Green Industrial Policy in a Globalized Economy [*†]

Sungwan Hong
University of Pittsburgh

July 2026

Abstract

Renewable-energy equipment is traded internationally and produced in a few countries, while the electricity it generates is largely local. I develop a dynamic multi-country trade model with traded renewable equipment and production learning, in which equipment trade shapes the cross-border incidence of subsidies. Theory characterizes when global welfare favors deployment support in an importing market and equipment support at a production hub, and when national incentives select the reverse assignment. At a common fiscal cost in the calibrated economy, global consumption welfare favors deployment support in the importing United States and equipment support at the Chinese production hub. At unilateral preferred rates, national welfare favors the reverse assignment: U.S. equipment support and Chinese deployment support. The two globally favored policies have opposite foreign-emissions effects: U.S. deployment support raises them, whereas Chinese equipment support lowers them. Removing equipment trade reverses the global-consumption ranking of the two U.S. instruments and eliminates most of the global-consumption advantage of locating equipment support in China rather than in the United States.

JEL: F13, F18, O25, Q42, Q54, Q56

Keywords: Green Industrial Policy, Inflation Reduction Act, Energy Transition, renewable-energy equipment trade

[*]I am grateful to Fernando Parro, Jingting Fan, and Stephen Yeaple for their guidance. I appreciate suggestions and comments from Kai-Jie Wu, Seula Kim, Yan Bai, Xiang Ding, Sharat Ganapati, Costas Arkolakis, Haruka Takayama, David Hemous, Konstantin Kucheryavyy, Johannes Boehm, Eugenie Dugoua, David Atkin, George Alessandria, Allan Hsiao, and seminar and conference participants. I am especially grateful to Jevan Cherniwchan, Dongkyu Yang, and Holger Görg for valuable discussions. All errors are mine.

[†]Correspondence by email: sungwan.hong@pitt.edu

1 Introduction

Governments can subsidize either renewable deployment or domestic production of renewable equipment. The U.S. Inflation Reduction Act (IRA) combines both approaches, and related policies operate in Europe, China, India, and Japan.[1] Their effects depend on the instrument and on whether support comes from a large importing market or an established production hub. The policy choice therefore helps determine where production, learning, and emissions reductions occur.

That choice matters because renewable electricity and renewable equipment have different economic geographies. Electricity is generated and consumed mostly within national markets, whereas solar and wind equipment is traded internationally and its production is concentrated in a few countries. China manufactures roughly 80 percent of the world's solar equipment, and the United States imports about 86 percent of what it uses. A deployment subsidy immediately expands renewable generation in the installing country, but much of the resulting equipment demand and production experience may accrue abroad. An equipment subsidy instead shifts production toward the supporting country and builds origin-specific experience there, but it also lowers current prices for foreign buyers and, through learning, their future installation costs.[2]

I study these forces in a dynamic multi-country trade model with traded renewable equipment, local electricity, forward-looking investment in renewable capital, and learning-by-doing. Equipment productivity depends on both world experience and production accumulated in each origin. Policies change current sourcing and income as well as the future geography of comparative advantage and the cost of renewable deployment. Fossil-fuel trade creates conventional carbon leakage through the world fossil price. Equipment trade adds a second leakage margin: deployment support can initially bid up traded equipment prices and slow renewable investment abroad, although the production it induces subsequently lowers costs through learning.

The analytical section characterizes the incidence and assignment of the two subsidies when carbon pricing is unavailable. A deployment subsidy raises equipment demand in one destination and allocates the resulting production and learning across origins according to sourcing shares. An equipment subsidy raises production at one origin—partly by displacing rival suppliers—and lowers purchaser prices wherever buyers source from that origin. Comparing welfare returns per fiscal dollar across instruments and locations shows when global welfare favors deployment support in the importing country and equipment support at the production hub. Because national governments value only domestic incidence, they can prefer the reverse assignment.

The main quantitative result has this importer–hub form: at a common fiscal cost, global consumption welfare favors deployment support in the importing United States and equipment support

[1] Examples include the EU's Green Deal Industrial Plan, India's Production-Linked Incentive scheme for solar modules, and Japan's Green Transformation plan.

