Environmental Ambition and Economic Protectionism: The Design of Border Carbon Adjustments

Eunseong Park* Sebastian Rausch† Valerie J. Karplus‡

March 2026

Abstract

As governments pursue climate ambition while preserving industrial competitiveness, border carbon adjustments (BCAs) have moved from theory to practice—yet jurisdictions are diverging on their fundamental architectures. Using a plant-level multi-region equilibrium model of the global steel industry, we evaluate the trade-off between mass-based BCAs, which price absolute embedded emissions, and rate-based BCAs, which price deviations from an intensity benchmark. To formalize this divergence, we introduce the effective signal ratio to measure the transmission of empirical carbon incentives. We show that rate-based designs mechanistically dilute this signal, unlocking structural reshuffling and the vertical leakage of carbon-intensive intermediates. In the EU context, a rate-based design equates to a 64% discount on the carbon price, paradoxically yielding higher national welfare losses than a mass-based system. Without a domestic price anchor—as proposed in the US—signal dilution reaches 90%. Ultimately, the choice of BCA design is not merely technical, but a strategic choice between downstream protectionism and upstream environmental integrity.

*ZEW Mannheim and University of Mannheim
†ZEW Mannheim and Heidelberg University
‡Carnegie Mellon University

1 Introduction

The intersection of global climate targets and industrial reality has created a new frontier in trade architecture: the border carbon adjustment (BCA). As jurisdictions move toward aggressive net-zero pathways, the environmental ambition of domestic carbon pricing increasingly collides with the competitive reality of global markets. This collision creates a systemic vulnerability known as carbon leakage, where emissions-intensive production relocates to regions with weaker regulations. To bridge this gap, BCAs have emerged as the definitive instrument for aligning cross-border incentives, charging imported goods for their embedded carbon to ensure a level playing field (Böhringer et al., 2012). However, beneath the shared objective of environmental integrity lies a deeper, unresolved conflict over the design of the border instrument itself.

We argue that the choice of BCA design is not merely a technical decision but a choice between distinct modes of economic protectionism targeting different segments of the industrial supply chain (Grossman and Helpman, 1994; Fischer, 2017). Governments face a fundamental trade-off between upstream defense via mass-based pricing and downstream relief via rate-based benchmarks. Mass-based BCAs, such as the European Union's Carbon Border Adjustment Mechanism (CBAM), levy a charge on the total absolute quantity of embedded emissions. This insulates domestic integrated producers from foreign volume but passes high carbon costs through to downstream users. Conversely, rate-based BCAs—the logic of which underpins U.S. data initiatives like the PROVE IT Act and proposed tariffs such as the Foreign Pollution Fee Act—charge only the intensity gap relative to a benchmark. This design acts as a hybrid tax-and-subsidy instrument (Fischer and Fox, 2012; Fowlie and Reguant, 2022; Kotchen and Maggi, 2025) that buffers manufacturers and consumers from inflationary pressure but, as we show, mechanistically dilutes the carbon signal and exposes the upstream industrial base. This design divergence reflects a systemic environmental bias in trade policy (Shapiro, 2021), where the political demand to safeguard clean downstream industries often comes at the expense of the decarbonization signal intended for the upstream sector.

This paper examines how the choice between these policy architectures—absolute versus relative trade measures—mediates the transmission of carbon prices, the distribution of national welfare, and the integrity of the environmental signal. We ask: How do mass-based and rate-based designs alter the effective transmission of the carbon price, and what are the resulting trade-offs for climate and competitiveness? We study these questions across two dominant policy paradigms: the symmetric paradigm, where the BCA anchors a domestic price, and the asymmetric paradigm, where the BCA operates as a standalone border measure without a domestic price anchor. By identifying the incentive wedge created by these designs, we reveal that competitiveness-friendly BCAs can inadvertently neuter the carbon signal by as much as 90%, leading to higher national welfare losses than their mass-based counterparts.

We quantify these effects in the global steel industry, which serves as the definitive laboratory for testing climate-trade policy (Mathiesen and Mæstad, 2004). Steel production features two characteristics that make standard trade-climate assumptions fail: extreme technological hetero-

geneity and deep vertical linkages (Collard-Wexler and De Loecker, 2015; Duscha et al., 2019; Cohen and Valacchi, 2022). The industry is split between carbon-intensive integrated routes and cleaner, scrap-based electric arc furnaces, creating massive opportunities for reshuffling—where foreign producers game benchmarks by reallocating clean output without physical abatement. Furthermore, the vertical separation between iron-making and steel-making allows us to identify vertical leakage, where protection at the steel stage is hollowed out by unpriced carbon in the iron-ore chain. Using a plant-level global equilibrium framework, we show that these micro-foundations of production are essential to capturing the true incidence and effectiveness of border carbon adjustments.

To formalize this mechanism, we introduce the effective signal ratio (λ), which measures the fraction of the nominal carbon price actually transmitted to producers. In a mass-based system, the signal is unitary ($\lambda = 1$). In contrast, rate-based designs inherently weaken this signal by implicitly subsidizing output via intensity benchmarks (Böhringer et al., 2017). Because steel production is characterized by highly polarized asset types, rate-based designs also unlock a structural reshuffling margin. Foreign producers can reallocate their "clean" scrap-based production to the regulating market while redirecting "dirty" integrated production elsewhere, gaming the benchmark without physical abatement. Our quantitative results show this dilution is severe: under a symmetric EU-style policy, a rate-based design acts as a 64% discount on the carbon price. Under an asymmetric US-style policy lacking a domestic price anchor (Clausing and Wolfram, 2023), λ plummets to nearly 10%, rendering the border adjustment a low-cost trade defense tool rather than a viable climate instrument.

Furthermore, the multi-stage nature of steel production introduces a second vulnerability: the vertical leakage trap. Standard aggregate models treat sectors as monolithic, missing how intensity benchmarks applied to final manufacturing (e.g., steel) create a tax shield for carbon-intensive upstream inputs (e.g., iron). We find that while this downstream relief successfully preserves the scale margin of production and sustains consumer surplus, it comes at a massive cost to the technology margin of decarbonization (Antweiler et al., 2001). By failing to price upstream embedded carbon, the rate-based design forces a larger abatement burden onto domestic producers while subsidizing foreign competitors. Consequently, the rate-based policy paradoxically yields higher national welfare losses ($1.38 billion) than an aggressive mass-based system ($0.89 billion), locking in long-term structural dependencies on unpriced foreign carbon.

Our findings offer three critical implications for the evaluation and design of climate-trade policy. First, we establish that border adjustments and domestic carbon pricing function as institutional complements. When deployed without a domestic price anchor, a rate-based border measure operates primarily as a rent-shifting tariff: it shields domestic producers but transmits a deeply discounted abatement signal abroad, rendering it an inefficient instrument for global decarbonization. Second, we demonstrate that the micro-structure of the border instrument—specifically benchmark stringency and the treatment of intermediate inputs—governs this price transmission. Policies designed to protect downstream users through generous exemptions predictably generate vertical leakage and accounting reshuffling. These mechanisms are not unique to steel; they apply directly

to other energy-intensive, trade-exposed sectors with substantial technological heterogeneity, such as cement and aluminum (Fowlie et al., 2016). Finally, our results emphasize the methodological necessity of incorporating plant-level heterogeneity and vertical supply chain linkages into trade policy analysis. Standard aggregate models that assume representative firms and monolithic sectors naturally omit the reshuffling margin and the vertical leakage trap, thereby systematically overestimating the environmental efficacy and welfare benefits of performance-based border benchmarks.

We advance three active literatures at the intersection of international trade and environmental economics (Copeland and Taylor, 2004; Felbermayr et al., 2025; Fontagné and Schubert, 2023). First, we bridge the gap between the macroeconomic theory of optimal climate-trade policy and the micro-design of real-world border instruments. A frontier theoretical literature explores how unilateral carbon pricing and border adjustments can operate as optimal macro-tariffs to deter free-riding, foster climate cooperation, and shift environmental rents (Copeland, 1996; Keen and Kotsogiannis, 2014; Nordhaus, 2015; Weisbach et al., 2023; Beaufils et al., 2024; Farrokhi and Lashkaripour, 2025; Itskhoki and Mukhin, 2025; Keen and Kotsogiannis, 2025; Kortum and Weisbach, 2026). While recent work proposes alternative border instruments to bypass the severe information asymmetries inherent in calculating true foreign carbon intensities (Campolmi et al., 2025), we build on these aggregate frameworks by analyzing the political-economic reality of instrument choice: the divergence between mass-based and rate-based architectures. By formally evaluating this signal dilution, we demonstrate how the micro-structure of benchmark exemptions unravels the macroscopic efficiency gains predicted by optimal tariff models.

Second, we provide a structural resolution to an emerging empirical puzzle regarding carbon leakage. While earlier retrospective studies found limited evidence of production relocation under initial carbon pricing regimes (Dechezleprêtre et al., 2022; Akerman et al., 2024), recent firm-level evidence has uncovered significant supply chain adaptation and carbon off-shoring (Ben-David et al., 2021; Coster et al., 2024). We reconcile this discrepancy by showing that leakage in hard-to-abate sectors is heavily skewed toward intermediate inputs. Because standard aggregate and sectoral studies fail to capture this granular supply-chain integration, they systemically under-count the true magnitude of carbon leakage, which simply shifts upstream in the supply chain.

Finally, we fundamentally expand the role of firm heterogeneity in environmental trade models (Shapiro and Walker, 2018; Sogalla et al., 2024; Sogalla, 2025). While recent work leverages plant-level data primarily to evaluate distributional incidence (Clausing et al., 2025) or explore intra-industry reallocation via intermediate imports (Stillger, 2025), we show that this profound structural variance is not merely an adjustment mechanism—it is the structural engine of international policy evasion. By embedding detailed micro-foundations of production within a multi-region general equilibrium framework, we reveal how intra-industry variance creates massive opportunities for accounting reshuffling. This integration of vertical iron-to-steel linkages into a multi-sector environment allows us to quantify evasion mechanisms that remain entirely invisible to both aggregate macro-models and partial equilibrium industry studies alike.

The rest of the paper is organized as follows. Section 2 provides background on the global steel industry and its vertical structure. Section 3 develops the conceptual framework and characterizes the trade-offs between mass-based and rate-based BCAs. Section 4 develops the quantitative model and characterizes the effective signal ratio. Section 5 discusses the data and plant-level calibration. Section 6 describes the policy scenarios for the symmetric and asymmetric paradigms. Section 7 presents our quantitative results on signal dilution, welfare, and leakage. Section 8 concludes with policy implications.

2 Industry Background

The global steel sector is a canonical emissions-intensive, trade-exposed industry: it accounts for roughly 7% of global greenhouse gas (GHG) emissions and nearly 25% of global steel production crosses international borders each year (World Steel Association, 2023). These features make steel a natural setting in which to study how border carbon adjustments (BCAs) transmit carbon costs through international supply chains.

Three industry characteristics are especially important for interpreting the policy trade-offs we quantify in Section 7. First, steel production is technologically heterogeneous, implying large dispersion in emissions intensity across plants and regions. Second, steel is produced through a vertically linked chain in which carbon-intensive intermediate inputs can be traded separately from downstream steelmaking. Third, there is a strong regional correlation between production costs and emissions intensities. Together, these features determine how mass-based versus rate-based BCAs translate into effective incentives at the margin.

2.1 Technologies, Inputs, and Emissions Heterogeneity

Steelmaking relies on multiple production routes with different raw material inputs, technologies, and emissions intensities. The blast furnace–basic oxygen furnace (BF-BOF) route accounts for around 70% of global production. It uses iron ore and coke to produce hot metal (or pig iron, when cooled) in a blast furnace before refining it in a basic oxygen furnace. This integrated process is efficient at scale but rigid in raw material choices; the energy savings of molten charging ensures that the majority of iron production is captive to specific steelmaking sites, with international trade largely restricted to solidified merchant iron units. The BF-BOF route averages roughly 2.2 tonnes of CO_2 per tonne of crude steel. While scrap can be added to BOF to reduce emissions, its use is limited (up to around 20%) because impurities affect final steel properties and excessive scrap lowers furnace temperature.

In contrast, electric arc furnace (EAF) steelmaking relies primarily on scrap metal and accounts for about 30% of global production. In the US, abundant scrap availability, relatively low electricity prices, and continued technological advances have driven EAFs to account for around 70% of production. When combined with relatively clean electricity, EAFs can achieve emissions as low as 0.3 tonnes of CO_2 per tonne of steel. In practice, many EAF producers blend primary iron inputs—

such as pig iron or direct reduced iron (DRI)—with scrap to meet higher-grade requirements. One notable variant is the DRI-EAF route, which averages about 1.4 tonnes of CO_2 per tonne of steel and is prevalent in regions with abundant, low-cost natural gas.

Although scrap-based EAFs offer comparatively low emissions, relying solely on scrap faces two constraints. First, global scrap supply is limited by the long service life of steel products (often around 40 years), so scrap is insufficient to meet total demand. Second, scrap quality varies, making it difficult to produce some advanced or high-strength grades without adding purer iron inputs.

2.2 Supply Chain Integration and the Merchant Iron Market

Steel products vary by grade and form, typically grouped into flat and long products with distinct applications. From an economic perspective, the key feature is that steel is not a single-stage process. While the high thermal energy required to keep iron molten ensures that most ironmaking occurs on-site at integrated facilities, steelmaking also utilizes merchant iron inputs purchased on the global market. In our setting, the most relevant intermediates are pig iron and DRI, which can be used as primary iron units in downstream steelmaking alongside steel scrap. Integrated BF-BOF production relies on molten iron charging and therefore has limited scope to substitute away from on-site ironmaking, while EAF producers can flexibly blend scrap with merchant pig iron or DRI to meet quality and grade requirements.

This vertical structure matters for BCAs because it creates multiple policy touchpoints: a border instrument can be applied to finished steel, to upstream iron inputs, or to both. In turn, different designs can shift incentives across stages of the supply chain and interact with firms' ability to substitute across iron inputs and sourcing locations.

2.3 Regional Clustering and Comparative Costs

Production technologies and cost structures vary significantly by region, driven by resource availability and industrial history. Figure 1 illustrates the geographic clustering of these technologies. With the exception of China, long products are mainly produced by EAF mills. Flat products are predominantly produced by BF-BOF plants, with exceptions including the US, India, and some Western Asian producers where EAF shares are substantial (Figure 1). The rise of EAFs in the US is documented by Collard-Wexler and De Loecker (2015).

In the United States, a larger share of costs comes from third-party scrap purchases in both long and flat production, reflecting widespread EAF use and a well-developed scrap recycling infrastructure (Figure 2). In contrast, regions dominated by BF-BOF face higher costs for iron ore and reductant precursors such as coking coal. Figure 3 extends this to emissions intensity: while North American and European plants have higher production costs, they achieve lower CO_2 intensities due to cleaner technologies and tighter regulations. Conversely, plants in Asia combine lower costs with higher emissions. These joint cost–emissions differences underpin the scale–composition trade-off we quantify: stringent policies can shift production toward cleaner

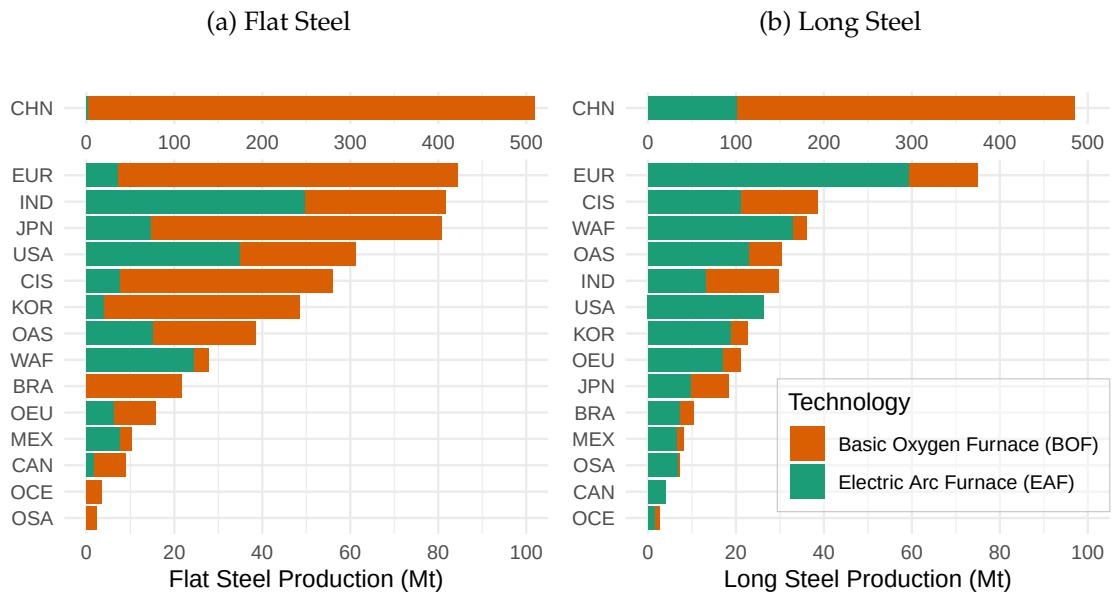


Figure 1: Steel Production by Technology and Region

Note: Bars report production by region, with fill indicating technology (BOF vs. EAF); production is in Mt. Region codes are defined in Section 5.1.
Source: Own calculation based on CRU Steel Cost Service and World Steel Association Statistical Yearbooks.

(but higher-cost) technologies and regions, but the magnitude of global emissions reductions also depends on the extent to which demand is suppressed.

2.4 Implications for Border Carbon Adjustments

These technological and regional disparities are central to BCA design because they determine both (i) the dispersion of embedded emissions that a policy could in principle target and (ii) the opportunities for strategic reallocation within multinational supply chains. Mass-based BCAs price total embedded emissions and therefore preserve a direct link between the border charge and the full emissions content of traded goods. Rate-based BCAs, in contrast, price only the gap relative to a benchmark; the resulting marginal wedge depends on where the benchmark is set and how much emissions-intensity dispersion exists at that stage of production.

Finally, heterogeneity in emissions intensity across plants also creates scope for reshuffling: exporters can disproportionately allocate their cleanest output to the regulated destination while diverting dirtier production to unregulated markets. This implies that measured reductions in embedded emissions in the regulated market may overstate true global abatement, a point that is difficult to assess without a framework that tracks plant-level production and trade flows. These industry features motivate the key modeling choices in Section 4: a plant-level representation of heterogeneous technologies, explicit treatment of the merchant intermediate markets alongside finished steel, and limited substitution across product categories.

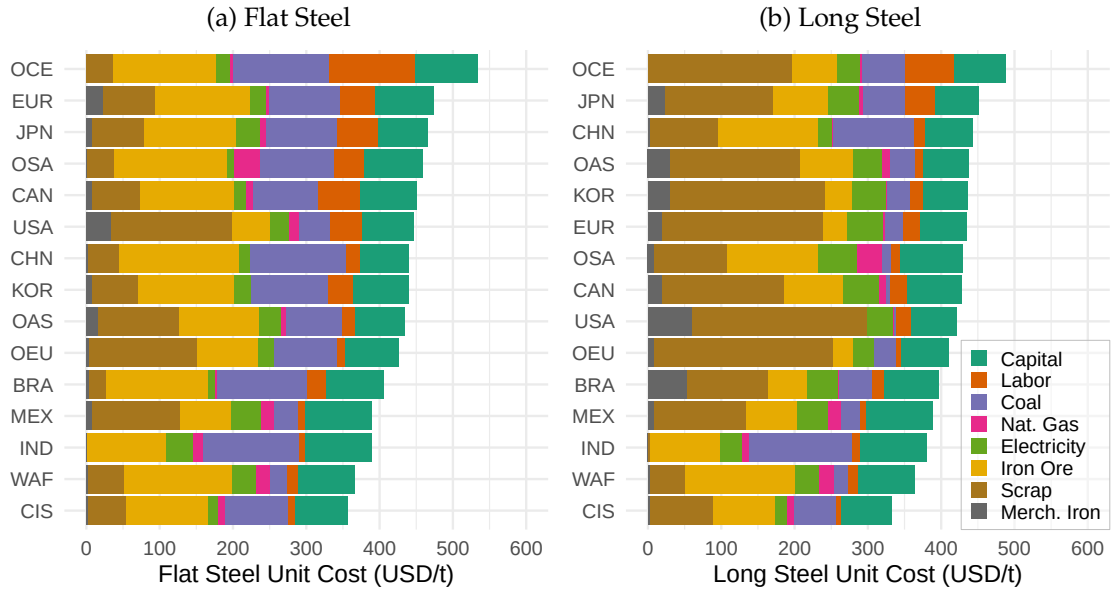


Figure 2: Average Steel Production Cost Structure by Region

Note: Bars decompose unit costs (USD/t) by input category for each region. Region codes are defined in Section 5.1.
Source: Own calculation based on CRU Steel Cost Service.

