal disclosure plan of the ex ante game in the interim game. Suppose the optimal disclosure plan of the ex ante game is denoted by $D : [\underline{e}, \bar{e}] \to S$. In the interim game, consider a candidate equilibrium in which:

1. Disclosure follows the ex ante disclosure $D(e)$.

2. For every possible signal s, the tariff schedule $\tau(s)$ maximizes the domestic welfare.

3. For $s \in S$, the belief $\mu(s)$ is given by G conditional on s. For any $s \notin S$ (that is, whenever s is an off-equilibrium signal), the belief $\mu(s)$ assigns probability 1 to type $\bar{e}(s)$.

4. For every possible signal s, $p(s)$ maximizes $\max_p (\delta v - p - \tau(s)) \Pr_{\mu(s)}[c - e \leq p]$.

The constructed profile is a perfect Bayesian equilibrium: on the equilibrium path, beliefs follow Bayes' rule and each party best responds by construction; off the equilibrium path, a type e deviating to a message $s \notin S$ faces the belief concentrated on $\bar{e}(s) = \sup s$, receives a price offer of at most $c - \bar{e}(s)$, and hence a payoff of at most $\max\{e - \bar{e}(s), 0\} = 0$ (since $e \in s$ implies $e \leq \bar{e}(s)$), which does not exceed its equilibrium payoff. We can show that the constructed strategy profile forms a foreign-firm-preferred equilibrium. By contradiction, suppose $(\tau_1(\cdot), D_1(\cdot), \mu_1(\cdot), p_1(\cdot))$ is an equilibrium that dominates $(\tau(\cdot), D(\cdot), \mu(\cdot), p(\cdot))$ among F types in the Pareto sense based on their interim payoffs. Let $S_1 = \{D_1(e) : e \in [\underline{e}, \bar{e}]\}$ and consider an ex ante disclosure $D_1 : [\underline{e}, \bar{e}] \to S_1$. Since interim disclosure strategies are subject to the monotonicity restriction of Assumption 1, $D_1(\cdot)$ is a monotone partition plan and hence an admissible disclosure plan of the ex ante game. Since $\mu_1(\cdot)$ is obtained by applying Bayes' rule on the equilibrium path, F's expected payoff under the disclosure D_1 of the ex ante game is equal to its expected payoff under the $(\tau_1(\cdot), D_1(\cdot), \mu_1(\cdot), p_1(\cdot))$ equilibrium of the interim game. It then follows that F's expected payoff under the disclosure plan $D_1(\cdot)$ is strictly higher than the one under the disclosure plan $D(\cdot)$ in the ex ante game, contradicting the fact that $D(\cdot)$ is an optimal disclosure plan of the ex ante game.

Before we move to the second part of the proof, we need to formalize the definition of Grossman-Perry-Farrell equilibrium. Denote by $U(e, s, \mu(s))$ the foreign firm's utility if its type is e, it sends a

message s, and the domestic government and the domestic firm best respond to the belief $\mu(s)$: the government's tariff $\tau(s)$ maximizes domestic welfare and the domestic firm quotes an optimal price given $\mu(s)$ and $\tau(s)$. For any interval message $s \subseteq [\underline{e}, \bar{e}]$ (including off-equilibrium messages), denote by μ_s the actual distribution of e conditional on $e \in s$. As in Bertomeu and Cianciaruso (2018), we define a Grossman-Perry-Farrell equilibrium by ruling out the existence of self-signaling sets.

Definition 3. *A pure-strategy perfect Bayesian equilibrium of the interim disclosure game* $(\tau(\cdot), D(\cdot), \mu(\cdot), p(\cdot))$ *is called a "Grossman-Perry-Farrell equilibrium" if there does not exist a self-signaling set, which is defined as a non-empty interval* $s \subset [\underline{e}, \bar{e}]$ *such that:*

$$s = \{e \in s : U(e, s, \mu_s) > U(e, D(e), \mu(D(e)))\}. \tag{A17}$$

We now show that the constructed strategy profile above also forms a Grossman-Perry-Farrell equilibrium of the interim game — indeed, we show that no self-signaling set exists at all.