[2] If electricity were freely traded, deployment support could also change foreign generation directly through power exports, weakening this separation.

at the Chinese production hub. Deployment support in the United States produces a positive global consumption gain and the smaller U.S. loss. U.S. equipment support shifts production toward a relatively small domestic base and has a slightly negative global consumption effect. In China, deployment support produces the larger domestic gain, but equipment support produces the larger global consumption gain and the larger discounted emissions reduction. The assignment is consistent with comparative advantage and the geography of learning: U.S. deployment support routes demand through the existing supplier network, whereas Chinese equipment support reinforces the established hub and generates more global learning than Chinese deployment support.

Equipment trade is central to this ranking because it governs which foreign suppliers expand or contract and which buyers receive a policy-induced price benefit. The emissions response also differs across instruments and locations. At a common fiscal cost, deployment support lowers emissions in the policy country but raises them abroad in both the U.S. and Chinese experiments. Deployment support can generate leakage through lower world fossil prices and, initially, higher renewable-investment costs abroad. The direct equipment-price effect of equipment support works in the opposite direction by lowering delivered prices in destinations that source from the subsidized origin. Foreign emissions rise slightly under U.S. equipment support but fall under Chinese equipment support. U.S. support passes relatively little price relief abroad and partly expands domestic production by displacing established foreign suppliers, whereas Chinese support displaces less foreign production and transmits a larger price reduction through its export network.

The covered equipment sector is small relative to the world economy, and the aggregate consumption effects are correspondingly small. The substantive question is whether this limited trade margin nevertheless changes instrument rankings and the relative return to policy locations.

Shutting down equipment trade while retaining world-experience learning reverses the U.S. global-consumption-welfare ranking: equipment support becomes more valuable than deployment support. It also reduces by about three quarters the global-consumption-welfare gap between locating equipment support in China and in the United States. The pattern reflects the two cross-border margins removed by the shutdown: U.S. equipment support no longer reallocates production away from foreign suppliers, while Chinese equipment support no longer lowers foreign investment costs through its export network. Equipment trade therefore affects both instrument choice in an importing market and the value of locating production support at an export hub.

National and global rankings need not select the same instrument. The unilateral preferred rates are about 3 percent for U.S. deployment support, 9 percent for U.S. equipment support, and 16 percent for either Chinese instrument. At these rates, national welfare favors U.S. equipment support and Chinese deployment support. Yet at both the equal-budget and unilateral preferred rate sets, the opposite assignment—U.S. deployment with Chinese equipment support—delivers the largest global consumption gain among the four fixed-rate assignments. Incidence helps explain the divergence: the national choices concentrate production or direct price benefits at home, while the global assignment

routes U.S. demand through the supplier network and spreads the price and learning benefits of Chinese equipment support abroad.

A stylized IRA-scale package combining U.S. deployment and equipment subsidies reduces discounted world emissions by 1.85 $GtCO_2$, but lowers U.S. and global consumption welfare before avoided climate damages are counted. Valuing the emissions decline at \$210 per tCO_2 makes the global effect positive. This does not make industrial subsidies a substitute for carbon pricing. The subsidies target learning in a small traded-input sector; a carbon price acts on fossil use throughout the economy.

The quantitative implementation covers 15 countries and regions, with trade in renewable equipment, fossil fuel, and other goods. Solar and wind electricity are produced locally, installed capacity evolves over time, and the model traces the transition through 2100. I construct the 2022 benchmark from input–output accounts, bilateral equipment trade, renewable capacity and generation, fossil use, and emissions. For solar, the decline in world module prices disciplines total learning, while bilateral exporter performance disciplines the origin-specific component. The estimates imply a 27.5 percent price decline per doubling of cumulative production and an origin-specific learning weight of about 0.12; variation in exporters' exposure to foreign feed-in-tariff expansions provides corroborating evidence. Wind learning is calibrated rather than estimated, informed by shorter price histories and manufacturer-location proxies.

Taken together, the results show that equipment trade is central to the assignment of green subsidies even though its aggregate consumption effects are small. The globally preferred instrument depends on trade patterns, local abatement, international-price incidence, and marginal learning values; the instruments favored by national welfare need not coincide with that assignment.

The model and results bring together four strands of research.