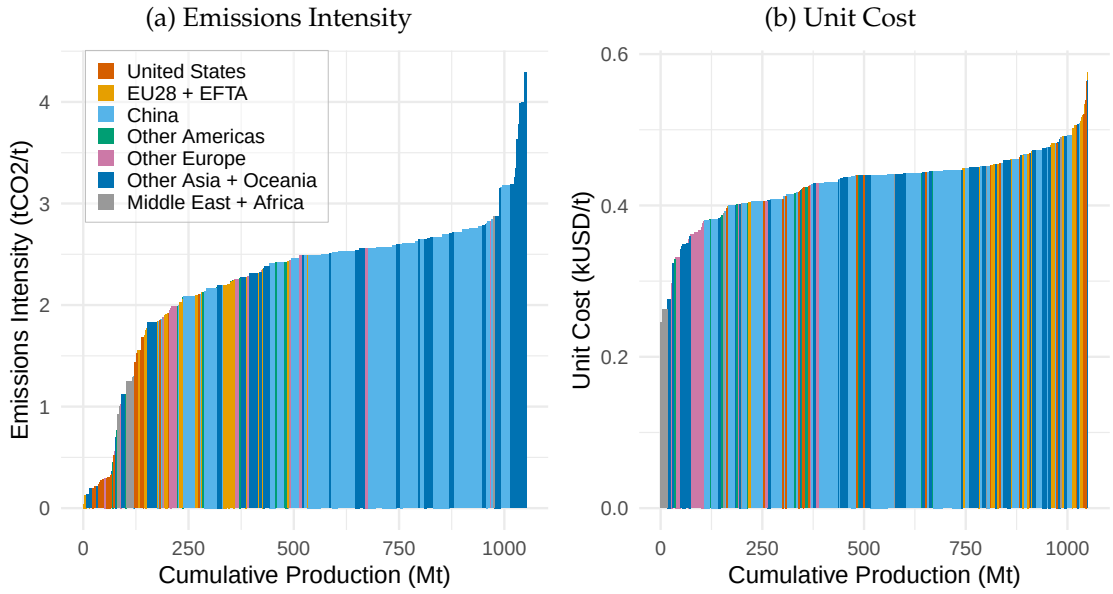


Figure 3: Cost and Emissions Intensity Curve for Flat Steel

Note: Each panel plots flat-steel plants ordered by cumulative production. Bars are colored by region code (see legend), bar width equals plant production, and cumulative production is in Mt. The left panel reports emissions intensity (tCO_2/t) and the right panel reports unit cost (USD/t). Region codes are defined in Section 5.1.
Source: Own calculation based on CRU Steel Cost Service.

3 Conceptual Framework

We provide a stylized framework to characterize the economic mechanisms distinguishing mass-based and rate-based BCAs. This framework formalizes the strategic trade-offs inherent in the choice of BCA design, mapping instrument parameters directly to abatement channels, supply chain incidence, and leakage dynamics. By reordering the theoretical analysis to mirror the efficiency, mechanism, and distribution margins, we provide a structural roadmap for the quantitative simulation results. Formal mathematical derivations for all lemmas and propositions are provided in Appendix A.

3.1 Setup

We consider a vertically linked economy with two sectors: an upstream sector and a domestic downstream sector. The upstream sector involves two regions, home (h) and foreign (f). To capture the full range of policy proposals, we distinguish between the domestic carbon price $\tau_h \geq 0$ faced by home producers and the border adjustment rate $\tau_f > 0$ applied to imports from the foreign region. This framework encompasses the symmetric paradigm ($\tau_h = \tau_f$, modeled on the EU) and the asymmetric paradigm ($\tau_h = 0$, modeled on the US).

Upstream Sector Each regional upstream sector is characterized by a representative producer managing a portfolio of clean (c) and dirty (d) technologies with emissions intensities $\epsilon_c < \epsilon_d$ and base unit production costs $c_c > c_d$. The producer's sectoral unit cost of production $c(s_{c,i})$ depends on the regional aggregate clean technology share $s_{c,i} \in [0,1]$ and is strictly increasing and convex ($c' > 0, c'' > 0$), reflecting the weighted average of base technology costs and the increasing marginal costs of technology integration and site-specific constraints.

Home upstream producers supply volume $q_{h,h}$ exclusively to the regulated domestic market, while foreign upstream producers serve both the regulated market ($q_{f,h}$) and an unregulated foreign domestic market ($q_{f,f}$). As a benchmark case, we consider inelastic foreign domestic demand $q_{f,f}$ to isolate the policy-driven trade-offs between technology adoption and trade scale within the regulated market. Let $\omega_{i,j} = q_{i,j}/(q_{i,h} + q_{i,f})$ denote the share of destination $j \in \{h, f\}$ in the total production of region i, where $\omega_{h,h} = 1$ and $\omega_{h,f} = 0$. The border charge $t(\cdot)$ applies only to the volume sold in the regulated market ($j = h$) based on its specific intensity $\epsilon_{i,h} = s_{c,i,h}\epsilon_c + (1 - s_{c,i,h})\epsilon_d$. The representative producer in region i chooses the clean technology shares for each market ($s_{c,i,h}, s_{c,i,f}$) to minimize the unit cost of the variety sold in the regulated market:

$$\min_{s_{c,i,h}, s_{c,i,f}} \{c(s_{c,i}) + t(\epsilon_{i,h})\} \tag{1}$$

where $s_{c,i} = \sum_j \omega_{i,j} s_{c,i,j}$ is the regional clean share. The instrument $t(\epsilon)$ is determined by the policy design: $t^M(\epsilon) = \tau\epsilon$ under a mass-based design, and $t^R(\epsilon) = \tau(\epsilon - \bar{\psi})$ under a rate-based design, where $\bar{\psi}$ denotes the intensity benchmark. Economically, the rate-based design is equivalent to a

mass-based tax combined with a specific output subsidy $\tau\bar{\psi}$. In equilibrium, the price $p_{i,h}$ of the regional variety in the regulated market is $p_{i,h} = c(s_{c,i}) + t(\epsilon_{i,h})$.

Downstream Sector Domestic downstream firms combine domestic ($q_{h,h}$) and foreign ($q_{f,h}$) upstream varieties to produce a composite intermediate Q. The aggregation follows a linearly homogeneous and concave function $Q = F(q_{h,h}, q_{f,h})$. Let $P(p_{h,h}, p_{f,h})$ denote the associated dual price index, representing the minimum cost of producing one unit of Q. This index is strictly increasing ($\partial P/\partial p_{i,h} > 0$) and concave ($\partial^2 P/\partial p_{i,h}^2 \leq 0$) in input prices. Downstream firms choose Q to maximize producer surplus $\Pi_d = \max_Q\{G(Q) - P(p_{h,h}, p_{f,h})Q\}$, where $G(Q)$ represents the value-added of the upstream input in downstream production. By Shephard's Lemma, the demand for each variety is given by $q_{i,h} = Q \cdot \partial P/\partial p_{i,h}$. The concavity of P ensures that demand for each variety is strictly decreasing in its own price $p_{i,h}$. This general structure can be specialized to a Constant Elasticity of Substitution (CES) functional form, consistent with standard Armington trade assumptions.

The core economic distinction between the mass-based and rate-based architectures lies in the implicit subsidy created by the intensity benchmark. Rather than acting as a simple carbon price, the rate-based design functions as a hybrid instrument that pairs a tax on emissions with a specific subsidy for production. The following lemma formalizes the resulting ranking of equilibrium prices, which serves as the foundation for the subsequent analysis.

Lemma 1 (Price Ranking and Signal Dilution)**.** *The equilibrium price under a rate-based design is strictly lower than under an equivalent mass-based design: $p_{i,h}^R < p_{i,h}^M$. Let $\lambda = (1 - \bar{\psi}/\epsilon_{i,h})$ denote the effective signal ratio. The effective carbon price faced by imports under a rate-based design is $t^R = \lambda\tau\epsilon_{i,h}$, where $\lambda < 1$ represents the inherent dilution of the carbon signal created by the intensity benchmark.*

Proof. See Appendix A.1. □

This price wedge fundamentally distorts the transmission of the carbon price, effectively decoupling technology adoption from the consumption margin. By lowering the equilibrium price of carbon-intensive imports, the benchmark ensures that trade volumes remain high, preventing the market from utilizing demand reduction as a low-cost path to global emissions targets.

3.2 Global Efficiency and Environmental Equivalence

The efficiency of any climate policy depends on its ability to internalize the social cost of carbon across all margins of the economy. A core challenge in comparing mass-based and rate-based designs is that they create inherently different levels of stringency for a given nominal carbon price. We formalize this difference by defining λ as the *effective signal ratio*, which measures the fraction of the full carbon cost that is actually applied at the border.

Proposition 1 (Global Cost-Effectiveness)**.** *For a fixed global emissions target $\bar{E}$, the mass-based BCA is strictly more cost-effective than the rate-based BCA. The rate-based design induces a distorted allocation of*

abatement by neutralizing the consumption margin, forcing all abatement onto the technology margin where marginal costs are increasing and convex.

Proof. See Appendix A.2. □

Comparing the first-order conditions for technology adoption and consumption reveals why the rate-based design is strictly less efficient. Because the output subsidy prevents the equalization of marginal abatement costs across margins, the technology margin is forced to bear a disproportionate share of the burden. The proof relies on the fact that the rate-based subsidy necessitates a higher marginal carbon price to hit the same global target, pushing technology switching into a region of sharply increasing marginal costs while leaving low-cost abatement from reduced consumption unexploited.

3.3 Abatement Mechanisms

Following the established efficiency gap, we characterize how production and emissions intensity respond to the policy. Abatement occurs along two primary margins: the scale margin (total volume) and the composition margin (technology mix). We focus on how the BCA design determines whether a target is met by shrinking the industry or by cleaning its technological composition.

Proposition 2 (Foreign Composition–Scale Substitution)**.** *Consider a comparison between a rate-based design and a mass-based design achieving the same global emissions reductions. For a fixed domestic emissions target, the rate-based design induces a strictly larger compositional shift in the foreign sector (higher $s_{c,f}$). Because the rate-based design preserves trade volumes more effectively than a mass-based design of equivalent stringency, it must compensate with a deeper reduction in the technological composition of production to satisfy the global emissions budget.*

Proof. See Appendix A.3. □

The derivation of this substitution effect relies on the constant global emissions budget. By establishing that the rate-based instrument sustains a larger trade volume, the results show that the average sectoral intensity must be strictly lower to satisfy the same aggregate cap, which in turn necessitates a higher clean technology share. While technology switching is the ultimate response, compositional shifts are not the only way for a foreign producer to clean up its exports. For a producer managing a portfolio of plants, the first and least expensive option is to reallocate existing clean units to the regulated market. This reshuffling behavior represents a major challenge for border standards.

Lemma 2 (Reshuffling Priority)**.** *In a perfectly competitive equilibrium, foreign producers will always prioritize the allocation of upstream production from clean technology plants to the regulated market to minimize the border charge. This priority rule holds whenever the carbon price makes the tax savings from the intensity gap outweigh the costs of technology integration.*

Proof. See Appendix A.4. □

Cost-minimizing producers will prioritize their cleanest units for the regulated market whenever the carbon price makes the tax savings from the intensity gap outweigh the costs of technology integration. This establishes that exporters will first exhaust their lowest-intensity production units to satisfy border standards. However, the efficacy of this easy abatement depends on the total volume of exports. If trade scale remains high, the existing pool of clean units may prove insufficient to meet the intensity target for the entire export volume, leading to physical constraints on the reshuffling margin.

Proposition 3 (Reshuffling and Capacity Pressure)**.** *For a fixed import emissions intensity target $\bar{\psi}_{f,h}$, the rate-based BCA drives a larger expansion of clean technology capacity in the foreign sector. Because the rate-based design sustains higher export volumes $q_{f,h}$, satisfying the intensity constraint at the border requires a larger absolute quantity of clean production, which exceeds the reshuffling potential of the baseline technology mix.*

Proof. See Appendix A.5. □

Identifying the absolute quantity of clean production required for a specific border standard shows that higher volumes eventually outstrip the capacity to reshuffle. This volume-driven pressure creates a physical limit on the reshuffling margin, eventually forcing the producer into real technology switching to meet the standard. The proof demonstrates that an intensity-based target becomes harder to satisfy as total volume grows, eventually forcing the producer to invest in new clean plants rather than simply reallocating existing ones.

3.4 Distributional Incidence and Leakage

The physical mechanisms of abatement described above build a causal chain toward specific distributional outcomes. The strategic choice of a BCA architecture determines how the economic burden of climate action is shared between upstream sectors, downstream firms, and governments. We begin by characterizing the fiscal implications of these designs, which reveal strategic trade-offs in the regional distribution of economic surplus.

Proposition 4 (Regional Incidence and Terms-of-Trade)**.** *The choice of BCA design determines the distribution of economic surplus between regions. For a fixed marginal carbon price τ_f, the mass-based design improves the Home region's terms-of-trade by capturing scarcity rents from foreign producers, while the rate-based design foregoes this fiscal revenue to minimize downstream costs.*

Proof. See Appendix A.6. □

Decomposing regional welfare into producer surplus, consumer surplus, and fiscal revenue reveals that the loss from the benchmark subsidy typically outweighs the gain in downstream surplus for nations where domestic industry value-added is small. For an implementing region, this choice also dictates the locus of economic protection within the domestic supply chain. The intensity benchmark acts as a direct transfer of scarcity rents from the government to the downstream sector,

effectively shielding the consumer but shrinking the competitive umbrella that protects domestic upstream output.

Proposition 5 (Industrial Incidence)**.** *The choice of BCA design determines the distribution of economic surplus between the upstream sectors and the domestic downstream sector:*

(i) Upstream: Mass-based designs maximize the competitive shield for domestic producers by increasing import prices $p_{f,h}$, thereby increasing their producer surplus: $\Pi_{u,h}(p^M_{f,h}) > \Pi_{u,h}(p^R_{f,h})$.

(ii) Downstream: Rate-based designs minimize the aggregate input price index $P(p_{h,h}, p_{f,h})$, thereby increasing the surplus of downstream firms: $\Pi_d(P^R) > \Pi_d(P^M)$.

Proof. See Appendix A.7. □

Applying the Envelope Theorem to the dual price index reveals how this benchmark subsidy lowers input costs for users while simultaneously increasing competitive pressure on domestic integrated mills. These findings highlight a fundamental conflict between competing notions of economic protection: a mass-based design protects upstream output, while a rate-based design protects downstream prices. In the asymmetric paradigm, where the lack of domestic pricing turns the BCA into a purely protectionist tariff, this can even lead to reshoring effects.

Proposition 6 (Reshoring and Diluted Abatement)**.** *In the asymmetric paradigm ($\tau_h = 0, \tau_f > 0$), the BCA induces reshoring of domestic production. The net impact on global emissions is determined by the trade-off between the abatement effect in the foreign sector and the substitution effect toward dirtier domestic varieties. Global emissions reductions are diluted if the reshoring of unabated domestic production is large relative to the technology response of foreign producers.*

Proof. See Appendix A.8. □

While the baseline setup considers a single-stage upstream sector, we now extend this logic to a multi-stage production chain to characterize how the potential for abatement varies across products. Rate-based designs rely on sufficient technological variance to generate a continuous relative abatement signal across firms; however, if the underlying technology is too uniform, this benchmark-driven mechanism can fail. This is particularly relevant for homogeneous intermediate products like iron, where technological variance is low.

Lemma 3 (The Knife-Edge Property of Homogeneous Sectors)**.** *In a rate-based BCA, the implicit marginal penalty δ^R becomes infinitely sensitive to the benchmark level $\bar{\psi}$ as the technology-specific emissions intensities ϵ_c and ϵ_d converge. For a sufficiently homogeneous sector ($\epsilon_c \to \epsilon_d$), any benchmark set below the physical limit of the technology leads to widespread compliance failure driving the penalty to its maximum ($\delta^R \to \infty$), while any benchmark set above it neutralizes the price signal ($\delta^R \to 0$).*

Proof. See Appendix A.9. □

The extreme sensitivity of the effective border penalty arises because, in homogeneous sectors, there is no technological margin available to absorb changes in the benchmark. This creates a binary outcome: if the standard is set even slightly below the physical limits of the technology, widespread compliance failure drives the implicit penalty to its maximum, whereas a standard set above this limit neutralizes the border signal. In the following proposition, we formalize how this lack of a price signal at the border induces leakage along the vertical production chain.

Proposition 7 (Vertical Leakage)**.** *Consider a multi-stage production chain where an intermediate input is technologically homogeneous ($\epsilon_c \approx \epsilon_d \approx \bar{\psi}$). A rate-based BCA applied to the intermediate induces vertical leakage. Domestic producers of the intermediate face the full cost of domestic carbon pricing τ_h, while foreign producers face a near-zero effective border tax. This price wedge incentivizes downstream firms to substitute domestic intermediates with imports, shifting the carbon-intensive stage of production offshore.*

Proof. See Appendix A.10. □

The derivation of the vertical leakage result demonstrates that when technological variance is low, the compliance gap $(\epsilon - \bar{\psi})$ collapses toward zero. This effectively turns off the border tax, leaving domestic producers exposed to the full cost of domestic carbon pricing while foreign competitors face negligible charges. This creates a price wedge that specifically penalizes domestic upstream production, leading downstream firms to substitute domestic inputs with higher-carbon imports. Together, these theoretical results build a causal chain from the initial price subsidy to the final distributional outcomes, providing the structural foundation for the plant-level simulation results presented in the subsequent sections.

4 Quantitative Model

We develop a multiregional, multisector general equilibrium model that accounts for the granular structure of the global steel sector. Production occurs across multiple regions and plants, each characterized by specific technology portfolios and emissions intensities. This separation allows us to analyze how climate policies drive technology adoption and emissions abatement at the production unit level. The model also incorporates product differentiation within the steel sector, limiting substitution across steel products to capture demand-side rigidities and the downstream propagation of carbon costs.

Figure 4 illustrates the high-level structure of the model, mapping the circular flow of factors and goods across regions. The framework connects households, which supply factors and demand final goods, to a granular production layer. This layer distinguishes between ironmaking and steelmaking stages, capturing the choice between integrated production and merchant input markets. Border carbon adjustments are implemented at the trade interface, mediating the price signals transmitted between global markets and individual production units. In the following subsections, we formalize the nested functional forms and equilibrium conditions that govern these linkages.

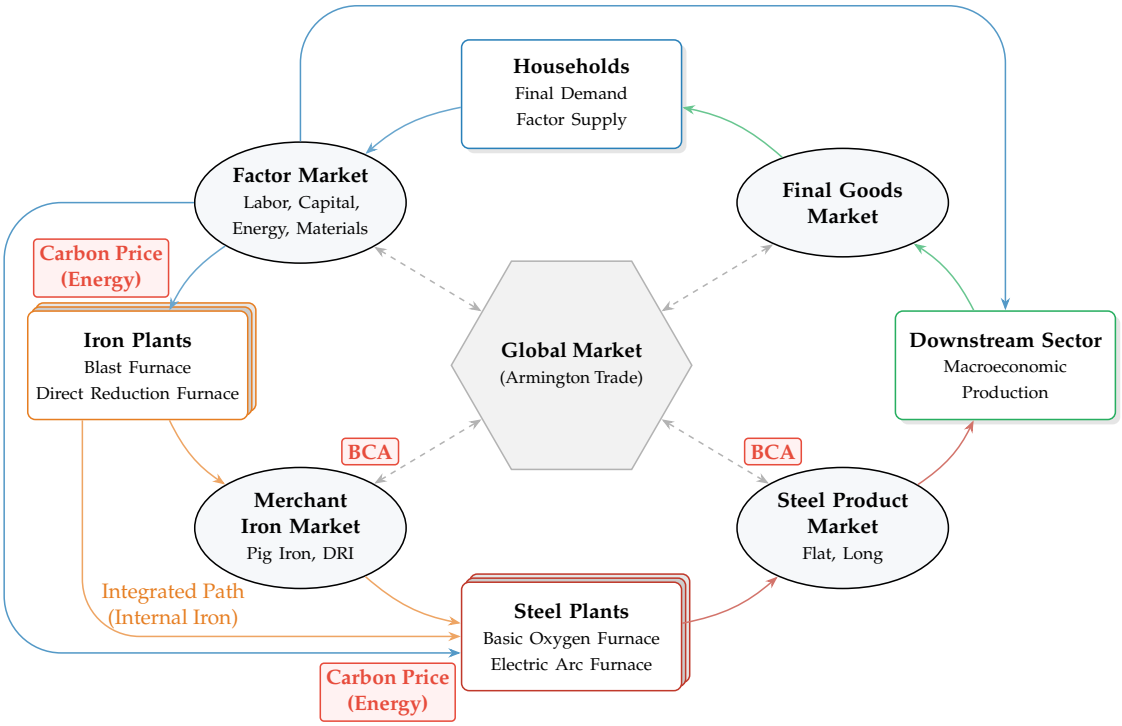


Figure 4: High-level structure of the multi-regional, multi-sector general equilibrium model.