We use three structural properties of the candidate equilibrium plan established in our baseline analysis: trade occurs with probability one conditional on each trading category (Proposition 4), so that a type e's equilibrium payoff is

$$U(e, D(e), \mu(D(e))) \;=\; e - a(e) \;\geq\; 0,$$

where $a(e) \leq e$ denotes the lower endpoint of e's trading category, and zero for types that do not trade; the planner's response conditional on any category contained in $[\hat{e}, \bar{e}]$ is no trade (Step 1 in the proof of Proposition 4); and the trading types form an interval $[\underline{e}, T]$ with $T \geq \hat{e}$, so that every trading category begins strictly below $\hat{e}$ (by the splitting arguments in the proofs of Propositions 2–3 and Proposition 8(ii), a plan excluding a positive measure of types below $\hat{e}$ from trade cannot be optimal).

Suppose, toward a contradiction, that a self-signaling interval s_0 with infimum x and supremum y exists, so that every type in s_0 strictly prefers the deviation outcome to its equilibrium outcome. In the deviation outcome, the domestic government best responds to the truthful conditional belief μ_{s_0}: it implements the domestic planner solution conditional on $e \in s_0$, inducing a marginal type $m \in [x, y]$ such that the domestic firm's price is $c - m$ and trade occurs for $e \geq m$; a deviating type $e \in s_0$ therefore receives $e - m$ if $e \geq m$ and zero otherwise. First, the deviation must induce trade: otherwise every member of s_0 receives zero, no more than its equilibrium payoff. In particular $x < \hat{e}$, since a message s_0 with $x \geq \hat{e}$ asserts $e \in [\hat{e}, \bar{e}]$, conditional on which the planner induces no trade. Second, $m = x$: if $m > x$, the types in $(x, m) \subset s_0$ receive zero, no more than their equilibrium payoffs, contradicting the self-signaling property. Third, consider members of s_0 arbitrarily close

to its infimum x. Since $x < \hat{e} \leq T$, the point x lies in an equilibrium trading category $[a^*, b^*)$ with $a^* \leq x$, and every $e \in (x, \min\{b^*, y\})$ belongs to that same category, so its equilibrium payoff satisfies

$$U(e, D(e), \mu(D(e))) \;=\; e - a^* \;\geq\; e - x,$$

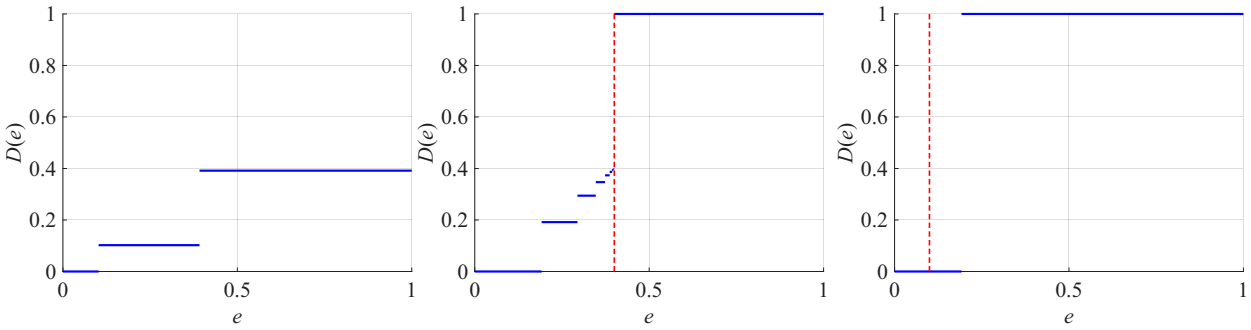
which weakly exceeds the deviation payoff $e - m = e - x$: such types do not strictly gain. Since s_0 is an interval with infimum x, it contains such types unless $s_0 = \{x\}$, in which case type x itself receives $x - x = 0 \leq U(x, D(x), \mu(D(x)))$. Either way, some member of s_0 fails to strictly gain, contradicting the requirement that every member of s_0 strictly gains. No self-signaling interval exists, so the constructed profile is a Grossman-Perry-Farrell equilibrium. □

B Appendix: Additional Results

B.1 Normally Distributed Environmental Externalities

We analyze the case where externalities follow a truncated normal distribution. We consider optimal disclosure plans chosen by the foreign government both absent interventions and in the presence of the optimal domestic tariff scheme. Importantly, we again examine how households' exposure to externalities shapes equilibrium outcomes.