Industrial policy and production learning. Classical analyses study infant-industry protection and production externalities (Baldwin, 1969; Bardhan, 1971; Young, 1991; Melitz, 2005), while recent empirical and quantitative work examines temporary protection and sector targeting (Juhász, 2018; Liu, 2019; Lane, 2025; Goldberg et al., 2024); see Juhász et al. (2024) for a survey and Rodrik (2014) for the green-industrial-policy case. Related evidence documents production learning in semiconductors, shipbuilding, batteries, and renewable technologies (Irwin and Klenow, 1994; Thornton and Thompson, 2001; Barwick et al., 2025b,a; Way et al., 2022). This paper places these forces in a multi-country setting where producer support expands world experience while reallocating production across suppliers. Because the targeted good is renewable equipment, these production and learning effects also shape renewable investment and carbon emissions across countries.

Clean growth and climate policy. A large literature studies clean growth, renewable deployment, and climate policy, mostly in closed or effectively single-country settings (Acemoglu et al., 2012, 2016; Fried, 2018; Bistline et al., 2023; Casey et al., 2026; Donald, 2024; Gentile, 2024; Aghion et al., 2025; Hahn et al., 2026). Related open-economy work examines unilateral environmental

policy, clean comparative advantage, and international subsidy spillovers (Hémous, 2016; Kotchen and Maggi, 2025). This paper distinguishes destination-based deployment support from origin-based equipment support and traces their cross-border incidence through equipment trade and production learning. The closest quantitative antecedent is Arkolakis and Walsh (2023), who combine world-experience learning with a spatial electricity system; relative to their framework, I model international trade in renewable equipment, distinguish world experience from origin-specific experience, and compare the two subsidy instruments.

Dynamic quantitative trade. Conditional on inherited states, the production and sourcing block follows multi-sector gravity models with input–output linkages (Eaton and Kortum, 2002; Caliendo and Parro, 2015). Renewable equipment plays the role of a traded investment good (Mutreja et al., 2018), while capacity accumulation and the transition solution build on dynamic quantitative trade models (Eaton et al., 2016; Ravikumar et al., 2019; Caliendo et al., 2019; Kleinman et al., 2023). As in the classic learning-and-trade literature and recent models of international technology diffusion, current trade patterns can change future comparative costs (Krugman, 1987; Young, 1991; Bai et al., 2024). Here, a temporary policy reshapes equipment trade; production at each origin accumulates into experience, which changes subsequent trade and investment.

Trade and the environment. Research on trade and the environment studies leakage, border adjustment, and the location of emissions-intensive production (Copeland et al., 2021; Shapiro, 2021; Kortum and Weisbach, 2021; Caliendo et al., 2024; Farrokhi and Lashkaripour, 2025), while related work examines clean innovation and international environmental policy (Xiang, 2023; Hsiao, 2024). Recent quantitative models study power investment with internationally traded fossil fuels and the adjustment of capital and labor to climate policy (Abuin, 2026; Shin, 2026). This paper instead studies trade in renewable-energy equipment, a capital input that determines future local abatement costs. Such trade changes current international incidence, where production learning accumulates, and the future cost of renewable deployment. Environmental effectiveness consequently depends on the joint response of terms of trade, learning, and fossil use.