4.1 The Economic Environment

The global economy is represented as a static multi-region equilibrium consisting of a discrete set of regions $r, s \in \mathcal{R}$. Within each region, we distinguish between a granular iron and steel industry and an aggregate macroeconomic sector that provides intermediate inputs and represents the broader economic environment.

Factor Endowments and Elastic Supply Each region r is endowed with a set of primary factors $f \in \mathcal{F}$, including labor, capital, and natural resources. We distinguish between factors in fixed supply and tradable raw materials (e.g., iron ore, coking coal, and steel scrap) that exhibit price-elastic supply. For elastically supplied inputs f, the regional market price $\Omega_{r,f}$ is determined by a nested CES price index:

$$\Omega_{r,f} = \left[\theta^R_{r,f}(\omega_{r,f})^{1-\sigma^R_f} + (1-\theta^R_{r,f})(\Omega_{r,K})^{1-\sigma^R_f}\right]^{\frac{1}{1-\sigma^R_f}}, \tag{2}$$

where $\omega_{r,f}$ is the return to the fixed regional resource, $\Omega_{r,K}$ is the regional price of capital, $\theta^R_{r,f}$ is the benchmark value share, and σ^R_f is the substitution elasticity. This structure implies a regional supply elasticity $\eta_{r,f} = \sigma^R_f((1-\theta^R_{r,f})/\theta^R_{r,f})$.

Regional Households and Consumption A representative household in each region r receives income I_r from factor returns, the benchmark trade imbalance $\bar{\Delta}_r$, and redistributed policy revenues:

$$I_r = \sum_{f \in \mathcal{F}} \omega_{r,f} E_{r,f} + \bar{\Delta}_r + TR_r, \tag{3}$$

where $E_{r,f}$ is the resource endowment and TR_r represents total transfers from carbon pricing and border adjustments. Households spend this income on the aggregate final consumption bundle $Q_{r,C}$:

$$P_{r,C} Q_{r,C} = I_r, \tag{4}$$

where the consumer price index $P_{r,C}$ is endogenously determined by the underlying macro-production technology.

Input Mobility and Plant-Level Allocation While primary factors and tradable intermediate inputs (e.g., energy carriers and raw materials) are cleared at the regional level, we allow for idiosyncratic price dispersion at the plant level to reflect local market frictions. Regional aggregate supplies—whether primary endowments or Armington composites—are transformed into plant-specific inputs ($\omega_{i,a,f}$) via a CET frontier:

$$\omega_{i,a,f} = \Omega_{r,f} \left[\theta^F_{i,a,f} \left(\frac{q_{i,a,f}}{Q_{r,f}} \right) \right]^{\frac{1}{\eta^F_f}}, \tag{5}$$

where $q_{i,a,f}$ is the quantity of input f allocated to plant i and η^F_f is the transformation elasticity governing mobility. This captures the *sluggishness* in reallocating specialized resources, including localized constraints in logistics or contracting for energy and materials, between specific industrial sites.

The Industrial Layer: Plants, Technologies, and Varieties The iron and steel sector is composed of a set of individual production facilities $i \in \mathcal{I}$. Each plant i operates a portfolio of available technologies $a \in \mathcal{A}_i \subseteq \mathcal{A}$. We explicitly partition the metallurgical supply chain into ironmaking and steelmaking stages:

$$\mathcal{A} = \mathcal{A}_I \cup \mathcal{A}_S, \quad \text{where} \quad \mathcal{A}_I = \{\text{BF}, \text{DRF}\} \quad \text{and} \quad \mathcal{A}_S = \{\text{BOF}, \text{EAF}\}. \tag{6}$$

The corresponding outputs are partitioned into intermediate iron varieties $v \in \mathcal{V}_I = \{\text{Pig Iron}, \text{DRI}\}$ and finished steel varieties $v \in \mathcal{V}_S = \{\text{Flat}, \text{Long}\}$. Plants act as price-takers under perfect competition, with their output levels determined by site-specific marginal costs.

4.2 Production and Pricing

The iron and steel industry is a classic hard-to-abate sector, characterized by a heavy reliance on carbon-intensive inputs and limited short-term substitutability between fossil fuels and clean capital. To capture these rigidities, production at the plant level is governed by a multi-tier nested CES cost structure. This formulation enforces the physical limits of metallurgy while allowing for price-induced substitution across energy, capital, labor, and materials.

Primary Production and Unit Costs At the plant level i using technology $a \in \mathcal{A}$, producers combine a material composite with a capital-labor-energy (KLE) bundle to produce semi-finished iron or steel. Under the assumption of perfect competition and constant returns to scale, the producer price of semi-finished products is equal to its unit cost $c_{i,a}(\boldsymbol{\omega}_{i,a,f}, p_{i,a,e})$. The top-level unit cost is defined as:

$$c_{i,a} = \left[\theta_{i,a}^{\text{Mat}}(p_{i,a}^{\text{Mat}})^{1-\sigma_a} + (1-\theta_{i,a}^{\text{Mat}})(p_{i,a}^{\text{KLE}})^{1-\sigma_a}\right]^{\frac{1}{1-\sigma_a}} \tag{7}$$

where $p_{i,a}^{\text{Mat}}$ is the price of the material composite, $p_{i,a}^{\text{KLE}}$ is the cost of the KLE bundle, $\theta_{i,a}^{\text{Mat}}$ is the benchmark material value share, and σ_a is the top-level elasticity of substitution, calibrated to a near-Leontief value to reflect metallurgical mass balance. The cost of the KLE bundle, $p_{i,a}^{\text{KLE}}$, aggregates value-added and energy:

$$p_{i,a}^{\text{KLE}} = \left[\theta_{i,a}^{\text{VA}}(p_{i,a}^{\text{VA}})^{1-\sigma_a^{\text{KLE}}} + (1-\theta_{i,a}^{\text{VA}})(p_{i,a}^{\text{Ene}})^{1-\sigma_a^{\text{KLE}}}\right]^{\frac{1}{1-\sigma_a^{\text{KLE}}}} \tag{8}$$

where $p_{i,a}^{\text{VA}}$ is the value-added cost, $p_{i,a}^{\text{Ene}}$ is the composite energy price, $\theta_{i,a}^{\text{VA}}$ is the benchmark value-added share, and σ_a^{KLE} is the substitution elasticity between energy and value-added. The value-added cost $p_{i,a}^{\text{VA}}$ is a composite of plant-specific factor prices $\boldsymbol{\omega}_{i,a,f}$ (e.g., labor and capital), while the energy bundle cost $p_{i,a}^{\text{Ene}}$ aggregates energy carriers $e \in \mathcal{E}$, where $\mathcal{E} = \{\text{coal}, \text{natural gas}, \text{electricity}\}$.

Input demand can be derived by differentiating the unit cost function with respect to input prices according to Shephard's Lemma:

$$q_{i,a,f} = Q_{i,a}\frac{\partial c_{i,a}}{\partial \omega_{i,a,f}}, \tag{9}$$

where $Q_{i,a}$ is the total activity level of the plant. Crucially, the user cost of each energy carrier $p_{i,a,e}$ reflects the regional carbon price τ_r:

$$p_{i,a,e} = \omega_{i,a,e} + \tau_r \epsilon_{i,a,e}, \tag{10}$$

where $\omega_{i,a,e}$ is the plant-specific pre-tax cost of energy carrier e (governed by the mobility frontier defined in 5), τ_r is the regional carbon price, and $\epsilon_{i,a,e}$ is the carbon intensity per unit of energy (measured in tCO_2 per physical unit of energy carrier e). Note that we distinguish between the plant's emissions intensity per energy source, $\epsilon_{i,a,e}$, and its product-level intensity, $\epsilon_{i,a,v}$ (measured

in tCO_2 per tonne of output), which accounts for the full metallurgical process including upstream iron production. The limited substitutability within these nests ($\sigma_a^{\text{KLE}}, \sigma_a^{\text{Ene}}$) formally represents the sector's hard-to-abate nature, as plants cannot easily switch away from carbon-intensive fuels in response to τ_r.

Transformation into Intermediate and Final Varieties The semi-finished product is transformed into variety-specific outputs $v \in \mathcal{V}$ according to a CET frontier. For ironmaking technologies $a \in \mathcal{A}_I$, this yields intermediate iron varieties ($v \in \mathcal{V}_I$), while steelmaking technologies $a \in \mathcal{A}_S$ produce finished steel varieties ($v \in \mathcal{V}_S$). The supply of a specific variety v from plant i is given by:

$$q_{i,a,v} = \theta_{i,a,v} Q_{i,a} \left(\frac{p_{i,a,v}}{c_{i,a}} \right)^{\tau_a} \tag{11}$$

where $q_{i,a,v}$ is the quantity of variety v produced, $Q_{i,a}$ is the total activity level of the plant, $p_{i,a,v}$ is the variety price, τ_a is the transformation elasticity, and $\theta_{i,a,v}$ is the benchmark volume share of variety v in technology a's output.

Vertical Supply Chains and Iron Sourcing A defining feature of the model is the vertical linkage between the iron and steel stages. Because carbon emissions are primarily concentrated in the ironmaking stage, steelmaking technologies ($a \in \mathcal{A}_S$) face a strategic supply chain choice: produce iron intermediate inputs $v \in \mathcal{V}_I$ internally (onsite integration) or purchase them from the regional merchant market. The effective user cost $p_{i,a,v}^{\text{Iron}}$ for a steelmaker is a CES composite:

$$p_{i,a,v}^{\text{Iron}} = \left[\theta_{i,a,v}^{\text{Iron}} (p_{i,a',v})^{1-\sigma_v^{\text{Iron}}} + (1 - \theta_{i,a,v}^{\text{Iron}}) (P_{r,v})^{1-\sigma_v^{\text{Iron}}} \right]^{\frac{1}{1-\sigma_v^{\text{Iron}}}} \tag{12}$$

where $p_{i,a',v}$ is the cost of iron produced onsite using technology $a' \in \mathcal{A}_I$, $P_{r,v}$ is the regional price of the Armington iron composite, $\theta_{i,a,v}^{\text{Iron}}$ is the plant-specific benchmark onsite share, and σ_v^{Iron} is the substitution elasticity between onsite and merchant iron. This structure implicitly captures the energetic and logistical advantages of onsite *hot metal* use versus merchant *cold* iron. We calibrate these parameters to reflect technology-specific rigidities; for example, BOF units are tightly integrated with onsite blast furnaces (low σ_v^{Iron}), while EAF units exhibit greater flexibility in sourcing various metallic charges from the market (higher σ_v^{Iron}). This formulation formally tracks the *carbon supply chain*, explicitly capturing how upstream carbon costs are transmitted to downstream steel varieties.

4.3 Trade and Macro-Production

The second component of the model defines the trade hierarchy and the macroeconomic bundling process that generates the derived demand for steel.

Trade in Steel and Raw Materials In each region r, the total supply of steel varieties and tradable raw materials (factors $f \in \mathcal{F}_{\text{tradable}}$) is formed by aggregating domestic and imported quantities through a nested CES (Armington) structure. For each commodity $j \in \mathcal{V} \cup \mathcal{F}_{\text{tradable}}$, the regional price index $P_{r,j}$ is:

$$P_{r,j} = \left[\theta^A_{r,j}(P^D_{r,j})^{1-\sigma^A_j} + (1-\theta^A_{r,j})(P^M_{r,j})^{1-\sigma^A_j}\right]^{\frac{1}{1-\sigma^A_j}}, \tag{13}$$

where $P^D_{r,j}$ and $P^M_{r,j}$ are the prices of the domestic and imported composites, respectively. The domestic composite price $P^D_{r,j}$ aggregates local supply from plants $i \in \mathcal{I}_r$ for steel, or local resource extraction for raw materials. The imported bundle price $P^M_{r,j}$ similarly aggregates bilateral imports across foreign origins $s \neq r$ using the substitution elasticity σ^M_j:

$$P^M_{r,j} = \left[\sum_{s \neq r} \theta^M_{s,r,j}(p_{s,r,j})^{1-\sigma^M_j}\right]^{\frac{1}{1-\sigma^M_j}}, \tag{14}$$

where $p_{s,r,j}$ is the price of commodity j exported from region s to region r, accounting for trade costs and applicable border taxes. This Armington structure ensures that the model captures the idiosyncratic trade patterns and cross-border price transmission for both industrial goods and upstream inputs.

Variety Demand The demand for a specific plant-level variety $q_{i,r,a,v}$ in destination region r is derived from the Armington composites via Shephard's Lemma. For a given total regional demand $Q_{r,v}$, the quantity sourced from plant i in region s is:

$$q_{i,r,a,v} = Q_{r,v}\frac{\partial P_{r,v}}{\partial p^x_{i,r,a,v}}. \tag{15}$$

Macro-Production of the Consumption Good The aggregate macroeconomic sector (*Mac*) produces the intermediate macro-good $Q_{r,Mac}$ by bundling primary factors, macro-intermediates, and Armington steel varieties $v \in \mathcal{V}_S$. The unit cost $P_{r,Mac}$ is the dual of the nested CES production function:

$$P_{r,Mac} = \mathcal{C}_{r,Mac}(\Omega_{r,f}, P_{r,v}, \dots), \tag{16}$$

and the resulting demand for each steel variety v is $Q_{r,Mac,v} = Q_{r,Mac}(\partial P_{r,Mac}/\partial P_{r,v})$. Finally, the consumption bundle $Q_{r,C}$ is produced using the macro-good as the primary input, with its price $P_{r,C}$ serving as the numeraire or being endogenously determined.

4.4 Climate Policy and Emissions

Climate policy is characterized by two distinct instruments: a regional carbon price on production inputs and border carbon adjustments on traded products.

Regional Carbon Pricing on Energy A regional carbon price τ_r is levied on the carbon content of energy carriers $e \in \mathcal{E}$ used in production. As established in (7), this enters the producer's cost directly through the energy price $p_{i,a,e}$. Total regional emissions from the steel sector $Q_{r,CO2}$ are:

$$Q_{r,CO2} = \sum_{i \in \mathcal{I}_r} \sum_{a \in \mathcal{A}} \sum_{e \in \mathcal{E}} \epsilon_{i,a,e} q_{i,a,e}, \tag{17}$$

where $\epsilon_{i,a,e}$ is the carbon intensity per unit of energy carrier e.

Border Carbon Adjustments on Embodied Carbon While regional pricing applies to inputs, BCAs are applied to the total carbon embodied in finished varieties v. The embodied emissions intensity $\epsilon_{i,a,v}$ accounts for direct process emissions and energy-related emissions across the metallurgical supply chain. Under a mass-based BCA, the destination region r levies a charge on imports from plant i in region s:

$$p^x_{i,r,a,v} = p_{i,a,v} + \tau_r \epsilon_{i,a,v}. \tag{18}$$

Under a rate-based BCA, the destination region sets an intensity target $\bar{\psi}_{r,v}$. Exporters must surrender permits covering their actual emissions $\epsilon_{i,a,v}$ while receiving credits based on the target $\bar{\psi}_{r,v}$. The export price is:

$$p^x_{i,r,a,v} = p_{i,a,v} + \nu_{r,v}(\epsilon_{i,a,v} - \bar{\psi}_{r,v}), \tag{19}$$

where $\nu_{r,v}$ is the shadow price of the intensity constraint. The value of $\nu_{r,v}$ is determined endogenously by the permit market-clearing condition.

Revenue Redistribution All policy-generated revenues, including carbon tax proceeds and border adjustment wedges, are redistributed to the regional representative household. Total regional transfers TR_r are given by:

$$TR_r = \tau_r Q_{r,CO2} + \sum_{v \in \mathcal{V}} \sum_{s \neq r} \sum_{i \in \mathcal{I}_s} \sum_{a \in \mathcal{A}} (p^x_{i,r,a,v} - p_{i,a,v}) q_{i,r,a,v}, \tag{20}$$

where the first term represents domestic carbon pricing revenue and the second term represents the net proceeds (or "feebate" balance) from BCAs on all imported varieties.

4.5 Market Clearing

The global equilibrium is closed by a set of market-clearing identities that ensure the balance of supply and demand across all regional and international markets.

Tradable Commodities For each steel variety and tradable raw material $j \in \mathcal{V} \cup \mathcal{F}_{\text{tradable}}$, the total quantity supplied by region r must equal the sum of domestic demand and aggregate export

demand:

$$Q_{r,j}^{\text{sup}} = q_{r,r,j}^{D} + \sum_{s \neq r} q_{r,s,j}^{M}, \tag{21}$$

where $Q_{r,j}^{\text{sup}}$ is the total regional production (summed across plants for steel, or local resource extraction for materials), $q_{r,r,j}^{D}$ is domestic demand, and $q_{r,s,j}^{M}$ denotes aggregate bilateral exports from region r to destination s. The global equilibrium requires that these bilateral exports strictly equal the sum of import demands from the destination region. For steel varieties v, this bilateral trade identity is:

$$q_{r,s,v}^{M} = \sum_{i \in \mathcal{I}_r} \sum_{a \in \mathcal{A}} q_{i,s,a,v}, \tag{22}$$

which formally maps the plant-to-region derived demand from (15) into the regional export flow.

Macroeconomic Good The regional market for the macroeconomic good Mac clears when the total output $Q_{r,Mac}$ equals the sum of intermediate demand from the steel sector and final consumption demand from the household:

$$Q_{r,Mac} = \sum_{i,a} q_{i,a,Mac} + Q_{r,Mac}^{C}, \tag{23}$$

where $Q_{r,Mac}^{C}$ is the quantity of the macro-good used to produce the final consumption bundle $Q_{r,C}$.

Primary Factors Regional factor markets clear when the price-elastic factor supply $Q_{r,f}$ (from (2)) equals the aggregate demand from all local production activities:

$$Q_{r,f} = \sum_{i,a} q_{i,a,f} + q_{r,Mac,f}, \tag{24}$$

where $q_{i,a,f}$ is the plant-level factor demand and $q_{r,Mac,f}$ is the demand from the macro-sector.

Intensity Permits In regions implementing rate-based BCAs, the permit market clears when the net embodied emissions of imports satisfy the intensity target. The permit price $\nu_{r,v}$ is non-negative and is strictly positive only if the aggregate intensity of imports from all regions $s \neq r$ binds at the target $\bar{\psi}_{r,v}$:

$$\sum_{s \neq r} \sum_{i \in \mathcal{I}_s} \sum_{a \in \mathcal{A}} (\epsilon_{i,a,v} - \bar{\psi}_{r,v}) q_{i,r,a,v} \leq 0. \tag{25}$$

4.6 Equilibrium

For a given set of climate policy instruments—characterized by regional carbon prices τ_r, intensity targets $\bar{\psi}_{r,v}$, and permit-clearing rules—a competitive equilibrium is defined as a vector of prices $\mathcal{P} = \{p_{i,a,v}, P_{r,j}, \Omega_{r,f}, \omega_{r,f}, P_{r,Mac}, P_{r,C}, \nu_{r,v}\}$, quantities $\mathcal{Q} = \{q_{i,r,a,v}, Q_{r,j}, Q_{r,f}, Q_{r,Mac}, Q_{r,C}\}$, and regional incomes $\mathcal{I} = \{I_r\}$ such that all agents optimize and markets clear. Specifically, optimization

requires that households maximize utility subject to their budget constraint (3), determining the demand for the final consumption bundle $Q_{r,C}$. Production facilities minimize costs at the plant level, with unit costs $c_{i,a}$ and input demands determined by the nested CES system (7)–(9), while the aggregate macroeconomic sector minimizes costs for the intermediate macro-good $Q_{r,Mac}$ according to (16).