Figure 6 considers a parameterization with a relatively low exposure of households to externalities. Absent interventions, the foreign government chooses a fairly coarse disclosure plan with three potential signals (see Panel (a)). In contrast, in the presence of government interventions (Panel (b)), the optimal disclosure plan becomes increasingly granular as externalities approach $\hat{e} \approx 0.4$, the level where global surplus turns negative. Interestingly, in this case the anticipation of government tariffs and the associated reduction in gains from trade between the domestic and foreign firm necessitate more informative disclosures to avoid inefficient trade breakdowns. This result obtains specifically when the domestic sector has a relatively low exposure to externalities (here $\lambda = 1.5$). In contrast, Panel (c) of the same figure illustrates the case where households are more strongly exposed to externalities, with the exposure parameter increased to $\lambda = 3$. In this case, government interventions reduce the precision of disclosures. The dashed red lines mark this threshold in both intervention panels: trade stops exactly at $\hat{e}$ under moderate exposure, while the single trading rating in Panel (c) extends well beyond $\hat{e}$, sustaining globally excessive trade.



(a) No intervention **(b)** Intervention & $\lambda = 1.5$ **(c)** Intervention & $\lambda = 3$

Figure 6: The figure illustrates the optimal disclosure plan chosen by the foreign government with and without government intervention. Externalities follow $e \sim N(0.5, 1)$, truncated from below at 0, and from above at 1, with $v - c = 0.2$ and $\delta = 1$. In Panels (b) and (c), the dashed red line marks the efficiency threshold $\hat{e}$: trade stops exactly at $\hat{e}$ under moderate exposure ($\lambda = 1.5$) and extends beyond it under severe exposure ($\lambda = 3$).

B.2 Unconditional Tariff Schemes

With an unconditional tariff, we can rewrite the domestic firm's payoff conditional on trade occurring as follows:

$$u_D = (\delta v - \tau) - p. \tag{A18}$$

Note that the domestic firm and the foreign firm in this case solve problems that are equivalent to the ones absent government interventions, except that one needs to replace the variable δv by $(\delta v - \tau)$. As a result, when the foreign government chooses an optimal disclosure plan, trade occurs if and only if the bilateral surplus of the two firms, after accounting for tariffs, is positive, that is, when:

$$u_D + u_F \geq 0 \quad \Leftrightarrow \quad e \geq c - \delta v + \tau. \tag{A19}$$

As a result, unconditional tariffs ($\tau > 0$) lead to the rationing of *low-externality* (that is, relatively green) transactions relative to the case absent government interventions. If those greener transactions are desirable from the perspective of the domestic country, the unconditional tariff backfires, since it only suppresses trade in those cases.

Comparing the unconditional tariff scheme to the disclosure-conditional tariff schedules studied in the main text, conditioning weakly increases domestic surplus for any *given* disclosure plan.

Proposition 11 (Conditioning weakly dominates plan-by-plan)**.** *Fix any disclosure plan $D(\cdot)$ with signal set S, and let $W(D, \tau)$ denote the expected domestic surplus under plan D when the tariff equals τ irrespective of the disclosure. Let $W_{cond}(D)$ denote the expected domestic surplus when*

the domestic government sets its optimal tariff $\tau^*(\tilde{s})$ conditional on each disclosure $\tilde{s} \in S$, and let $W_{uncond}(D) \equiv \max_{\tau \in \mathbb{R}} W(D,\tau)$ denote the expected domestic surplus under the best constant tariff. Then

$$W_{cond}(D) \geq W_{uncond}(D), \tag{A20}$$

with strict inequality whenever the optimal conditional schedule $\tau^*(\tilde{s})$ cannot be replaced by a constant tariff without changing some induced marginal type.