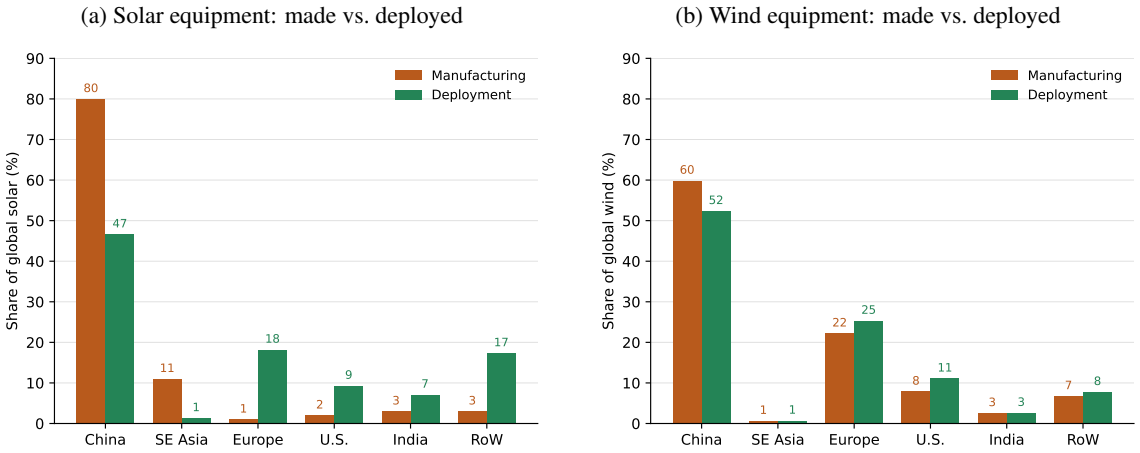
The remainder of the paper proceeds as follows. Section 2 documents the geography of renewable equipment production, deployment, trade, and cost decline. Section 3 presents the model, and Section 4 derives its incidence results. Section 5 describes the calibration. Section 6 reports the policy comparisons and mechanism diagnostics. Section 7 concludes.

2 Background

2.1 Production, Deployment, and Trade in Renewable-energy Equipment

Renewable electricity is generated locally, but the equipment that builds generating capacity is traded internationally and produced in a small number of countries. Figure 1 compares the geography of equipment manufacturing with the geography of deployment. These regional shares provide a

Figure 1: Production and Deployment of Solar and Wind Equipment, 2022



Note: Bars report regional shares of 2022 equipment manufacturing and capacity additions. Manufacturing shares are calibration targets. Appendix B.3 gives the sources.

descriptive motivation for equipment trade; they do not themselves measure bilateral trade in final equipment or components.[3]

For solar (panel (a)), China and Southeast Asia together produce about 90% of global equipment but account for only about 48% of deployment. Europe, the United States, and the rest of the world deploy far more than they manufacture; the United States installs about 9% of global capacity while producing about 2%. Trade therefore connects production in a few locations to deployment across many markets. For wind (panel (b)), manufacturing and deployment are more closely aligned: China and Europe produce about 82% of global equipment and deploy a comparable share, while the United States produces about 8% and installs about 11%. These aggregate shares locate production and installation; bilateral sourcing—of blades, nacelles, towers, and other components—is measured separately by the equipment-expenditure matrices that discipline the model's trade patterns.

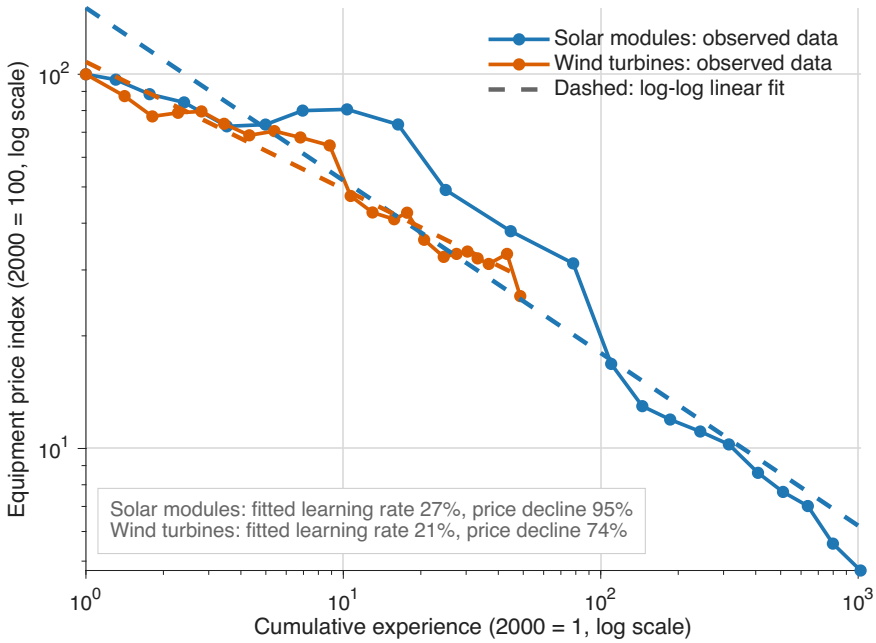
Direct trade data sharpen the distinction between local electricity and tradable equipment. Gross cross-border electricity imports amounted to 2.8% of total electricity supplied in 2021, so renewable generation remains overwhelmingly local.[4] Renewable equipment, by contrast, is internationally sourced. Imports account for 53% of world solar-equipment absorption in 2022, and the United States imports 86% of the solar equipment it installs. Wind equipment appears more home-biased within the product coverage used here: imports account for 35% of world absorption, and the United States sources about 45% of its wind equipment domestically.[5] The relevant distinction is therefore

[3] For details on the data used in this section, see Appendix B.3.