Perfect competition ensures that a zero-profit condition holds, rendering all producer prices equal to their corresponding marginal costs, including variety prices $p_{i,a,v}$ and Armington price indices $P_{r,j}$. Regional incomes I_r account for factor returns, the benchmark trade imbalance $\bar{\Delta}_r$, and net policy transfers TR_r (20). Simultaneously, market-clearing identities are satisfied globally for all tradable commodities (21), with global bilateral trade quantities satisfying the identity (22), and regionally for the macroeconomic good (23) and primary factors (24). Finally, the permit price $\nu_{r,v}$ clears the intensity permit market according to (25). Welfare changes across policy scenarios are evaluated using the Hicksian equivalent variation (EV), which calculates the change in regional income at baseline prices required to achieve the same utility level as in the counterfactual equilibrium.

5 Data and Parameters

The quantification of the structural model involves mapping plant-level engineering data to the benchmark variables $\{Q_{i,a}, q_{i,a,f}, q_{i,s,a,v}, \epsilon_{i,a,v}\}$ that define the global equilibrium. We quantify the model using a harmonized dataset that integrates high-resolution micro-data for the industrial layer with global trade and input-output statistics for the macroeconomic environment. The baseline for the industrial data is 2019, while the macroeconomic environment is anchored to a 2017 baseline.

5.1 Data Sources

The empirical foundation of the model is built upon three pillars that bridge the gap between industrial engineering and national accounting.

Industrial Micro-Data Data for the industrial layer ($i \in \mathcal{I}, a \in \mathcal{A}$) is primarily sourced from the CRU Steel Cost Service, which provides operational profiles for approximately 350 steel mills across 50 countries. This includes production volumes ($Q_{i,a}$), technology portfolios, and granular input coefficients for energy and materials. The dataset covers 53% of world production, with even higher coverage of 67.8% outside China. Notably, the coverage of the BF-BOF segment outside China is exceptional, representing 96% of global production in this category. We validate the representativeness of this sample against the Global Steel Plant Tracker (GSPT), which covers 96.1% of global steelmaking capacity (Global Energy Monitor, 2023). As shown in Figure A1, the distribution of plant capacities in the CRU dataset closely mirrors the GSPT, although it shows a slight under-representation of smaller plants with capacities between 1–2 Mt per year. The detailed process modeling within the CRU data allows us to distinguish between the technologies used

for different steel products. Finally, we reconcile aggregate regional totals with the World Steel Association Statistical Yearbooks.

Trade and Margins Bilateral trade flows for finished steel and intermediate iron varieties ($v \in \mathcal{V}$) are sourced from UN Comtrade at the HS 6-digit level. Trade in primary raw materials (e.g., iron ore, scrap, coal, natural gas) is reconciled using the BP Statistical Review for gas, the U.S. Energy Information Administration (EIA) for metallurgical coal, and UN Comtrade for ore and scrap. To model the extensive and intensive margins of trade, we utilize firm-level trade statistics from the OECD, including Structural and Demographic Business Statistics (SDBS) for firm counts and Trade by Enterprise Characteristics (TEC) for exporting behavior. Transport frictions and CIF/FOB margins are derived using the OECD International Transport and Insurance Costs (ITIC) dataset.

Macroeconomic Framework The aggregate macroeconomic sector (Mac) and primary factor markets ($f \in \mathcal{F}$) are parameterized using the GTAP 11 Database (Aguiar et al., 2022). This provides the general equilibrium consistency required to capture the feedback between the steel industry and the broader economy. The model includes 15 regions ($\mathcal{R}$): the European Union and the European Free Trade Association (EUR), the United States (USA), China (CHN), India (IND), Japan (JPN), South Korea (KOR), Other Asia (OAS), Oceania (OCE), Canada (CAN), Mexico (MEX), Brazil (BRA), Other South America (OSA), Turkey and Other Europe (OEU), CIS and Ukraine (CIS), and Western Asia and Africa (WAF).

5.2 Quantitative Implementation

The implementation process transforms raw data into the structural parameters and benchmark states that define the global equilibrium. Technical details of the reconciliation and imputation procedures are provided in Appendix B.

Benchmark Construction We populate the plant-level benchmark quantities through a systematic reconciliation procedure. Plant-level output $Q_{i,a}$ is scaled by regional adjustment factors $\alpha_{r,s}$ to match worldsteel control totals, ensuring that the model exactly reproduces observed regional production scales. Input quantities $q_{i,a,f}$ are determined by mapping CRU's engineering coefficients to the physical requirements of metallurgy. Plant-level bilateral trade flows $q_{i,s,a,v}$ for steel are constructed via a Ricardian selection mechanism, where export status is determined by the unit-cost rankings ($c_{i,a}$) of facilities within each region. For tradable input factors $f \in \mathcal{F}_{\text{tradable}}$, bilateral flows $Q_{r,s,f}$ are reconciled with EIA and BP production totals through a quadratic programming optimization that ensures global material balance. Finally, emissions intensities $\epsilon_{i,a,v}$ are calculated bottom-up and corrected by a byproduct factor (ϕ^{bp}) to account for the physical energy flows of internally generated process gases.

Structural Parameterization The behavioral responses of the model are governed by a set of structural parameterizations summarized in Table A1. Because exact econometric estimates for highly granular intra-plant substitution are scarce, we combine standard macroeconomic values with calibrations grounded in metallurgical constraints. Macroeconomic trade parameters, such as the Armington elasticities (σ_j^A, σ_j^M) and the consumption substitution elasticity (σ^C), are drawn from the GTAP database. For the plant-level nests, elasticities ($\sigma_a, \sigma_a^{\text{KLE}}, \sigma_a^{\text{VA}}, \sigma_a^{\text{Ene}}$) are set low to reflect the stoichiometric and thermal rigidities of heavy industrial production. For instance, substitution in the material nest (σ_a^{Mat}) is tightly bounded in blast furnaces to reflect physical maximums on scrap-to-ore ratios, while operating with greater flexibility in Electric Arc Furnaces (EAF). Similarly, the vertical integration elasticity (σ_v^{Iron}) is parameterized to reflect the severe logistical and energetic frictions (e.g., thermal penalties) associated with unbundling integrated molten-iron sites.

Beyond the core production nests, the model calibrates several parameters governing supply-side and demand-side heterogeneity. The transformation elasticity (τ_a) governs the supply-side capability to shift from semi-finished steel into specific product varieties (flat vs. long). On the demand side, the within-domestic variety substitution (σ_v^D) captures how downstream users switch between domestic plants. Finally, factor market rigidities are captured by the price elasticity of factor supply ($\eta_{r,f}$), distinguishing between the relatively inelastic structural constraints of circular scrap recovery and the highly elastic supply of primary resources, as well as transformation elasticities restricting plant-level mobility of capital and labor (η_f^F). To ensure our core findings are not artifacts of these point estimates, we subject the most critical structural and behavioral parameters to an extensive sensitivity analysis in Section 7.5.

6 Design of Counterfactual Experiments

We evaluate the structural properties of border carbon adjustments through a series of counterfactual experiments. These experiments are designed to isolate the efficiency and distributional consequences of a shift from absolute carbon pricing to performance-based standards across two distinct regulatory environments.

6.1 Policy Paradigms

The analysis is structured around two paradigms: a *symmetric paradigm*, where the BCA complements an existing domestic carbon price, and an *asymmetric paradigm*, where the BCA operates as a standalone instrument in a region without a domestic price. Within these paradigms, we simulate eight specific policy configurations:

Reference Scenarios The *Global Cap* scenario serves as a first-best reference point where a unified global carbon price is applied across all regions to achieve a target emissions reduction. For the symmetric paradigm, the *EU Cap* represents the standard unilateral policy case (modeling the EU

Emissions Trading System), where domestic carbon pricing (τ_{EUR}) is implemented at stringency levels ranging from 0% to 50% domestic abatement without any border measures.

Symmetric Paradigm: EU Policy Designs We compare two primary border designs for the EU. In the *EU M-BCA* scenario, the region combines domestic pricing with a mass-based BCA, extending τ_{EUR} to the full carbon content of imports ($\epsilon_{i,a,v}$) as defined in (18). The *EU R-BCA* scenario replaces the mass-based tariff with a rate-based standard (tradable performance standard), where intensity targets $\bar{\psi}_{r,v}$ are set to the import intensities observed in the corresponding M-BCA case. This ensures that the permit price $\nu_{r,v}$ is calibrated to the same technological baseline. To isolate the efficiency loss inherent in the rate-based design, we also model an *EU DM-BCA* (Discounted Mass-based) scenario, where the absolute carbon price is scaled down to yield the same global emissions reduction as the R-BCA.

Asymmetric Paradigm: standalone BCAs In the asymmetric paradigm, we model standalone border measures for the United States, which currently lacks a federal carbon price. The *US M-BCA* implements a mass-based tariff on imports at a rate matching the EU ETS price, while the *US R-BCA* implements a standalone intensity standard. The target $\bar{\psi}_{USA,v}$ is set to the import intensity under the US M-BCA to maintain comparability. Finally, the *US DM-BCA* scales the standalone tariff to match the global abatement of the rate-based alternative.

We select a 30% reduction in EU domestic steel emissions as our primary point of comparison. This stringency level is aligned with the European steel industry's decarbonization pathway under the Fit for 55 package, which aims for a 55% reduction by 2030 relative to 1990 levels—approximately 30% below our 2019 baseline (European Steel Association, 2020).

6.2 Accounting Scope and Facility Boundaries

A critical detail in the design of any BCA is the scope of emissions covered ($\epsilon_{i,a,v}$). We adopt a comprehensive accounting boundary that captures direct process emissions, indirect electricity emissions (Scope 2), and emissions from upstream intermediate production (e.g., ironmaking). For electricity, we utilize exogenous regional grid intensities to determine the carbon cost component of energy carriers $p_{i,a,e}$. While current implementations such as the EU CBAM initially limit Scope 2 coverage, our comprehensive scope evaluates the long-run structural properties of these instruments, consistent with the goals of the US Foreign Pollution Fee Act.

The effectiveness of this coverage depends on the facility boundary. In both rate-based and mass-based designs, the tax is assessed using the exporting plant's declared intensity. By explicitly modeling trade in both finished steel ($v \in \mathcal{V}_S$) and intermediate iron ($v \in \mathcal{V}_I$), we evaluate whether firms exploit regulatory boundaries to reclassify emissions. For instance, a non-integrated producer that purchases merchant iron may declare a lower facility-level intensity if upstream emissions are not perfectly tracked, a phenomenon we formalize as *vertical leakage*.

6.3 The Discounted Price Approach

A challenge in comparing mass-based and rate-based BCAs is that they inherently yield different global abatement levels at the same domestic stringency. Rate-based designs typically generate lower global reductions because they lack the scale-effect pricing on the total volume of imports. To ensure a fair comparison of cost-effectiveness, we introduce a *Discounted Price Approach*. We solve for a scaling parameter $\lambda \in [0, 1]$ such that a mass-based BCA with price $\lambda\tau_r$ produces identical total global emissions as the corresponding rate-based design. This ensures that the environmental outcomes are held constant, allowing us to compare the regional welfare burden (EV_r) and industrial protection outcomes on a level playing field. This approach provides a rigorous quantitative measure of the "efficiency premium" or "discount" inherent in performance-based policy designs.

7 Quantitative Results

In this section, we evaluate how the choice of BCA design mediates the tension between climate objectives and the economic costs of adjustment. We structure our analysis around three key dimensions: the efficiency of the policy paradigm, the structural mechanisms of abatement, and the distributional consequences for industries and nations.

To preview the mechanisms, we show that rate-based designs tend to shift adjustment toward composition and reallocation rather than scale. In our steel setting, this manifests in two salient ways. First, when benchmarks are defined in technologically homogeneous upstream stages (such as ironmaking), the intensity-gap design can generate an implicit subsidy that keeps intermediate imports near benchmark levels and facilitates unbundling of the production chain (vertical leakage). Second, we document reshuffling: some exporters can lower measured embedded emissions in the regulated market by shipping their cleanest output while diverting dirtier production elsewhere. We develop these mechanisms in Section 7.2, quantify their distributional implications in Section 7.3, and synthesize the resulting design considerations for policy in Section 7.4.

7.1 The Cost of Ambition: National vs. Global Perspectives

The design of a border carbon adjustment shapes how much of the carbon cost is passed through to producers and consumers, and how much is absorbed through implicit subsidies to foreign competitors. We evaluate the choice between mass-based and rate-based designs across two policy paradigms: a symmetric paradigm (domestic carbon pricing combined with a BCA, modeled on the EU) and an asymmetric paradigm (a standalone BCA without domestic pricing, modeled on the US). The mass-based design generates more revenue and abatement, while the rate-based design reduces downstream price shocks but is associated with higher national welfare losses and greater leakage.

We begin by evaluating these trade-offs from a national perspective. Figure 5 (Panel a) traces

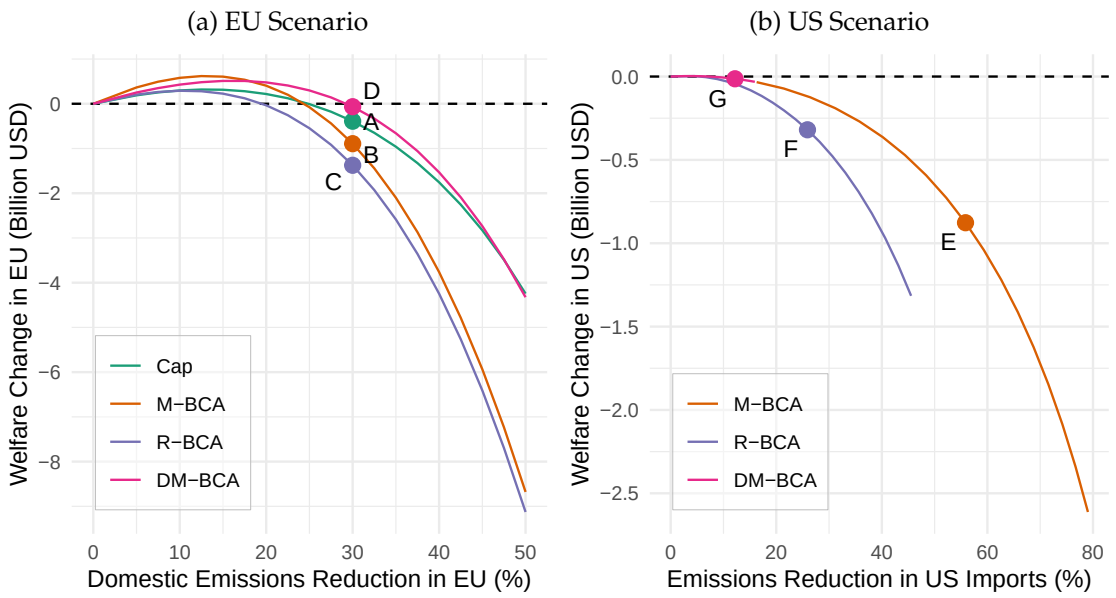


Figure 5: National Welfare Changes

Note: Panel (a) plots EU welfare change (billion USD) against domestic emissions reduction in the EU (%). Panel (b) plots US welfare change (billion USD) against emissions reduction in US imports (%). Lines trace policy designs across stringency, and labeled points (A–G) mark the comparison points used in the analysis.

the welfare implications of EU policy choices. We evaluate welfare as the regional equivalent variation, representing the net economic burden of the policy relative to the baseline. At lower stringency levels, EU policies yield positive welfare gains as terms-of-trade benefits outweigh domestic deadweight losses. As stringency increases, these gains erode and become negative. Under the rate-based design (Point C), welfare losses are larger: \$1.38 billion compared to \$0.89 billion for the mass-based alternative. Point D (discounted mass-based BCA) achieves the same global emissions reduction as the rate-based design at a lower cost ($-$\$0.07 billion). Solving for environmental equivalence yields an implicit discount factor of $\lambda \approx 0.36$. As established in Section 3, this indicates that the rate-based design is economically equivalent to a tariff discounted by 64%, reflecting the signal dilution inherent in intensity-based benchmarks.

In the US context (Panel b), where there is no domestic carbon price to anchor the policy, the trade-off shifts. As shown in Table 1, the mass-based BCA (Point E) provides stronger industrial protection (output increases by $+1.11\%$) but imposes a higher welfare cost ($-$\$0.88 billion) due to consumer price impacts. The rate-based design (Point F) yields a smaller national loss ($-$\$0.32 billion) but reduced environmental effectiveness. As shown by Point G, the implicit discount factor drops to $\lambda \approx 0.12$. This indicates that the asymmetric paradigm implies an effective carbon price nearly 90% lower than the reference price, limiting incentives for technology switching or fuel substitution in the steel sector.

Comparing these strategies on the global stage reveals the efficiency gap. Figure 6 illustrates

Table 1: Summary of Policy Impacts at Comparable Stringency Levels

Scenario	Global Emissions (%)	(Reverse) Leakage (%)	Domestic Welfare ($bn)	Domestic Output (%) Flat	Long	Macro	Carbon Price ($/tCO$_2$)
Panel (a): EU Scenarios							
A. Cap	−0.69	56.44	−0.39	−17.65	−6.68	−0.062	69.71
B. M-BCA	−1.34	16.05	−0.89	−16.25	−4.60	−0.126	87.90
C. R-BCA	−1.02	36.18	−1.38	−17.02	−6.19	−0.081	74.94
D. DM-BCA	−1.02	36.18	−0.07	−17.12	−5.79	−0.087	76.91
Panel (b): US Scenarios							
E. M-BCA	−0.26	(23.05)	−0.88	1.21	1.11	−0.018	87.90
F. R-BCA	−0.06	(14.93)	−0.32	0.40	0.86	−0.004	
G. DM-BCA	−0.06	(15.72)	−0.01	0.11	0.19	−0.003	10.19

Note: Scenarios correspond to the points in Figure 5. Panel (a) presents EU policy options, all of which include the domestic ETS cap. Panel (b) presents US policy options, which are standalone border measures without a domestic price anchor. Scenarios are compared at a stringency level corresponding to a carbon price of approximately \$88/tCO$_2$ (equivalent to a 30% domestic emissions reduction for the EU). Leakage rates in Panel (b) represent reverse leakage (onshoring).

the global cost-effectiveness frontier. The EU mass-based BCA (EU M-BCA) achieves the highest global emissions reduction (1.34%) for a given welfare cost. In contrast, the US scenarios (Points E, F, G) exhibit a nearly vertical cost curve. Even with a mass-based tariff (Point E), the US policy achieves limited global abatement (0.26%). Shifting to a rate-based design (Point F) yields similarly small results (0.06%). While the magnitude is constrained by the smaller volume of US imports, the vertical cost curve indicates that without a domestic price anchor, border measures alone provide limited leverage for global decarbonization.

These patterns empirically validate the global efficiency results established in Section 3. As predicted by Proposition 1, the rate-based design's implicit subsidy neutralizes the consumption margin as a channel for abatement. This forced reliance on the technology margin explains the efficiency gap observed in Figure 6, where the EU R-BCA (Point C) incurs higher costs for lower abatement compared to the mass-based alternative (Point D). Similarly, the vertical cost curve in the asymmetric paradigm reflects the high marginal costs of forcing abatement exclusively through technology switching when the scale of imports is preserved.