Proof. Holding the disclosure plan $D(\cdot)$ fixed, the domestic government's problem separates across realized disclosures: conditional on $\tilde{s}$, the tariff $\tau(\tilde{s})$ affects payoffs only through the marginal type it induces within $D^{-1}(\tilde{s})$. A constant schedule $\tau(\tilde{s}) \equiv \tau$ is a feasible element of the class of conditional schedules, so optimizing signal-by-signal yields weakly higher expected domestic surplus, and strictly higher surplus whenever no constant tariff induces the optimal marginal type in every category simultaneously. □

In equilibrium, the comparison is more subtle, because the foreign government re-optimizes its disclosure plan against each tariff regime. Two features of the uniform example of Section 3.4 illustrate the possibilities. Under severe exposure ($\lambda > 2$ with the pooling constraint binding), the foreign government extends the "green" trading pool until the domestic sector is exactly indifferent between trading with the pool and not trading at all, so equilibrium domestic surplus equals its no-trade level of zero *despite* the optimal conditional policy; a constant tariff can do no better (its best use is to shut trade down), so conditioning confers no equilibrium gain to the domestic sector in this regime — regulation protects, but does not profit. Under moderate exposure ($\lambda < 2$), by contrast, the conditional-tariff equilibrium delivers strictly positive domestic surplus, while an unconditional tariff cannot: raising a constant tariff rations the *greenest* suppliers first (as shown above) and leaves the pooled brown segment trading. As a numerical observation for the parameterization of Figure 3 ($e \sim U[0,1]$, $\delta = 1$, $v - c = 0.2$, $\lambda = 1.25$): under a constant tariff τ, the continuation game is the laissez-faire game with effective valuation $v - \tau$, so the foreign government's optimal plan is the plan of Proposition 1 for that effective valuation, and domestic surplus can be computed in closed form category by category. Evaluating over a fine grid of τ (including subsidies), expected domestic surplus is strictly negative for every τ below the prohibitive level $v - c + \bar{e} = 1.2$ — e.g., -0.155 at $\tau = 0$ — and equals zero at shutdown; the best constant tariff thus attains zero, while the conditional-tariff equilibrium attains ≈ 0.034. Conditioning on disclosures is thus strictly valuable to the domestic sector precisely in the moderate-exposure regime, where fine-tuned tariffs retain their potency.

References

Acharya, V. V., DeMarzo, P., Kremer, I., 2011. Endogenous Information Flows and the Clustering of Announcements. American Economic Review 101, 2955–2979.

Admati, A. R., Pfleiderer, P., 2000. Forcing Firms to Talk: Financial Disclosure Regulation and Externalities. Review of Financial Studies 13, 479–519.

Aghamolla, C., An, B.-J., 2023. Mandatory vs. voluntary esg disclosure, efficiency, and real effects. Working Paper, Nanyang Business School Research Paper No. 22-41.

Ali, S. N., Lewis, G., Vasserman, S., 2023. Voluntary Disclosure and Personalized Pricing. The Review of Economic Studies 90, 538–571.

Antràs, P., Helpman, E., 2004. Global sourcing. Journal of Political Economy 112, 552–580.

Antràs, P., Staiger, R., 2012. Offshoring and the role of trade agreements. American Economic Review 102, 3140–83.

Azarmsa, E., Shapiro, J. D., 2025. The market for ESG ratings. Journal of Finance, forthcoming .

Barwick, P. J., Li, S., Lin, L., Zou, E. Y., 2024. From fog to smog: The value of pollution information. American Economic Review 114, 1338–1381.

Bergemann, D., Brooks, B., Morris, S., 2015. The Limits of Price Discrimination. American Economic Review 105, 921–957.