[4] World Meteorological Organization and World Trade Organization (2024, p. 22) measure international electricity trade as the sum of countries' gross imports and report a 2021 total of 809 TWh.

[5] The import shares are measured across individual countries. Aggregating to the model's 15 regions nets out intra-regional—especially intra-European—trade and yields calibrated import shares of 47% for solar and 14% for wind. The wind matrix covers generating sets and selected parts but omits other components and production networks, so it is a partial

Figure 2: Equipment Prices and Cumulative Experience



Note: Prices are indexed to 100 in 2000. Solar experience is cumulative global PV module production from IEA PVPS (IEA PVPS, 2024); wind experience is cumulative global installed wind capacity. Dashed lines are fitted log–log relationships.

between locally generated electricity and internationally sourced capital equipment, with sourcing patterns allowed to vary by technology.

2.2 Equipment Cost Declines and Learning-by-Doing

Unlike fossil generation, solar and wind generation require no purchased fuel, and much of their cost is embodied in the equipment installed at the outset. Declines in equipment cost have therefore been central to the falling cost of renewable electricity. Figure 2 plots equipment price indexes against cumulative experience. Solar module prices fell by 95% between 2000 and 2021, while wind turbine prices fell by 74%. A simple log–log fit implies price declines per doubling of cumulative experience of about 27% for solar modules and 21% for wind turbines.

Experience curves summarize the empirical regularity that unit costs tend to fall as cumulative output expands. In equipment manufacturing, repeated production can improve line design, yields, worker proficiency, supplier coordination, and the organization of production. This is the learning-by-doing channel studied in the renewable-energy literature.[6] The relationship in Figure 2 is nevertheless descriptive: cumulative output also moves with R&D, input prices, scale, and

measure of wind-supply-chain openness.

[6] See, among others, Baker et al. (2013), Bolinger et al. (2022), Covert and Sweeney (2024), and Arkolakis and Walsh (2023).

other experience-linked improvements. I therefore use the price–experience relation to discipline the model's learning process, rather than interpret the fitted slope as a causal estimate of learning alone. Appendix B.5 describes the estimation and supporting evidence.

The geographic reach of learning is a separate question. Some improvements are codified in designs, machinery, and common suppliers and can diffuse across producing countries. Other gains depend on production-line experience, worker skills, and local supplier relationships and may remain tied to the origin where production occurs. The model allows output to build both a world experience stock, which lowers future costs at all origins, and an origin-specific experience stock, which improves the relative productivity of the producing origin. The first component spreads learning gains internationally; the second gives the location of manufacturing persistent value.

These observations motivate the model's dynamic channel. A subsidy changes current prices and production, where production experience accumulates, and who benefits from the resulting cost decline. The next section formalizes that link between equipment trade, production experience, and future productivity.

3 A Dynamic Trade Model of Green Industrial Policy

Countries install solar and wind capacity using equipment sourced from a global market. Equipment trade determines where equipment is produced and where production experience accumulates; accumulated experience raises future equipment productivity and changes comparative advantage. Electricity generation and fossil displacement occur locally, while renewable equipment, fossil fuel, and the non-energy composite are traded. Within each year, inherited states and forward-looking investment choices jointly determine prices, trade, production, and fossil use. Current investment updates installed capacity, which affects future capacity accumulation, while current real equipment output updates experience.