7.2 Abatement Margins: Scale, Composition, and Technique

Having established the comparative welfare impacts of rate-based designs, we now examine the structural mechanisms driving these outcomes. We decompose the global emissions response to identify leakage channels and isolate three distinct margins of abatement: the scale effect (changes in total production), the composition effect (shifts in technology mix), and the technique effect

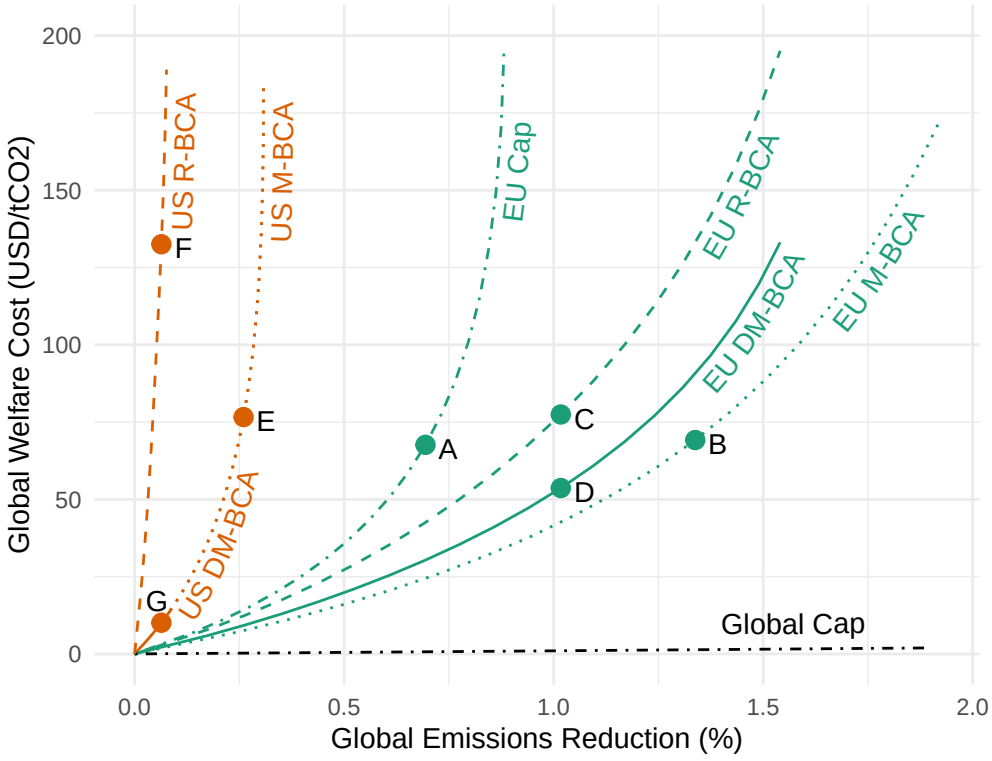


Figure 6: Global Cost-Effectiveness Frontier

Note: Global emissions reduction (%) is plotted against global welfare cost (USD/tCO_2). Points trace policy designs across stringency in the EU and US scenarios.

(within-technology efficiency improvements).

As discussed in Section 3, we apply these concepts at the intra-sectoral level. While standard trade theory typically defines the composition effect as a shift between sectors (e.g., manufacturing vs. services), we use it to capture the transition between production routes (e.g., from integrated BF-BOF to cleaner EAF). The technique effect captures efficiency gains within specific plants (e.g., through fuel switching or energy efficiency). While the stylized theoretical framework in Section 3 focused on the strategic trade-offs between scale and composition, the full quantitative model incorporates the technique margin as an additional channel for abatement.

Table 2 (Panel a) shows that the rate-based design results in a larger increase in emissions in the Rest of the World (22.50 Mt CO_2) than the mass-based alternative (9.98 Mt CO_2). A key driver is the shift in production processes across borders. Under the EU R-BCA scenario, emissions from merchant ironmaking (Column 4) increase by 5.36 Mt, whereas they remain nearly flat under the mass-based scenario (-0.04 Mt).

This pattern is driven by the technological homogeneity of ironmaking, which triggers the knife-edge sensitivity formalized in Section 3. As shown in Table A2, the standard deviation of emissions intensity for pig iron imports is low (0.31) compared to long steel (1.01), reflecting

Table 2: Mechanism Analysis of Emissions Changes in ROW

Panel (a): Sources of Emissions Change

Scenario	Steelmaking (i)	Ironmaking (ii) On-site	Ironmaking (iii) Merchant	Total
Cap	1.46	30.73	2.91	35.11
M-BCA	1.29	8.74	-0.04	9.98
R-BCA	2.57	14.57	5.36	22.50
DM-BCA	1.24	20.04	1.22	22.50

Panel (b): Decomposition of Steel and On-site Ironmaking Emissions (i + ii)

Scenario	Scale	Technique	Composition	Total
Cap	26.98	2.17	3.05	32.20
M-BCA	8.80	3.29	-2.06	10.03
R-BCA	17.24	2.88	-2.97	17.14
DM-BCA	17.29	2.66	1.34	21.28

Note: Values represent CO_2 emissions changes in the Rest of the World (ROW) relative to the baseline in million tonnes (Mt). All scenarios are EU policy options, with the domestic ETS cap in place and at a stringency level corresponding to a 30% emissions reduction target. Panel (a) distinguishes between direct steelmaking (i), on-site ironmaking (ii), and merchant ironmaking (iii). Panel (b) decomposes the combined integrated emissions (i + ii) into scale, technique, and composition effects.

chemical constraints of the BF that limit abatement even at high carbon prices. In our rate-based scenario, benchmarks are set to the emissions intensities of imports observed under the symmetric mass-based design. Because the intensity of pig iron remains rigid under these constraints, the calibrated benchmark (ψ) converges toward the actual intensity of all producers. As predicted by Lemma 3, this homogeneity renders the credit price highly sensitive to the benchmark level. In our simulation, the calibration lands on the zero-signal side of the knife-edge, effectively disabling the carbon price at the border. Table A2 confirms this asymmetry: the import price of pig iron rises to \$540/t under the mass-based design but remains at \$320/t under the rate-based design. This price wedge induces the vertical leakage predicted by Proposition 7, allowing import volumes to rebound from 1.73 Mt (M-BCA) to 3.50 Mt (R-BCA) as downstream steelmakers substitute domestic integrated iron with under-taxed imports.

In interpreting these findings, we acknowledge that the substitution of on-site hot metal with merchant pig iron involves significant engineering and logistical rigidities. Most integrated BOF plants are optimized for molten iron charging; utilizing cold merchant iron would incur a substantial thermal penalty and require specialized handling infrastructure. Furthermore, the global merchant pig iron market is currently small relative to total iron production. However, these physical constraints do not negate the underlying structural vulnerability created by the intensity standard. Rather, they suggest that vertical leakage is most likely to manifest through the long-term

reconfiguration of the production fleet—specifically by incentivizing the growth of EAF capacity that utilizes under-taxed imported pig iron as a high-purity feedstock. Our results thus illustrate the structural vulnerability created when performance standards are applied to rigid upstream stages. In this context, the benchmark required to match environmental outcomes inadvertently creates an implicit subsidy that permits the offshoring of carbon-intensive intermediates.

To understand the drivers of the remaining steel emissions, we decompose them using the Logarithmic Mean Divisia Index (LMDI) method. Following the identity $E = \sum_a (Y \cdot s_a \cdot e_a)$, where Y is total production, $s_a = Y_a/Y$ is the production share of technology a, and $e_a = E_a/Y_a$ is emissions intensity, we isolate three effects:

$$\Delta E = \underbrace{\sum_a w_a \ln\left(\frac{Y^1}{Y^0}\right)}_{\text{Scale Effect}} + \underbrace{\sum_a w_a \ln\left(\frac{s_a^1}{s_a^0}\right)}_{\text{Composition Effect}} + \underbrace{\sum_a w_a \ln\left(\frac{e_a^1}{e_a^0}\right)}_{\text{Technique Effect}} \tag{26}$$

where $w_a = \frac{E_a^1 - E_a^0}{\ln E_a^1 - \ln E_a^0}$ is the logarithmic mean weight, and superscripts 0 and 1 denote the baseline and policy scenario, respectively.

Panel B presents this decomposition. The rate-based design (Row 3) induces a stronger negative composition effect (-2.97 Mt) than the mass-based design (-2.06 Mt), indicating a more pronounced shift toward cleaner production methods like EAF for exports. However, this gain is offset by the scale effect. Because rate-based measures lower the effective border price relative to full tariffs, they do not suppress global demand as much. The scale effect under EU R-BCA contributes to an emissions increase of 17.24 Mt, nearly double the 8.80 Mt observed under the mass-based design. This decomposition confirms the composition–scale substitution established in Proposition 2: by neutralizing the scale margin as an abatement channel, the rate-based design forces the foreign sector to over-rely on the composition margin to satisfy the global emissions budget.

The technique effect, which reflects within-plant improvements such as fuel switching or energy efficiency, contributes an additional 1.2 Mt of abatement in both scenarios. While this margin is quantitatively smaller than the structural shifts in technology mix, its stability across designs suggests that border measures primarily influence the high-level technological configuration of the industry rather than the incremental efficiency of existing units.

In the US context, the absence of a domestic cap yields a mirror image. Figure 7 contrasts standard leakage in the EU (Panel a) with reverse leakage in the US (Panel b). Consistent with Brunel and Levinson (2024), the US tariff-only approach induces onshoring of pollution. By protecting domestic producers without a corresponding domestic carbon cost, the policy increases domestic production and partially offsets the reduction in imports. This empirically confirms the substitution effect predicted by Proposition 6: without a domestic price anchor, the border measure acts as a standard industrial tariff that reshapes production toward unabated domestic varieties, thereby limiting its effectiveness as a global climate instrument.

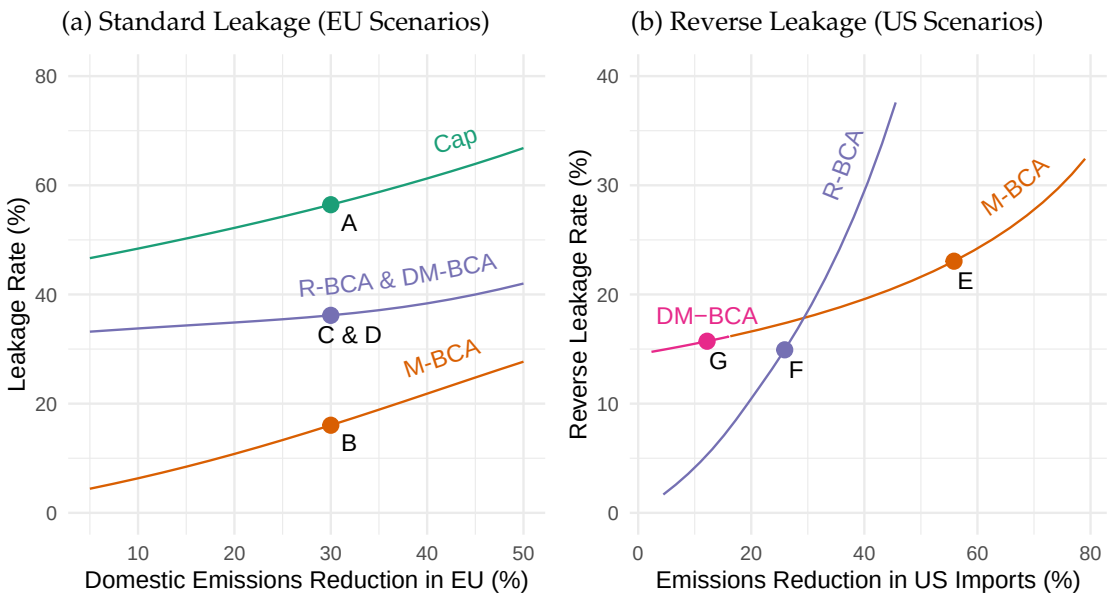


Figure 7: Leakage Dynamics

Note: Panel (a) plots the leakage rate (%) against domestic emissions reduction in the EU (%) for EU policy scenarios. Panel (b) plots the reverse leakage rate (%) against emissions reduction in US imports (%) for US policy scenarios. Lines trace policy designs across stringency, and markers indicate the reference stringency used for comparisons.

7.3 Distributional Impacts: Winners, Losers, and the Instrument Choice

The structural trade-offs identified above manifest as distributional differences. Because rate-based designs do not price the total carbon content of imports, they yield larger national welfare losses ($-$\$1.38 billion vs. $-$\$0.89 billion). This welfare gap is primarily driven by the foregone fiscal revenue identified in Proposition 4: by subsidizing imports via the benchmark, the EU forfeits the opportunity to capture terms-of-trade benefits from foreign producers.

Figure 8 visualizes this regional burden shift. The EU's rate-based design increases its own welfare loss compared to the mass-based alternative, while reducing losses for the US and the CIS region. Japan and Korea gain in both scenarios, and Other Asia (OAS) gains more under the rate-based design (\$0.16 billion compared to \$0.02 billion under mass-based). This pattern is driven by the intensity standard: by penalizing the highest-intensity producers (China and India) most, the policy confers relative advantages to mid-intensity producers such as the CIS region and to relatively cleaner producers such as Korea and OAS, allowing them to capture market share.

Trade volume analysis confirms this mechanism (Figure A2). Under the rate-based design, CIS pig iron exports rebound to 2.31 Mt, more than double the suppressed level under the mass-based design (1.03 Mt). This indicates that the policy does not suppress the vertical leakage identified in Section 7.2.

The surge in exports from regions like Japan and Korea is driven by resource reshuffling, a dynamic formalized in Section 3. We define the Reshuffling Index (R) as the ratio of the emissions

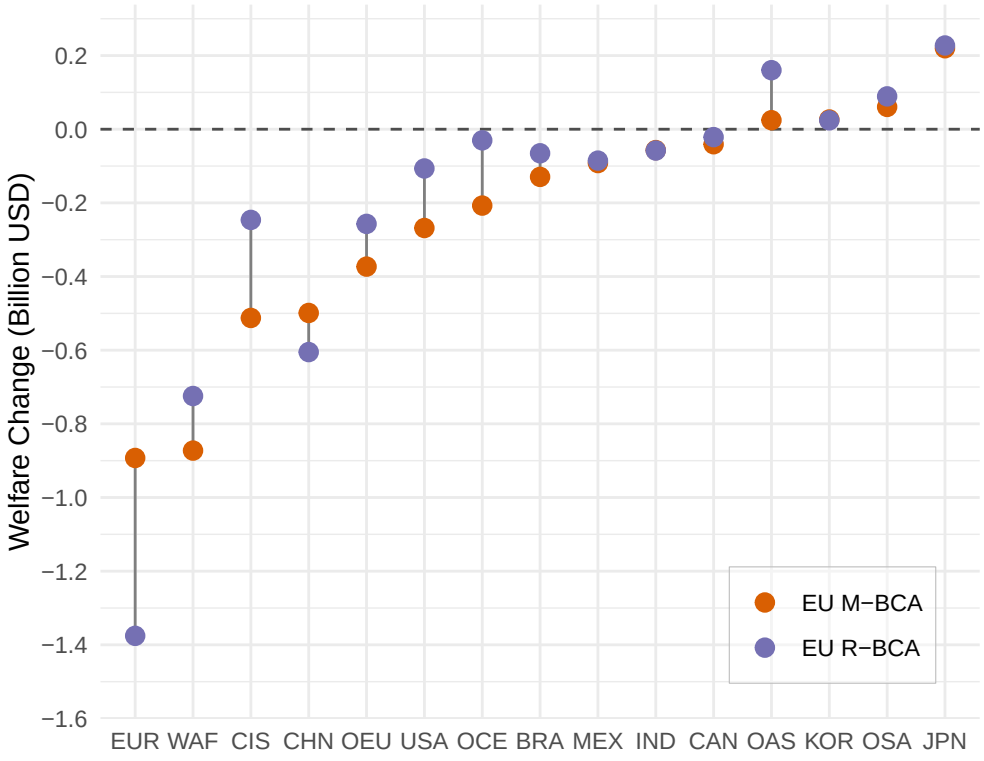


Figure 8: Regional Burden Sharing

Note: Bars report regional welfare changes (billion USD) under the EU policy scenarios. Region codes are defined in Section 5.1.

intensity of exports to the EU relative to the producer's domestic average intensity ($R = e_{exp}/e_{dom}$). As shown in Table A3, this strategic selection is restricted to sectors with high technological variance. The aggregate Reshuffling Index for the long steel sector falls from 1.08 to 0.55 under both border adjustment designs, while remaining nearly stagnant for the homogeneous pig iron sector (0.96). This disparity demonstrates that firms exploit technological heterogeneity to satisfy border standards without necessarily improving global production methods. At the regional level, this distortion is most pronounced for mixed-technology producers like Japan, where the index for long steel reaches 0.20, effectively diverting the clean capacity of the domestic market to satisfy the regulated market. This behavior empirically validates Proposition 3: at lower trade volumes, producers prioritize clean units to minimize taxes (Lemma 2), but as the rate-based design sustains higher volumes, this existing capacity is diluted, eventually forcing investment in new capacity to meet the standard.

Domestically, a central motivation for the rate-based design is downstream price protection. Table 3 shows that at 30% stringency, user prices for flat steel rise by 24.2% under the rate-based design versus 34.6% under the mass-based alternative. This price differential reflects the effective border tax rate of $\lambda \approx 0.36$ identified in Section 7.1, which acts as a 64% discount on the carbon cost.

Table 3: Import and User Price Changes in the EU

	Import Prices		User Prices		Consumer
Scenario	Flat	Long	Flat	Long	Price Index
Cap	0.69	-0.26	19.88	8.70	0.02
M-BCA	37.10	29.90	34.58	16.36	0.04
R-BCA	10.95	5.70	24.18	10.49	0.02
DM-BCA	13.80	11.47	25.61	11.85	0.03

Note: Values represent percentage changes relative to the benchmark. Import prices refer to the cost of imported steel products into the EU. User prices refer to the prices paid by end consumers within the EU. Consumer Price Index refers to the overall consumer price index within the EU.

Table 4: US Price Impacts of rate-based vs. Mass-based BCAs at Varying Stringency

	Flat		Long		Consumer	
Stringency	R-BCA	DM-BCA	R-BCA	DM-BCA	R-BCA	DM-BCA
Low	0.02	0.38	0.30	0.34	0.0001	-0.0021
Medium	0.65	1.05	1.60	0.94	0.0020	-0.0020
High	3.56	1.40	2.95	1.23	0.0081	-0.0019

Note: Values represent percentage changes relative to the benchmark. Comparison is between US rate-based BCA and an environmentally equivalent discounted mass-based BCA. The Low, Medium, and High stringency levels correspond to the carbon prices in the EU M-BCA scenario at 10%, 30%, and 50% domestic emissions reduction targets, respectively (approximately \$20, \$88, and \$241 per ton CO_2).

The lower input cost is reflected in macro outcomes: Table 1 shows a smaller decline in EU Macro Production under the rate-based scenario (-0.081%) than under the mass-based scenario (-0.126%). At the same time, domestic steel output declines more under the rate-based design; for long steel, output falls by 6.19% versus 4.60% under the mass-based design, as lower-cost imports exert greater competitive pressure on domestic integrated mills. These results highlight the fundamental trade-off between competing notions of economic protection: while rate-based designs insulate downstream consumers via lower prices, they do so at the cost of weaker protection for domestic upstream industrial output. This is the industrial incidence result predicted by Proposition 5.

Finally, Table 4 shows diminishing returns to downstream protection. In the US context, rate-based designs appear attractive at low stringency, minimizing welfare losses relative to a full tariff. However, this advantage erodes as climate ambition rises; the rate-based design is cheaper for users at low stringency (\$20/tCO_2), but becomes more expensive than the mass-based alternative at high stringency (\$241/tCO_2). This reversal occurs because the rate-based design lacks a revenue mechanism. A mass-based tariff generates fiscal revenue that can be recycled to offset economic distortions, whereas a rate-based standard does not. These differences underscore that the choice between mass-based and rate-based designs shapes both distributional outcomes and price impacts

as climate ambition increases.

7.4 Design Considerations for Policy

The structural trade-offs identified above provide several insights for border adjustment design. The vulnerability of rate-based BCAs to vertical leakage in technologically homogeneous sectors suggest that intensity standards may be less effective for upstream, homogeneous commodity markets where stronger price signals are needed to override arbitrage incentives. Conversely, the results indicate that while rate-based designs face constraints for rigid intermediates such as pig iron, they remain effective for heterogeneous finished goods.

These findings suggest that a product-differentiated approach could be a potential design strategy. Applying a mass-based BCA to homogeneous intermediates to mitigate vertical leakage, while retaining a rate-based design for finished steel to preserve downstream price protection, could balance competing objectives. Alternatively, a consumption-based charge (or climate contribution) could be paired with a rate-based BCA to restore the scale effect and generate revenue for recycling while preserving targeted technology incentives. Ultimately, the effectiveness of any BCA architecture depends on how well the instrument's design aligns with the technological and structural characteristics of the regulated sector.

7.5 Sensitivity and Robustness

To verify that our qualitative results are not artifacts of specific calibration choices, we conduct a comprehensive sensitivity analysis on seven critical structural and elasticity parameters in the model. Each parameter is tested at an inelastic (Low) and elastic (High) bound to examine the stability of the cost-effectiveness gap and the structural leakage channels.

The sensitivity analysis confirms the structural robustness of the policy efficiency gap. Table A4 demonstrates that the mass-based design consistently achieves greater global emissions reduction at a lower regional welfare cost than the rate-based alternative across all parameter variations. This persistent superiority is visually summarized in the global cost-effectiveness frontier (Figure A4), which shows the mass-based policy occupying a strictly dominant frontier relative to the intensity-gap design regardless of technological substitution limits or scrap availability.