Bertomeu, J., Cianciaruso, D., 2018. Verifiable disclosure. Economic Theory 65, 1011–1044.

Beyer, A., Cohen, D. A., Lys, T. Z., Walther, B. R., 2010. The financial reporting environment: Review of the recent literature. Journal of accounting and economics 50, 296–343.

Bolton, P., Kacperczyk, M., 2021. Do investors care about carbon risk? Journal of Financial Economics 142, 517–549.

Bonetti, P., En, Y., Kadach, I., Ormazabal, G., 2024. The role of supply chains in climate disclosure, working Paper, IESE Business School; SSRN No. 4850301.

Chen, H., Schneemeier, J., 2022. Responsible investors and stock market feedback. Working Paper.

Christensen, H. B., Hail, L., Leuz, C., 2021. Mandatory csr and sustainability reporting: Economic analysis and literature review. Review of Accounting Studies 26, 1176–1248.

Climate Action Tracker, 2023. Country assessments. http://climateactiontracker.org, accessed: 2023-11-11.

Cong, L. W., Qu, Y., Wang, G., 2024. Blockchains for environmental monitoring: Theory and empirical evidence from china. Available at SSRN 5001165 .

Copeland, B., Taylor, M., 2003. Trade and the Environment: Theory and Evidence. Princeton University Press.

Darendeli, A., Fiechter, P., Hitz, J.-M., Lehmann, N., 2022. The role of corporate social responsibility (csr) information in supply-chain contracting: Evidence from the expansion of csr rating coverage. Journal of Accounting and Economics 74, 101525.

DeMarzo, P., Duffie, D., 1999. A Liquidity-Based Model of Security Design. Econometrica 67, 65–100.

Döttling, R., Rola-Janicka, M., 2025. Too levered for Pigou: Carbon pricing, financial constraints, and leverage regulation. Journal of Financial Economics 172.

Elliott, J., Foster, I., Kortum, S., Munson, T., Pérez Cervantes, F., Weisbach, D., 2010. Trade and carbon taxes. American Economic Review 100, 465–469.

Fabra, N., Reguant, M., 2014. Pass-through of emissions costs in electricity markets. American Economic Review 104, 2872–2899.

Farrell, J., 1993. Meaning and credibility in cheap-talk games. Games and Economic Behavior 5, 514–531.

Fishman, M. J., Hagerty, K. M., 1990. The Optimal Amount of Discretion to Allow in Disclosure. The Quarterly Journal of Economics 105, 427–444.

Fishman, M. J., Hagerty, K. M., 2003. Mandatory Versus Voluntary Disclosure in Markets with Informed and Uninformed Customers. The Journal of Law, Economics, and Organization 19, 45–63.

Fowlie, M., Reguant, M., Ryan, S. P., 2016. Market-based emissions regulation and industry dynamics. Journal of Political Economy 124, 249–302.

Gal-Or, E., 1985. Information Sharing in Oligopoly. Econometrica 53, 329–343.

Gertner, R., Gibbons, R., Scharfstein, D., 1988. Simultaneous signalling to the capital and product markets. The RAND Journal of Economics pp. 173–190.

Glode, V., Opp, C. C., Zhang, X., 2018. Voluntary disclosure in bilateral transactions. Journal of Economic Theory 175, 652–688.

Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. Econometrica 82, 41–88.

Grossman, S. J., 1981a. The informational role of warranties and private disclosure about product quality. The Journal of law and Economics 24, 461–483.

Grossman, S. J., 1981b. The Informational Role of Warranties and Private Disclosure about Product Quality. Journal of Law and Economics 24, 461–483.

Grossman, S. J., Hart, O. D., 1980. Disclosure Laws and Takeover Bids. Journal of Finance 35, 323–334.

Grossman, S. J., Perry, M., 1986. Perfect sequential equilibrium. Journal of economic theory 39, 97–119.

Gupta, D., Starmans, J., 2024. Dynamic green disclosure requirements. Working Paper.

Guttman, I., Kremer, I., Skrzypacz, A., 2014. Not Only What but Also When: A Theory of Dynamic Voluntary Disclosure. American Economic Review 104, 2400–2420.