Furthermore, the mechanisms driving these differential outcomes remain robust. Figure A3 highlights that vertical leakage—specifically the surge in merchant pig iron imports under the rate-based design—is uniquely and heavily sensitive to the iron integration elasticity (σ_v^{Iron}). This confirms that upstream unbundling is a structural consequence of the intensity-gap incentive, bounded primarily by the technological friction of decoupling integrated facilities. Similarly, Table A6 shows that while parameters like scrap supply elasticity ($\eta_{r,\text{Scrap}}$) can constrain the composition margin in the rest of the world, the fundamental reliance on the scale margin under the rate-based policy persists, driving up the global cost of mitigation.

8 Conclusion

This paper provides a plant-level comparative analysis of mass-based and rate-based border carbon adjustments, uncovering the structural trade-offs inherent in these two policy paradigms. Our results demonstrate that while both instruments can be calibrated to achieve equivalent global abatement, their economic and distributional consequences differ profoundly.

The mass-based design, by pricing the total carbon footprint, functions as an efficient fiscal instrument that generates significant revenue and provides robust support for domestic industrial capacity. However, it also generates substantial price shocks for downstream consumers. In contrast, the rate-based design effectively insulates downstream industries from these shocks by pricing only the emissions intensity gap relative to a benchmark. Yet, this insulation involves a significant dilution of the carbon price signal—quantified as an implicit 64% discount in the symmetric paradigm—and higher national welfare losses as the regulating region forfeits fiscal revenue to foreign competitors.

Our analysis also highlights a structural sensitivity in performance-based standards within technologically homogeneous sectors. In the case of ironmaking, where chemical constraints limit the variance in emissions intensity, the benchmark required to match environmental outcomes can inadvertently create an implicit subsidy that permits the offshoring of carbon-intensive intermediates. This suggests that the effectiveness of intensity standards is highly product-dependent and may be undermined in rigid upstream segments of the value chain. Furthermore, we demonstrate that discriminatory border measures trigger aggressive resource reshuffling across both paradigms, as firms satisfy regulatory requirements by selectively diverting clean capacity to the regulated market rather than achieving global abatement. Finally, we find that the relative price advantage of rate-based measures faces diminishing returns as climate ambition scales, potentially reversing as the lack of a revenue-recycling mechanism imposes a heavier uncompensated burden on the economy.

Our analysis has several limitations that provide avenues for future work. First, we use a static equilibrium model, which may understate long-run technological learning and capital adjustment that rate-based designs are intended to induce. Over a multi-decade horizon, the technique effect from intensity standards could strengthen, potentially narrowing the stringency gap identified here. Second, while we focus on steel, interactions with other emissions-intensive trade-exposed sectors such as cement and aluminum warrant analysis. Finally, although our modeling of export source shifting captures trade substitution frictions via aggregate substitution elasticities, incorporating explicit logistical constraints and contractual rigidities could further moderate the extent of resource reshuffling. Furthermore, our framework abstracts from retaliatory trade measures and strategic policy interactions between regions, treating regulatory ambition as exogenous. Future research should examine coordinated BCA implementations and potential climate-club arrangements.

Ultimately, the choice between these instruments is not merely a technical decision but a fundamental strategic choice about which economic stakeholders to prioritize. Governments must weigh the benefits of downstream price mitigation (associated with rate-based designs) against

the objectives of upstream industrial defense and global climate effectiveness (both of which are more robustly realized through mass-based designs) as they navigate the transition to a low-carbon economy.

References

Aguiar, Angel, Maksym Chepeliev, Erwin Corong, and Dominique Van Der Mensbrugghe, "The Global Trade Analysis Project (GTAP) Data Base: Version 11," *Journal of Global Economic Analysis*, December 2022, *7* (2), 1–37.

Akerman, Anders, Rikard Forslid, and Ossian Prane, "Imports and the CO2 Emissions of Firms," *Journal of International Economics*, November 2024, *152*, 104004.

Antweiler, Werner, Brian R. Copeland, and M. Scott Taylor, "Is Free Trade Good for the Environment?," *American Economic Review*, 2001, *91* (4), 877–908.

Beaufils, Timothé, Joschka Wanner, and Leonie Wenz, "The Potential of Carbon Border Adjustments to Foster Climate Cooperation," 2024.

Ben-David, Itzhak, Yeejin Jang, Stefanie Kleimeier, and Michael Viehs, "Exporting Pollution: Where Do Multinational Firms Emit CO2?," *Economic Policy*, July 2021, *36* (107), 377–437.

Böhringer, Christoph, Edward J. Balistreri, and Thomas F. Rutherford, "The Role of Border Carbon Adjustment in Unilateral Climate Policy: Overview of an Energy Modeling Forum Study (EMF 29)," *Energy Economics*, December 2012, *34*, S97–S110.

_ **, Knut Einar Rosendahl, and Halvor Briseid Storrøsten**, "Robust Policies to Mitigate Carbon Leakage," *Journal of Public Economics*, May 2017, *149*, 35–46.

Brunel, Claire and Arik Levinson, "Carbon Tariffs 101," Technical Report w33024, National Bureau of Economic Research October 2024.

Campolmi, Alessia, Harald Fadinger, Chiara Forlati, Sabine Stillger, and Ulrich J. Wagner, "Designing Effective Carbon Border Adjustment with Minimal Information Requirements. Theory and Evidence," SSRN Scholarly Paper 4644941, Social Science Research Network July 2025.

Clausing, Kimberly A. and Catherine Wolfram, "Carbon Border Adjustments, Climate Clubs, and Subsidy Races When Climate Policies Vary," *Journal of Economic Perspectives*, August 2023, *37* (3), 137–162.

Clausing, Kimberly, Jonathan Colmer, Allan Hsiao, and Catherine Wolfram, "The Global Effects of Carbon Border Adjustment Mechanisms," Technical Report w33723, National Bureau of Economic Research April 2025.

Cohen, Francois and Giulia Valacchi, "The Heterogeneous Impact of Coal Prices on the Location of Cleaner and Dirtier Steel Plants," *The Energy Journal*, March 2022, *43* (2), 67–90.

Collard-Wexler, Allan and Jan De Loecker, "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*, January 2015, *105* (1), 131–171.

Copeland, Brian R, "Pollution Content Tariffs, Environmental Rent Shifting, and the Control of Cross-Border Pollution," *Journal of International Economics*, May 1996, *40* (3-4), 459–476.

Copeland, Brian R. and M. Scott Taylor, "Trade, Growth, and the Environment," *Journal of Economic Literature*, 2004, *42* (1), 7–71.

Coster, Pierre, Isabelle Mejean, and Julian di Giovanni, "Firms' Supply Chain Adaptation to Carbon Taxes," Technical Report 2024.

Dechezleprêtre, Antoine, Caterina Gennaioli, Ralf Martin, Mirabelle Muûls, and Thomas Stoerk, "Searching for Carbon Leaks in Multinational Companies," *Journal of Environmental Economics and Management*, March 2022, *112*, 102601.

Duscha, Vicki, Everett B. Peterson, Joachim Schleich, and Katja Schumacher, "Sectoral Targets to Address Competitiveness – a CGE Analysis with Focus on the Global Steel Sector," *Climate Change Economics*, February 2019, *10* (01), 1950001.

European Steel Association, "Successful Implementation of Bold New 2030 Climate Target Urgently Needs Tangible Framework," Press release September 2020.

Farrokhi, Farid and Ahmad Lashkaripour, "Can Trade Policy Mitigate Climate Change?," *Econometrica*, 2025, *93* (5), 1561–1599.

Felbermayr, Gabriel, Sonja Peterson, and Joschka Wanner, "Trade and the Environment, Trade Policies and Environmental Policies—How Do They Interact?," *Journal of Economic Surveys*, July 2025, *39* (3), 1148–1184.

Fischer, Carolyn, "Environmental Protection for Sale: Strategic Green Industrial Policy and Climate Finance," *Environmental and Resource Economics*, March 2017, *66* (3), 553–575.

_ **and Alan K. Fox**, "Comparing Policies to Combat Emissions Leakage: Border Carbon Adjustments versus Rebates," *Journal of Environmental Economics and Management*, September 2012, *64* (2), 199–216.

Fontagné, Lionel and Katheline Schubert, "The Economics of Border Carbon Adjustment: Rationale and Impacts of Compensating for Carbon at the Border," *Annual Review of Economics*, September 2023, *15* (1), 389–424.

Fowlie, Meredith L. and Mar Reguant, "Mitigating Emissions Leakage in Incomplete Carbon Markets," *Journal of the Association of Environmental and Resource Economists*, March 2022, *9* (2), 307–343.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan, "Market-Based Emissions Regulation and Industry Dynamics," *Journal of Political Economy*, February 2016, *124* (1), 249–302.

Global Energy Monitor, "Global Steel Plant Tracker," 2023.

Grossman, Gene M. and Elhanan Helpman, "Protection for Sale," *The American Economic Review*, 1994, *84* (4), 833–850.

Itskhoki, Oleg and Dmitry Mukhin, "The Optimal Macro Tariff," Technical Report w33839, National Bureau of Economic Research May 2025.

Keen, Michael and Christos Kotsogiannis, "Coordinating Climate and Trade Policies: Pareto Efficiency and the Role of Border Tax Adjustments," *Journal of International Economics*, September 2014, *94* (1), 119–128.

_ and _ , "Principles for Pareto Efficient Border Carbon Adjustment," Technical Report 2025.

Kortum, Samuel and David Weisbach, "Regional Carbon Policy in a Global Economy," Technical Report January 2026.

Kotchen, Matthew and Giovanni Maggi, "Carbon Taxes and Green Subsidies in a World Economy," Technical Report w34080, National Bureau of Economic Research July 2025.

Mathiesen, Lars and Ottar Mæstad, "Climate Policy and the Steel Industry: Achieving Global Emission Reductions by an Incomplete Climate Agreement," *The Energy Journal*, October 2004, *25* (4), 91–114.

Nordhaus, William, "Climate Clubs: Overcoming Free-Riding in International Climate Policy," *American Economic Review*, 2015, *105* (4), 1339–70.

Shapiro, Joseph S, "The Environmental Bias of Trade Policy*," *The Quarterly Journal of Economics*, May 2021, *136* (2), 831–886.

Shapiro, Joseph S. and Reed Walker, "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade," *American Economic Review*, 2018, *108* (12), 3814–54.

Sogalla, Robin, "Unilateral Carbon Pricing and Heterogeneous Firms," Technical Report 2025.

_ , Joschka Wanner, and Yuta Watabe, "New Trade Models, Same Old Emissions?," DIW Discussion Papers 2077, Deutsches Institut für Wirtschaftsforschung (DIW) 2024.

Stillger, Sabine, "The Role of Firm Heterogeneity and Intermediate Inputs in Carbon Leakage," Technical Report May 2025.

Weisbach, David A., Samuel Kortum, Michael Wang, and Yujia Yao, "Trade, Leakage, and the Design of a Carbon Tax," *Environmental and Energy Policy and the Economy*, January 2023, *4*, 43–90.

World Steel Association, "World Steel in Figures 2023," 2023.

A Formal Proofs

This appendix provides the formal mathematical derivations for the lemmas and propositions presented in the Conceptual Framework (Section 3).

A.1 Proof of Lemma 1 (Price Ranking)

Proof. The equilibrium price is the sum of production costs and the border charge: $p_{i,h} = c(s_{c,i}) + t(\epsilon_{i,h})$. Under a mass-based design, $t^M(\epsilon) = \tau\epsilon_{i,h}$. Under a rate-based design, $t^R(\epsilon) = \tau(\epsilon_{i,h} - \bar{\psi})$. We can express the rate-based charge as a fraction of the mass-based charge:

$$t^R = \tau\epsilon_{i,h}\left(1 - \frac{\bar{\psi}}{\epsilon_{i,h}}\right) = \lambda\tau\epsilon_{i,h} \tag{27}$$

where $\lambda = (1 - \bar{\psi}/\epsilon_{i,h})$. Since $\bar{\psi} > 0$ and $\epsilon_{i,h} > 0$ for carbon-intensive varieties, then $\lambda < 1$. The price ranking follows directly:

$$p^M_{i,h} - p^R_{i,h} = \tau\epsilon_{i,h} - \lambda\tau\epsilon_{i,h} = \tau\bar{\psi} > 0 \tag{28}$$

Thus, $p^M_{i,h} > p^R_{i,h}$. In the symmetric paradigm, λ represents the dilution of the domestic price anchor. In the asymmetric paradigm, λ represents the effective stringency relative to a mass-based reference price. □

A.2 Proof of Proposition 1 (Global Cost-Effectiveness)

Proof. A cost-effective achievement of the global target $\bar{E}$ requires minimizing the sum of production costs and lost consumer value for a given emissions constraint. Let Q^M and Q^R denote the equilibrium demand for the composite intermediate Q under the mass-based and rate-based designs, respectively. In a competitive equilibrium, the mass-based design internalizes the carbon constraint on both the technology adoption margin ($s_{c,i}$) and the consumption margin (Q), thereby achieving the cost-effective benchmark $Q^M = Q^*$.

The cost-effectiveness conditions require that the marginal cost of technology adoption equals the marginal cost of reducing consumption:

$$\frac{c'(s_{c,i})}{-\epsilon'(s_{c,i})} = \frac{G'(Q)}{\epsilon(s_{c,i})} = \tau \tag{29}$$

Under a rate-based BCA, the implicit subsidy $\sigma = \tau\bar{\psi}$ reduces the price index $P(p_{h,h}, p_{f,h})$ relative to the mass-based baseline. Consequently, downstream firms choose a larger volume $Q^R > Q^*$, neutralizing the consumption margin as a channel for abatement.

To satisfy the fixed global budget $\bar{E}$ at this larger scale, the sectoral intensity $\epsilon(s_c)$ must fall further than under the cost-effective benchmark, requiring a higher marginal carbon signal $\tau^R > \tau^M$. Since $c'' > 0$, the marginal cost of this additional technology switching is strictly higher than the

marginal cost of the avoided consumption reduction. This results in a deadweight loss, as the total economic cost of achieving the target $\bar{E}$ is higher under the rate-based architecture. □

A.3 Proof of Proposition 2 (Foreign Composition–Scale Substitution)

Proof. By construction, we compare two policy designs that achieve the same global emissions target $\bar{E}_{global} = E_h + E_f$. Given a fixed domestic emissions target $\bar{E}_h$ is achieved in both scenarios, the foreign sector must satisfy a fixed residual emissions budget:

$$E_f = \epsilon_{f,h} q_{f,h} + \epsilon_{f,f} q_{f,f} = \bar{E}_f \tag{30}$$

where $\bar{E}_f = \epsilon_f q_f$ is the sectoral emissions budget and $q_f = q_{f,h} + q_{f,f}$ is the total production volume.

From Lemma 1, the rate-based design results in a lower price index for the regulated variety ($p^R_{f,h} < p^M_{f,h}$) for any level of technological composition. For designs of equivalent stringency, the rate-based instrument sustains a larger trade volume $q^R_{f,h} > q^M_{f,h}$. Focusing on the benchmark case of inelastic foreign domestic demand established in the Setup, this implies a larger total production scale: $q^R_f > q^M_f$. This result is robust to elastic foreign domestic demand as long as the expansion of trade volume—enabled by the intensity benchmark—outweighs the contraction of the foreign domestic market caused by higher regional production costs.

To satisfy the fixed emissions budget $\bar{E}_f$, the higher production scale q^R_f must be offset by a lower average sectoral intensity ϵ^R_f:

$$\epsilon^R_f = \frac{\bar{E}_f}{q^R_f} < \frac{\bar{E}_f}{q^M_f} = \epsilon^M_f \tag{31}$$

Given the intensity definition $\epsilon_f = s_{c,f}\epsilon_c + (1 - s_{c,f})\epsilon_d$, the lower sectoral intensity $\epsilon^R_f < \epsilon^M_f$ necessitates a strictly higher clean technology share $s^R_{c,f} > s^M_{c,f}$. This formalizes the fundamental trade-off: by neutralizing the scale margin as an abatement channel, the rate-based design forces the foreign sector to rely more heavily on the composition margin to achieve a given environmental outcome. □

A.4 Proof of Lemma 2 (Reshuffling Priority)

Proof. The foreign producer allocates production from a technology portfolio $\{q_c, q_d\}$ to two markets $\{h, f\}$. The producer minimizes total tax-inclusive costs:

$$\min \sum_{j \in \{h,f\}} [c(s_{c,j}) q_j + t(\epsilon_j) q_j \cdot \mathbb{I}_{j=h}] \tag{32}$$

subject to $\sum_j q_{c,j} \leq q_c$ and $\sum_j q_{d,j} \leq q_d$. The first-order condition for allocating a clean unit to market h versus f yields a tax-saving term $\tau(\epsilon_d - \epsilon_c)$ that is strictly positive. Since costs are otherwise

identical across markets, the producer will always prioritize the clean variety for the export market h to minimize the border charge. □

A.5 Proof of Proposition 3 (Reshuffling and Capacity Pressure)

Proof. Consider a comparison between a rate-based and a mass-based design at the same marginal carbon price τ_f. Let $q_f = q_{f,h} + q_{f,f}$ be the total foreign production. By Lemma 2, producers prioritize the allocation of clean units to the regulated market. The observed intensity at the border, $\epsilon_{f,h}$, is the weighted average of the intensities of the clean ($q_{c,f,h}$) and dirty ($q_{d,f,h}$) units allocated to the export market:

$$\epsilon_{f,h} = \frac{q_{c,f,h}\epsilon_c + (q_{f,h} - q_{c,f,h})\epsilon_d}{q_{f,h}} \tag{33}$$

Solving for the volume of clean units required to hit the target $\bar{\psi}_{f,h}$ yields:

$$q_{c,f,h} = q_{f,h}\left[\frac{\epsilon_d - \bar{\psi}_{f,h}}{\epsilon_d - \epsilon_c}\right] \tag{34}$$

From Lemma 1, the rate-based design results in a lower price for the variety sold in the regulated market ($p^R_{f,h} < p^M_{f,h}$), which sustains a larger export volume $q^R_{f,h} > q^M_{f,h}$. At this higher export scale, the existing clean capacity $s_{c,f}q_f$ is diluted over a larger volume, making it more difficult to achieve the intensity target $\bar{\psi}_{f,h}$ through reshuffling alone.

To maintain the target $\bar{\psi}_{f,h}$ at the larger scale $q^R_{f,h}$, the required absolute clean production $q^R_{c,f,h}$ must satisfy:

$$q^R_{c,f,h} = q^R_{f,h}\left[\frac{\epsilon_d - \bar{\psi}_{f,h}}{\epsilon_d - \epsilon_c}\right] > q^M_{f,h}\left[\frac{\epsilon_d - \bar{\psi}_{f,h}}{\epsilon_d - \epsilon_c}\right] = q^M_{c,f,h} \tag{35}$$

Since the total volume of clean units must satisfy $s_{c,f}q_f \geq q_{c,f,h}$, the larger requirement $q^R_{c,f,h}$ necessitates an expansion of the total regional clean technology share $s_{c,f}$ in the foreign sector. Thus, by preserving trade scale, the rate-based design creates capacity pressure that forces deeper technology adoption to compensate for the dilution of reshuffled assets. □

A.6 Proof of Proposition 4 (Regional Incidence and Terms-of-Trade)

Proof. Let regional welfare be defined as the sum of producer surplus, downstream surplus, and fiscal revenue. Considering the impact of the border adjustment on Home welfare W_h and Foreign welfare W_f:

$$W_h = \Pi_{u,h} + \Pi_d + R_{BCA} \tag{36}$$

$$W_f = \Pi_{u,f} \tag{37}$$

where R_{BCA} is the fiscal revenue collected at the border. Under a mass-based design, $R^M_{BCA} = \tau_f\epsilon_{f,h}q_{f,h}$. Under a rate-based design, the intensity benchmark $\bar{\psi}$ creates an implicit subsidy $\sigma = \tau_f\bar{\psi}$,

such that $R^R_{BCA} = \tau_f(\epsilon_{f,h} - \bar{\psi})q_{f,h} = R^M_{BCA} - \sigma q_{f,h}$.

The difference in fiscal revenue collected by the Home government, $R^M_{BCA} - R^R_{BCA} = \sigma q_{f,h}$, constitutes a direct transfer of economic surplus to the Foreign export sector. Because the rate-based design lowers the effective border charge (Lemma 1), it increases Foreign producer surplus ($\Pi^R_{u,f} > \Pi^M_{u,f}$) and domestic downstream surplus ($\Pi^R_d > \Pi^M_d$). However, this gain for the downstream sector is offset by the loss of fiscal revenue. For a sufficiently small downstream value-added, $W^M_h > W^R_h$, as the mass-based design maximizes the extraction of scarcity rents from foreign producers, thereby improving the regional terms-of-trade. □

A.7 Proof of Proposition 5 (Industrial Incidence)

Proof. Let $\Pi_{u,h}$ denote the producer surplus of domestic upstream producers, which is strictly increasing in the price of the competing imported variety $p_{f,h}$. From Lemma 1, $p^M_{f,h} > p^R_{f,h}$, which implies higher surplus for domestic producers under the mass-based design: $\Pi_{u,h}(p^M_{f,h}) > \Pi_{u,h}(p^R_{f,h})$.

For the downstream sector, the surplus is $\Pi_d = \max_Q\{G(Q) - PQ\}$, representing the value-added accruing to fixed factors in downstream production. By the Envelope Theorem:

$$\frac{d\Pi_d}{dP} = -Q(P) < 0 \tag{38}$$

Since the dual price index P is increasing in input prices, $P(p_{h,h}, p^M_{f,h}) > P(p_{h,h}, p^R_{f,h})$. Therefore, downstream surplus is strictly higher under the rate-based design:

$$\Pi_d(P^R) > \Pi_d(P^M) \tag{39}$$

□

A.8 Proof of Proposition 6 (Reshoring and Diluted Abatement)

Proof. Total global emissions are defined as $E_{global} = \epsilon_{h,h}q_{h,h} + \epsilon_{f,h}q_{f,h} + \epsilon_{f,f}q_{f,f}$, where $\epsilon_{f,h}$ and $\epsilon_{f,f}$ denote the average intensities of foreign production allocated to the regulated and unregulated markets, respectively. Differentiating with respect to the border charge t:

$$\frac{dE_{global}}{dt} = \underbrace{(\epsilon_{h,h} - \epsilon_{f,h})\frac{dq_{h,h}}{dt}}_{\text{Substitution Effect}} + \underbrace{q_{f,h}\frac{d\epsilon_{f,h}}{dt} + q_{f,f}\frac{d\epsilon_{f,f}}{dt}}_{\text{Abatement/Reshuffling Effect}} \tag{40}$$

With $\tau_h = 0$, domestic intensity $\epsilon_{h,h}$ is constant. Since the BCA raises the price of the foreign variety, $\frac{dq_{h,h}}{dt} > 0$ (reshoring). If domestic production is dirtier than the foreign imports ($\epsilon_{h,h} > \epsilon_{f,h}$), the substitution effect is positive. While the foreign producer may clean up its exports ($\frac{d\epsilon_{f,h}}{dt} < 0$), if this cleanup is achieved primarily through reshuffling existing units from the domestic to the export

market rather than sectoral investment ($s_{c,f}$), then $d\epsilon_{f,f}/dt > 0$, and the net global abatement is diminished. □

A.9 Proof of Lemma 3 (The Knife-Edge Property)

Proof. In a rate-based design, the implicit marginal penalty δ^R is determined by the aggregate sectoral compliance condition: $\sum_i \bar{\psi} q_{i,h} = \sum_i \epsilon_{i,h} q_{i,h}$, which requires the aggregate sectoral intensity to equal the benchmark ($\bar{\epsilon} = \bar{\psi}$). Substituting the intensity definition $\epsilon_{i,h} = s_{c,i,h}\epsilon_c + (1 - s_{c,i,h})\epsilon_d$:

$$\bar{\epsilon} = \epsilon_d - s_{c,i,h}(\epsilon_d - \epsilon_c) = \bar{\psi} \tag{41}$$

Solving for the required clean technology share $s_{c,i,h}$ yields:

$$s^*_{c,i,h} = \frac{\epsilon_d - \bar{\psi}}{\epsilon_d - \epsilon_c} \tag{42}$$

As $\epsilon_c \to \epsilon_d$, the denominator approaches zero. If $\bar{\psi}$ is set below the physical limit of the technology ($\bar{\psi} < \epsilon_c \approx \epsilon_d$), the required share $s^*_{c,i,h}$ becomes mathematically undefined and exceeds 1, indicating that the target cannot be met at any price ($\delta^R \to \infty$). Conversely, if $\bar{\psi} > \epsilon_c \approx \epsilon_d$, the condition is always satisfied even at $s_{c,i,h} = 0$, causing the implicit penalty to collapse to zero ($\delta^R = 0$). Thus, the effective border penalty exhibits a knife-edge sensitivity to the benchmark level in homogeneous sectors. □

A.10 Proof of Proposition 7 (Vertical Leakage)

Proof. The price of the intermediate input in the domestic market for a domestic producer h is $p_h = c(s_{c,h}) + \tau_h \epsilon_h$. For a foreign producer f, the price inclusive of the border adjustment is $p_{f,h} = c(s_{c,f}) + t(\epsilon_f)$.

From Lemma 3, if the intermediate sector is homogeneous, $t(\epsilon_f) \approx 0$. Assuming production costs $c(\cdot)$ and intensities ϵ are approximately equal across regions for the homogeneous input, the resulting price wedge $\Delta p = p_h - p_{f,h}$ is:

$$\Delta p \approx [c(s) + \tau_h \epsilon] - [c(s) + 0] = \tau_h \epsilon \tag{43}$$

The domestic producer faces a competitive disadvantage equal to the full domestic carbon cost $\tau_h \epsilon$. Downstream firms, seeking to minimize the price index $P(p_h, p_{f,h})$, will substitute the domestic intermediate with imports $q_{f,h}$, shifting the emissions-intensive stage of production offshore. □

B Data Construction

This appendix details the construction of the global dataset, which harmonizes plant-level engineering estimates with aggregate trade and production statistics. The procedure ensures that the micro-level heterogeneity in technology and efficiency is preserved while satisfying macro-level constraints on global supply and demand.

B.1 Data Sources and Sample Selection

The core iron and steel production data is drawn from the CRU Steel Cost Service for the baseline year 2019. We extract key physical and economic parameters for each facility, including production capacities, operating rates, granular material inputs (e.g., kg per tonne of output), energy consumption vectors, and local unit costs for factors of production.

We apply a rigorous data cleaning and validation protocol to clean and harmonize the raw data. Plants are excluded if they report discontinuous operations, missing vital cost components, or negligible sales volumes (defined as less than 0.1 kt annually). Following the plant-level filtering, we map individual ISO-alpha3 country codes from plant locations into the model's distinct custom regions (e.g., USA, China, EU28/EFTA, Russia/CIS/Ukraine, and rest of world aggregates). This mapping balances computational tractability with the resolution required for accurate policy evaluation. We validate the representativeness of this cleaned sample by cross-referencing aggregate regional capacities against the Global Steel Plant Tracker (GSPT) and production totals against the World Steel Association (worldsteel).

B.2 Process Modeling and Carbon Accounting

Translating engineering cost accounts into economic production functions requires two major structural adjustments. First, we map granular sub-processes into the model's aggregated sectors. We incorporate the physical and economic profiles of upstream activities—specifically blast furnace pelletizing, direct reduction (DR) pelletizing, coke making, and sintering—and consolidate their costs and physical inputs recursively into the primary ironmaking functions (blast furnace, DRI, Corex, and custom ironmaking). Specific configurations, such as Conarc and custom iron or steelmaking setups, are standardized into representative BF or DRI functions. Crucially, the pipeline calculates net intermediate production to maintain mass-balance consistency. Pig iron and DRI outputs are dynamically split between captive consumption (used on-site for steelmaking) and merchant surplus (sold to third parties).

Second, we rigorously correct for byproducts to ensure accurate carbon accounting. Standard industry cost accounting frameworks typically treat process gases (e.g., blast furnace gas, coke oven gas) and surplus generated power or steam as revenue credits that artificially lower net operational costs. For a physical climate model, however, these must be treated as conserved energy flows to track emissions accurately. To resolve this, we calculate a mathematical correction factor ϕ^{bp} to

convert these financial credits back into gross physical input requirements:

$$\phi^{bp} = 1 - \frac{\text{Total Production}}{\text{Total Consumption}} \tag{44}$$

Here, "Total Consumption" refers to the absolute gross amount of an energy carrier a plant physically requires for its processes, while "Total Production" is the volume of that carrier the plant generates internally as a byproduct. The factor ϕ^{bp} calculates the share of the required energy that must be purchased externally. Multiplying the gross input requirement by ϕ^{bp} yields the net physical external input, preventing the double-counting of emissions from self-generated process gases and ensuring the model captures the true energy intensity of production. We also calculate and add the auxiliary energy requirements for desulfurization and secondary steelmaking to the core process definitions to capture the full, unabated energy footprint of semi-finished steel products.

B.3 Regional and Raw Material Balances

To reconcile the bottom-up microdata with official national accounts, we employ a systematic re-weighting procedure. Plant-level output is scaled to match authoritative worldsteel regional control totals for both intermediate inputs (pig iron, DRI) and final steel goods (flat and long products). We calculate a regional adjustment factor $\alpha_{r,s}$ for each region r and sector s:

$$\alpha_{r,s} = \frac{Y_{r,s}^{\text{worldsteel}}}{Y_{r,s}^{\text{CRU}}} \tag{45}$$

where Y denotes aggregate production. Each plant's output is multiplied by this factor, ensuring that the baseline simulation exactly reproduces observed regional distinctives in production scale without altering the underlying micro-level cost distributions or technological shares.

For primary raw materials, we construct physically consistent global balances with differing approaches depending on data availability. For iron ore and ferrous scrap, total consumption (C) is determined bottom-up from CRU plant requirements, while bilateral trade flows (exports X and imports M) are anchored to UN Comtrade statistics. Production (Y) is therefore derived as the residual ($Y = C + X - M$), with imports proportionally scaled down if they implausibly exceed total consumption. Missing Comtrade quantities are imputed using regional average prices. Natural gas bilateral trade (pipeline and liquefied natural gas) and production are sourced directly from the internally consistent BP Statistical Review, with consumption calculated as the residual.

In contrast to gas and ore, the dataset isolates the trade of metallurgical (coking) coal rather than thermal coal. Regional production and consumption data for metallurgical coal are sourced from the U.S. Energy Information Administration (EIA), but consistent bilateral trade flows are unavailable at the same coverage. Therefore, we assume bilateral flows follow the general coal trade patterns (HS code 270112) from UN Comtrade. To reconcile the discrepancies between the independent macroeconomic totals and the trade network, we frame the problem as a quadratic

programming optimization. This minimizes the squared deviation from the reported EIA totals subject to the global material balance constraint ($Y_r + M_r = C_r + X_r$, and $\sum_r X_r = \sum_r M_r$). After finding the optimal balanced macro totals, we apply Iterative Proportional Fitting (IPF) to adjust the Comtrade bilateral matrix to exactly match these new targets.

B.4 Trade Flow Imputation

Modeling trade with heterogeneous firms requires determining not just how much is traded between countries (intensive margin), but which specific plants participate in export markets (extensive margin). Since customs data (UN Comtrade) provides only aggregate bilateral flows, we impute plant-level trade linkages for final steel products (flat and long steel) through a three-step procedure that combines structural gravity estimation with our micro-level cost data. Intermediate iron products (pig iron and DRI) are allocated proportionally based on plant-level surplus production.

Calibration of exporting margins First, we calibrate the extensive margin—the share of firms in a region that export to a specific destination—using firm-level trade statistics from the OECD. We integrate multiple OECD datasets: Inter-Country Input-Output (ICIO) tables for trade flows (X_{ijs}), International Transport and Insurance Costs (ITIC) for CIF/FOB transport margins, Structural and Demographic Business Statistics (SDBS) for total firm counts (N_{iis}), and Trade by Enterprise Characteristics (TEC) for exporting firm counts (N_{is}). Following Sogalla et al. (2024), we estimate a gravity regression to predict the destination-specific exporting firm share s_{ijs}:

$$\ln(s_{ijs}) = \beta_0 + \beta_1 \ln(\lambda_{ijs}) + \beta_2 \ln(\text{dist}_{ij}) + \gamma_j + \epsilon_{ij} \tag{46}$$

where $\lambda_{ijs} = X_{ijs}/X_{iis}$ is the bilateral trade intensity, $\text{dist}_{ij} = 1/(1 - \text{ITIC}_{ij})$ is the transport cost friction, and γ_j are destination fixed effects. This estimation provides a region-pair specific target for the proportion of capacity that serves each export market.

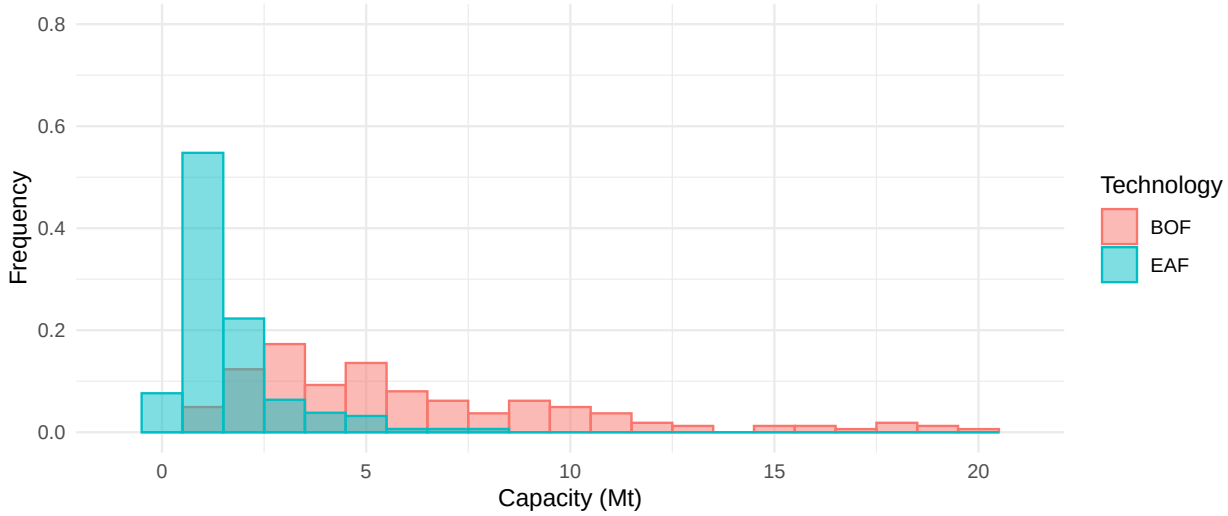
Ricardian selection Second, we identify the specific plants that fill these export quotas using a Ricardian selection mechanism. Consistent with trade theory, we assume that more productive (lower-cost) firms are more likely to overcome fixed trade costs. Within each region, we rank plants by their unit production costs. We determine the discrete plant count quota for exports as $\lceil s_{ijs} N_{is} \rceil$. We then assign export status to the lowest-cost plants until this quota is satisfied for each destination. This ranking mechanism ensures that the model preserves the theoretical correlation between productivity and export participation observed in microdata.

Bilateral flow reconciliation Finally, we determine the specific volume of trade for each plant-destination pair. Initial export volumes are allocated based on a gravity-style share equation where market share is inversely related to delivered costs (production plus transport), utilizing a constant

elasticity parameter. Because these initial allocations do not sum perfectly to observed national trade totals, we apply an IPF method to iteratively adjust the flows. This parallelized algorithm scales the plant-level exports until they simultaneously satisfy two constraints: (1) the sum of all exports and domestic sales from a plant exactly equals its modeled production capacity, and (2) the aggregation of plant-level flows between two regions exactly matches the harmonized bilateral trade volume reported in UN Comtrade.

C Additional Figures and Tables

(a) CRU Steel Cost Service (Number of Plants: 347)



(b) Global Steel Plant Tracker (Number of Plants: 932)

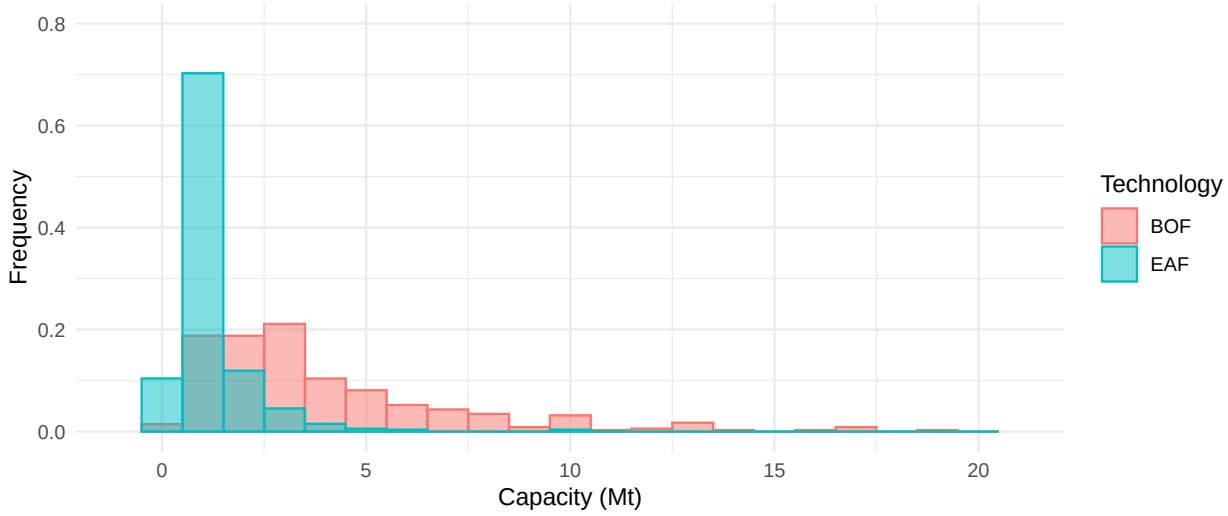


Figure A1: Histogram of Steel Plants by Production Capacity per Year

Note: Panels show the distribution of plant capacities; sample sizes are reported in the panel captions.
Source: CRU Steel Cost Service and Global Steel Plant Tracker.