Hodder, L. D., Schipper, K., 2025. Standard setting for sustainability reporting. Accounting and Business Research 55, 567–599.

Huang, Y. A., Weber, C. L., Matthews, H. S., 2009. Categorization of scope 3 emissions for streamlined enterprise carbon footprinting. Environmental Science & Technology 43, 8509–8515, pMID: 20028044.

Innes, R. D., 1990. Limited liability and incentive contracting with ex-ante action choices. Journal of Economic Theory 52, 45–67.

Inostroza, N., Tsoy, A., 2025. Optimal information and security design, working Paper, University of Toronto; SSRN No. 4093333.

Kim, E. H., Lyon, T. P., 2011. Strategic environmental disclosure: Evidence from the doe's voluntary greenhouse gas registry. Journal of Environmental Economics and Management 61, 311–326.

Lutz, N. A., 1989. Warranties as signals under consumer moral hazard. The Rand journal of economics pp. 239–255.

Lyon, T. P., Maxwell, J. W., 2011. Greenwash: Corporate Environmental Disclosure under Threat of Audit. Journal of Economics & Management Strategy 20, 3–41.

Macchiavello, R., 2022. Relational contracts and development. Annual Review of Economics 14, 337–362.

Matthews, S., Postlewaite, A., 1985. Quality Testing and Disclosure. RAND Journal of Economics 16, 328–340.

Milgrom, P., 2008a. What the seller won't tell you: Persuasion and disclosure in markets. Journal of Economic Perspectives 22, 115–131.

Milgrom, P., 2008b. What the seller won't tell you: Persuasion and disclosure in markets. Journal of Economic Perspectives 22, 115–131.

Milgrom, P., Roberts, J., 1986. Relying on the Information of Interested Parties. RAND Journal of Economics 17, 18–32.

Milgrom, P. R., 1981. Good News and Bad News: Representation Theorems and Applications. Bell Journal of Economics 12, 380–391.

Nachman, D. C., Noe, T. H., 1994. Optimal Design of Securities under Asymmetric Information. The Review of Financial Studies 7, 1–44.

Nordhaus, W., 1991. To slow or not to slow: The economics of the greenhouse effect. Economic Journal 101, 920–937.

Nunn, N., 2007. Relationship-specificity, incomplete contracts, and the pattern of trade. Quarterly Journal of Economics 122, 569–600.

Okuno-Fujiwara, M., Postlewaite, A., Suzumura, K., 1990. Strategic Information Revelation. Review of Economic Studies 57, 25–47.

Orlov, D., Zryumov, P., Skrzypacz, A., 2023. The design of macroprudential stress tests. The Review of Financial Studies 36, 4460–4501.

Peters, G., Minx, J., Weber, C., Edenhofer, O., 2011. Growth in emission transfers via international trade from 1990 to 2008. Proceedings of the National Academy of Sciences 108, 8903–8908.

Pindyck, R. S., 2013. Climate change policy: What do the models tell us? Journal of Economic Literature 51, 860–872.

Poterba, J. M., 1991. Tax policy to combat global warming: On designing a carbon tax. In: Dornbusch, R., Poterba, J. M. (eds.), *Global Warming: Economic Policy Responses*, MIT Press, Cambridge, MA.

Shin, H. S., 2003. Disclosures and Asset Returns. Econometrica 71, 105–133.

UNFCCC, 2015. The paris agreement. `https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement`.

Verrecchia, R. E., 1983. Discretionary disclosure. Journal of Accounting and Economics 5, 179–194.

Verrecchia, R. E., 2001. Essays on disclosure. Journal of accounting and economics 32, 97–180.

Wagenhofer, A., 2024. Sustainability reporting: A financial reporting perspective. Accounting in Europe 21, 1–13.

Weitzman, M. L., 1974. Prices vs. quantities. Review of Economic Studies 41, 477–491.

Xue, H., 2023. Esg disclosure, market forces, and investment efficiency. Working Paper.