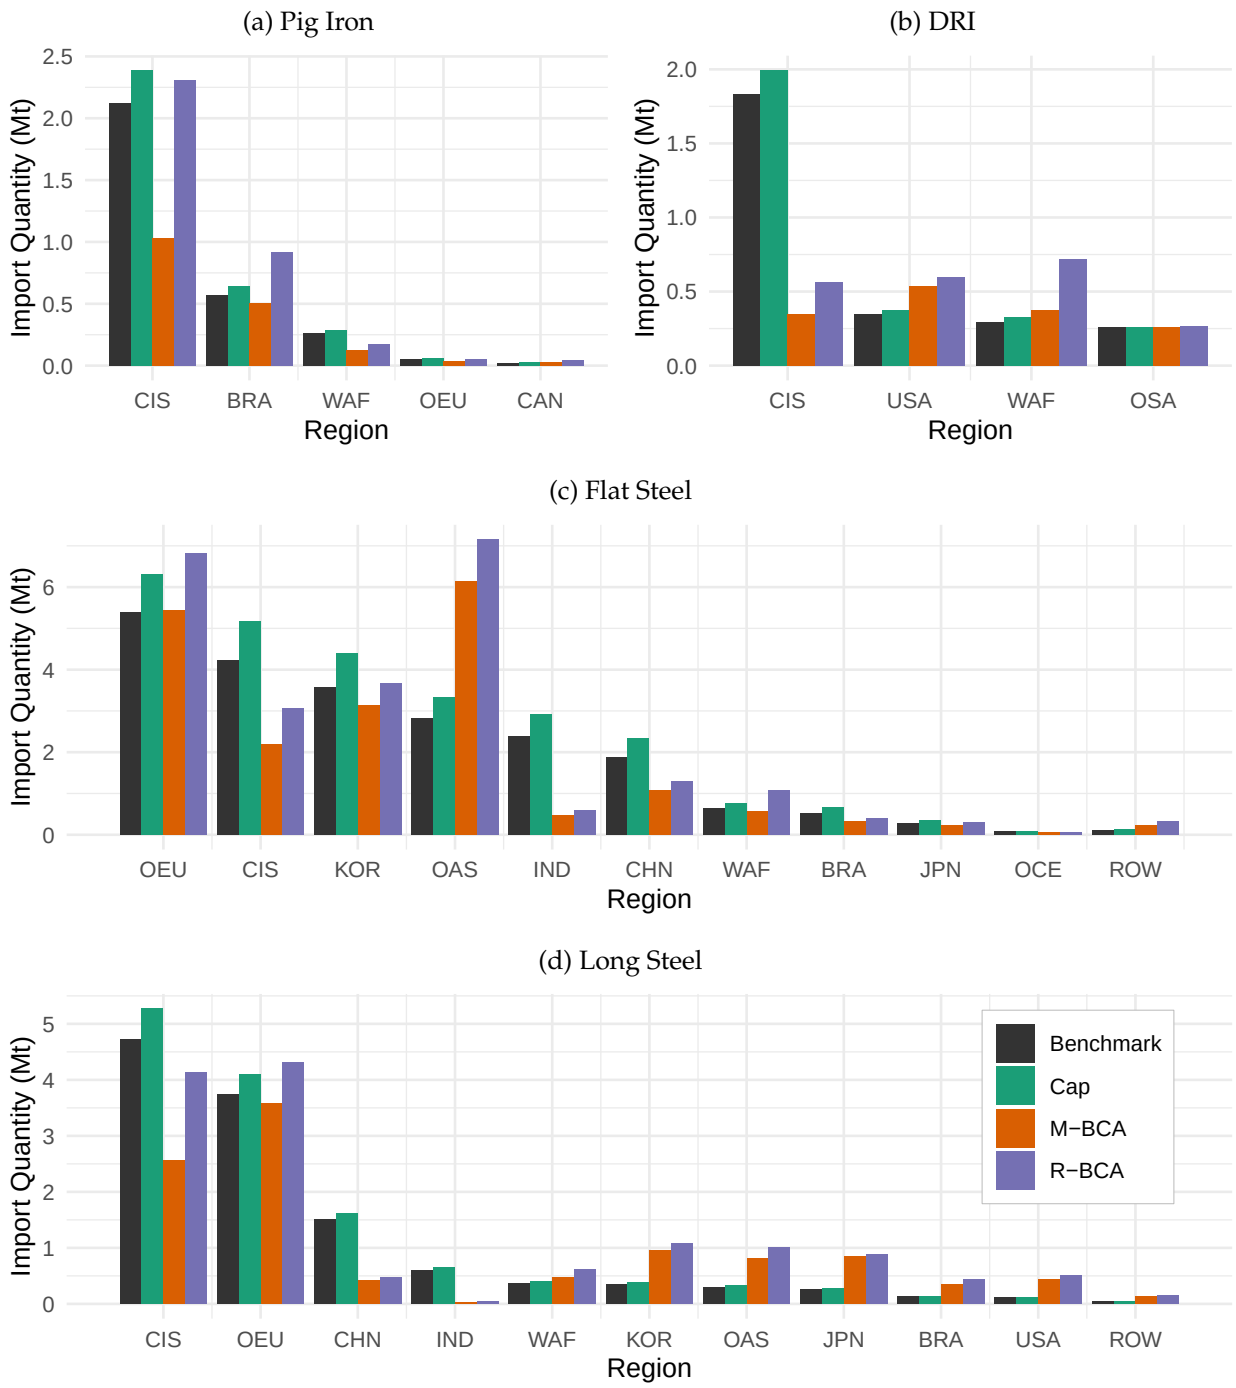


Figure A2: Import Quantity Changes by Product and Exporting Region

Note: Each panel illustrates the change in import quantities to the EU market for a specific steel product category under the mass-based (M-BCA) and rate-based (R-BCA) scenarios. The data is disaggregated by exporting region, highlighting the differential impacts of the two policy designs on trade flows. Region codes are defined in Section 5.1.

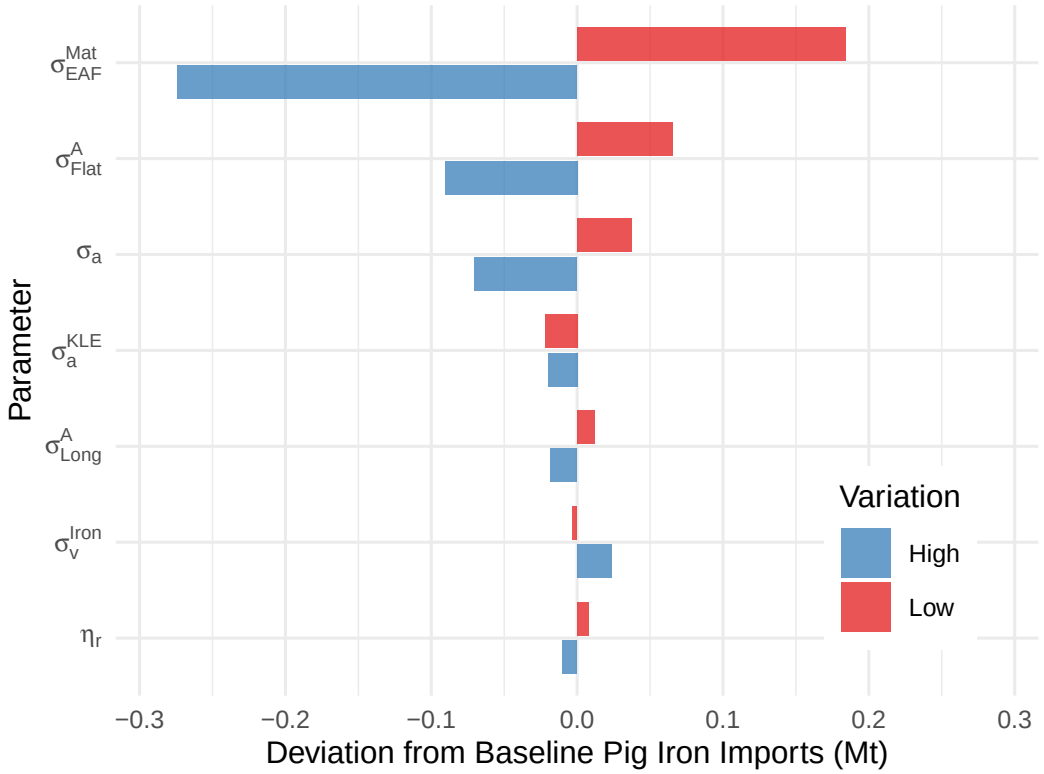


Figure A3: Sensitivity of EU Merchant Pig Iron Imports (Vertical Leakage)

Note: Bars represent the deviation from the baseline parameter volume of merchant pig iron imports (Mt) into the EU under the EU R-BCA policy design. The horizontal axis indicates the impact of moving each parameter from its calibrated baseline to either its Low (L) or High (H) setting.

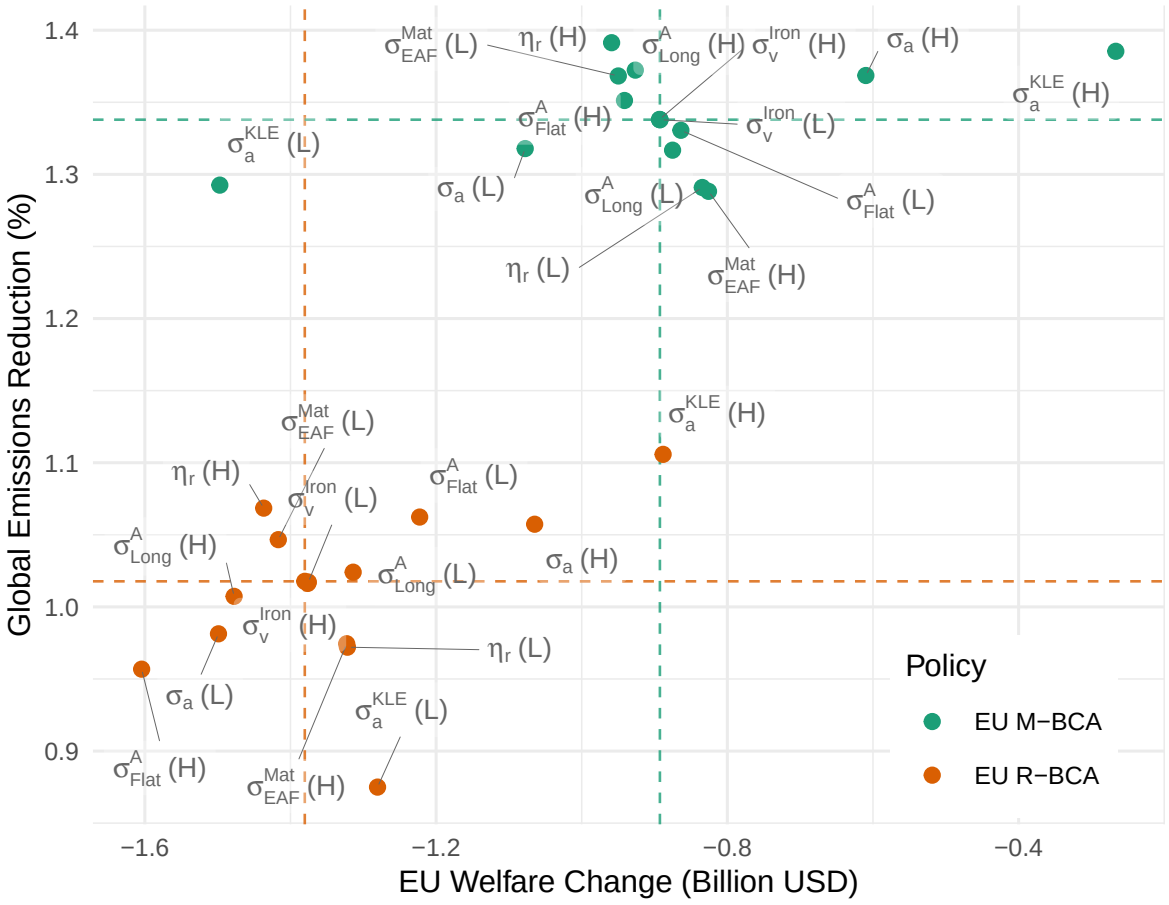


Figure A4: Sensitivity of Global Cost-Effectiveness Frontier

Note: The scatter plot maps global emissions reduction against regional welfare for all 28 sensitivity scenarios across two EU policy designs. Dashed lines mark the baseline parameter cost-effectiveness coordinates for the EU M-BCA and EU R-BCA. (H) and (L) denote High and Low parameter variations respectively.

C.1 Tables

Table A1: Calibrated Parameters

Panel A: Production Substitutions and Transformations

Symbol	Description	Value
Steel Production		
σ_a	Top-level	0.1
σ_a^{KLE}	KLE nest	0.4 (BF/DRF), 0.5 (BOF/EAF)
σ_a^{VA}	Value-added nest	1.0
σ_a^{Ene}	Energy nest	0.1 (BF/DRF), 0.2 (BOF/EAF)
σ_a^{Mat}	Material nest	0.3 (BOF), 0.4 (BF/DRF), 2.5 (EAF)
τ_a	Transformation	0.5
σ_v^{Iron}	Iron make-or-buy	0.1 (Pig iron), 0.5 (DRI)
Macro Production		
σ^C	Top-level	1.0
$\sigma_{\text{mac}}^{\text{VA}}$	Value-added nest	1.255
$\sigma_{\text{mac}}^{\text{Mat}}$	Material nest	0.3

Panel B: Trade and Final Demand Elasticities

Commodity	Domestic vs. Imported (σ^A)	Across Foreign Origins (σ^M)	Within Domestic (σ^D)
Flat steel	1.4	5.0	5.0
Long steel	1.4	8.0	8.0
Pig iron	1.4	10.0	10.0
DRI	1.4	10.0	10.0
Iron ore	0.9	1.8	–
Steel scrap	2.95	5.9	–
Coal	3.05	6.1	–
Natural gas	10.0	20.0	–
Macro good	2.46	4.92	–

Panel C: Factor Supply and Mobility Elasticities

Factor (f)	Factor Supply ($\eta_{r,f}$)	Plant Mobility (η_f^F)
Capital	–	0.1
Labor	–	2.0
Natural gas	2.0	2.9
Electricity	4.0	4.2
Iron ore	10.0	4.5
Coal	10.0	5.8
Steel scrap	0.5	13.8

Table A2: Market Summary Statistics for EU Scenarios

	Domestic (EUR → EUR)			Import (ROW → EUR)		
Scenario	Quantity	Intensity	Price	Quantity	Intensity	Price
Panel (a): Pig Iron						
Benchmark	4.93	2.12 (0.33)	0.35 (0.04)	3.03	2.67 (0.37)	0.32 (0.03)
Cap	3.50	1.90 (0.27)	0.49 (0.06)	3.41	2.67 (0.36)	0.32 (0.03)
M-BCA	3.45	1.87 (0.26)	0.52 (0.07)	1.73	2.52 (0.32)	0.54 (0.04)
R-BCA	3.51	1.90 (0.26)	0.50 (0.07)	3.50	2.52 (0.31)	0.32 (0.03)
DM-BCA	3.51	1.89 (0.26)	0.50 (0.07)	2.64	2.60 (0.34)	0.39 (0.03)
Panel (b): Long Steel						
Benchmark	60.71	0.58 (0.79)	0.45 (0.04)	12.11	2.07 (1.01)	0.40 (0.07)
Cap	58.54	0.44 (0.61)	0.49 (0.05)	13.32	2.06 (1.01)	0.40 (0.07)
M-BCA	60.47	0.42 (0.58)	0.50 (0.06)	10.59	1.04 (1.05)	0.49 (0.10)
R-BCA	58.88	0.43 (0.60)	0.49 (0.05)	13.65	1.04 (1.01)	0.39 (0.09)
DM-BCA	59.36	0.43 (0.60)	0.50 (0.05)	11.83	1.69 (1.12)	0.44 (0.08)

Note: Values represent the production-weighted mean and standard deviation (in parentheses) for domestic production (EU) and imports (ROW) across scenarios. Intensity is measured in tCO_2/ton of product. Quantity is measured in Mt. Price represents the delivered price in 1000 USD/ton.

Table A3: Reshuffling Index (R) for EU Scenarios

Product	Benchmark	Cap	M-BCA	R-BCA	DM-BCA
Pig Iron	1.02	1.02	0.96	0.96	1.00
DRI	0.70	0.70	0.66	0.66	0.69
Flat Steel	0.85	0.86	0.53	0.53	0.74
Long Steel	1.08	1.09	0.55	0.55	0.90

Note: The Reshuffling Index (R) is defined as the ratio of the emissions intensity of exports to the regulating region relative to the producer's domestic average intensity ($R = e_{exp}/e_{dom}$). Values less than 1 indicate that the producer is exporting cleaner-than-average output to the regulating region, while values greater than 1 indicate that dirtier-than-average output is being exported. The table compares R under EU scenarios.

Table A4: Robustness of Policy Efficiency

Param.	Var.	EU M-BCA		EU R-BCA		US M-BCA		US R-BCA	
		Welf. ($bn)	Red. (%)	Welf. ($bn)	Red. (%)	Welf. ($bn)	Red. (%)	Welf. ($bn)	Red. (%)
Baseline	Base	-0.89	1.34	-1.38	1.02	-0.87	0.26	-0.32	0.06
σ_v^{Iron}	Low	-0.89	1.34	-1.38	1.02	-0.88	0.26	-0.32	0.06
σ_v^{Iron}	High	-0.89	1.34	-1.38	1.02	-0.88	0.26	-0.32	0.06
$\eta_{r,\text{Scrap}}$	Low	-0.83	1.29	-1.32	0.97	-0.88	0.24	-0.31	0.06
$\eta_{r,\text{Scrap}}$	High	-0.96	1.39	-1.44	1.07	-0.86	0.28	-0.32	0.07
σ^A_{Flat}	Low	-0.86	1.33	-1.22	1.06	-0.75	0.25	-0.31	0.06
σ^A_{Flat}	High	-0.94	1.35	-1.60	0.96	-1.07	0.28	-0.34	0.07
σ^A_{Long}	Low	-0.88	1.32	-1.31	1.02	-0.85	0.26	-0.31	0.06
σ^A_{Long}	High	-0.93	1.37	-1.48	1.01	-0.92	0.26	-0.33	0.07
$\sigma^{\text{Mat}}_{\text{EAF}}$	Low	-0.95	1.37	-1.42	1.05	-0.68	0.22	-0.32	0.07
$\sigma^{\text{Mat}}_{\text{EAF}}$	High	-0.83	1.29	-1.32	0.97	-1.09	0.29	-0.33	0.06
σ_a^{KLE}	Low	-1.50	1.29	-1.28	0.87	-0.61	0.23	-0.21	0.05
σ_a^{KLE}	High	-0.27	1.39	-0.89	1.11	-0.74	0.26	-0.26	0.06
σ_a	Low	-1.08	1.32	-1.50	0.98	-0.90	0.26	-0.32	0.06
σ_a	High	-0.61	1.37	-1.06	1.06	-0.73	0.24	-0.26	0.06

Note: Results report regional welfare change (billion USD) and global emissions reduction (%) relative to the no-policy equilibrium. Scenarios are evaluated at a fixed stringency of $\zeta = 0.7$ (corresponding to a 30% reduction in domestic emissions for the EU). Parameters are varied by a factor of 0.5 (L) or 2.0 (H) relative to calibrated values, except for σ^I which uses a 5.0 factor for High (H) elasticity.

Table A5: Sensitivity of Leakage and Protection

Param.	Var.	EU M-BCA		EU R-BCA		US M-BCA	US R-BCA
		Leak. (%)	Pig I. (Mt)	Leak. (%)	Pig I. (Mt)	Rev. Lk. (%)	Rev. Lk. (%)
Baseline	Base	16.0	1.73	36.1	3.50	23.0	14.9
σ_v^{Iron}	Low	16.0	1.73	36.2	3.49	23.0	14.9
σ_v^{Iron}	High	16.0	1.72	36.2	3.52	23.1	14.9
$\eta_{r,\text{Scrap}}$	Low	19.0	1.73	39.0	3.50	24.8	16.5
$\eta_{r,\text{Scrap}}$	High	12.7	1.73	33.0	3.49	21.3	13.3
σ_{Flat}^{A}	Low	16.5	1.73	33.3	3.56	15.2	9.5
σ_{Flat}^{A}	High	15.2	1.72	40.0	3.41	31.8	23.4
σ_{Long}^{A}	Low	17.4	1.73	35.7	3.51	22.2	13.0
σ_{Long}^{A}	High	13.9	1.74	36.8	3.48	24.3	17.8
$\sigma_{\text{EAF}}^{\text{Mat}}$	Low	14.1	1.89	34.3	3.68	31.4	14.9
$\sigma_{\text{EAF}}^{\text{Mat}}$	High	19.2	1.54	38.9	3.22	14.7	15.1
σ_a^{KLE}	Low	18.9	1.58	45.1	3.47	21.5	12.6
σ_a^{KLE}	High	13.1	1.95	30.6	3.48	22.1	12.6
σ_a	Low	17.3	1.70	38.4	3.53	22.1	15.0
σ_a	High	14.1	1.79	33.7	3.43	24.0	13.3

Note: Leak. refers to the global leakage rate for the EU scenarios. Pig I. refers to the volume of merchant pig iron imports into the EU (Mt). Rev. Lk. refers to reverse leakage (onshoring) in the US scenarios.

Table A6: Abatement Margins in the ROW (EU R-BCA)

Param.	Var.	Scale (Mt)	Composition (Mt)	Technique (Mt)	Total (Mt)
Baseline	Base	37.77	-3.71	2.96	37.02
σ_v^{Iron}	Low	37.81	-3.70	2.96	37.07
σ_v^{Iron}	High	37.81	-3.70	2.97	37.08
$\eta_{r,\text{Scrap}}$	Low	38.84	-2.51	4.22	40.55
$\eta_{r,\text{Scrap}}$	High	36.64	-5.06	1.53	33.11
σ^A_{Flat}	Low	34.95	-3.98	2.99	33.97
σ^A_{Flat}	High	41.70	-3.41	2.90	41.19
σ^A_{Long}	Low	37.41	-3.64	2.95	36.72
σ^A_{Long}	High	38.37	-3.81	2.98	37.54
$\sigma^{\text{Mat}}_{\text{EAF}}$	Low	36.69	-3.50	1.93	35.11
$\sigma^{\text{Mat}}_{\text{EAF}}$	High	39.35	-3.99	4.41	39.76
σ_a^{KLE}	Low	47.65	-1.80	2.60	48.45
σ_a^{KLE}	High	30.78	-3.48	3.64	30.94
σ_a	Low	40.23	-3.61	2.98	39.61
σ_a	High	34.73	-3.23	3.14	34.64

Note: The table decomposes the rest of the world (ROW) emissions response (Mt CO_2) for the intensity-based design into scale, composition, and technique effects using an LMDI decomposition. Scale represents production volume shifts; Composition represents shifts between primary and secondary routes; Technique represents within-technology efficiency gains.