Within each period, production and sourcing follow a multi-sector Eaton–Kortum economy with input–output linkages (Eaton and Kortum, 2002; Caliendo and Parro, 2015). Renewable equipment is a traded input into renewable investment, connecting the model to the capital-goods-trade literature (Mutreja et al., 2018), and renewable capacity follows the constant-elasticity adjustment technology used in dynamic quantitative trade models (Ravikumar et al., 2019). Production experience is an additional state variable. Setting the origin-specific weight to zero recovers the world-experience learning specification in Arkolakis and Walsh (2023); a positive weight allows comparative advantage to evolve with the location of production.

3.1 Geography, Sectors, and Timing

There are N countries and discrete time $t = 1, 2, \ldots$. Renewable technologies are $r \in \mathcal{R} \equiv \{S, W\}$. The traded composites are $\mathcal{J}^T \equiv \{O, F, M^S, M^W\}$, where O is a non-energy composite, F is fossil

fuel, and M^S and M^W are solar and wind equipment. A superscript M^r therefore denotes a variable associated with the equipment sector for technology r. Solar and wind generation are local. The non-energy composite is used for consumption, production, investment, balance-of-system inputs, and operations and maintenance; energy services enter household demand and are used as an input by non-energy and equipment producers.

The beginning-of-period state is

$$\mathcal{S}_t = \left\{ K^D_{n,t}, K^S_{n,t}, K^W_{n,t}, Q^S_{n,t}, Q^W_{n,t} \right\}_{n=1}^{N}.$$

Here K^D is non-energy capital, K^S and K^W are renewable capacities, and the Q's are stocks of equipment-production experience. End-of-2022 capacities form the capital state inherited in 2023. The experience state inherited in 2023 combines the beginning-of-2022 stock with production during 2022, as specified below. Given the state and investment choices, wages, prices, sourcing, production, and absorption are jointly determined within the year; investment and real equipment output update the next state. Labor endowments are fixed, while fossil extraction follows a country-specific supply curve. World labor income is the annual numeraire, $\sum_n w_{n,t} L_n = 1$. Households, firms, and the national investment fund have perfect foresight over policy paths and equilibrium prices when making intertemporal choices.

Preferences. The representative household supplies labor inelastically. Consumption combines the non-energy composite and local energy services:

$$C_{n,t} = \left(\frac{C^O_{n,t}}{\mu^O_n} \right)^{\mu^O_n} \left(\frac{C^E_{n,t}}{\mu^E_n} \right)^{\mu^E_n}, \qquad \mu^O_n + \mu^E_n = 1,$$
$$P^C_{n,t} = (P^O_{n,t})^{\mu^O_n} (P^E_{n,t})^{\mu^E_n}.$$

Lifetime utility is

$$U_n = \sum_{t=1}^{\infty} \beta^{t-1} \log C_{n,t}. \tag{1}$$

Each period is one year; $t = 1$ is 2023, $T = 78$ is 2100, and $\beta = 0.96$.

3.2 Production and Bilateral Trade

The equipment sectors link deployment in each destination to production across domestic and foreign origins. The non-energy composite and fossil fuel enter the same multi-sector Eaton–Kortum structure, so wages, intermediate costs, terms of trade, and fossil prices adjust jointly with equipment sourcing. Sector-j productivity is Fréchet with location $A^j_{i,t}$ and trade elasticity θ^j. Equipment

producers use labor, the non-energy composite, and local energy services, with unit cost

$$c_{i,t}^{M^r} = B_i^r w_{i,t}^{\zeta_i^{L,r}} (P_{i,t}^E)^{\zeta_i^{E,r}} (P_{i,t}^O)^{1-\zeta_i^{L,r}-\zeta_i^{E,r}}, \qquad r \in \mathcal{R}.$$

Here B_i^r is a cost shifter, $\zeta_i^{L,r}$ is the labor share, and $\zeta_i^{E,r}$ is the energy share. The non-energy sector uses labor, domestic capital, local energy services, and its own composite as intermediates; $r_{i,t}^D$ denotes the rental on non-energy capital. Fossil extraction responds to the origin price relative to the wage according to the isoelastic supply curve

$$F_{i,t}^s = \varphi_i^F \left(\frac{p_{i,t}^F}{w_{i,t}} \right)^{\varepsilon_F},$$

where φ_i^F is calibrated to 2022 extraction. The benchmark $(\varepsilon_F, \eta_F) = (1, 1/2)$ admits a competitive labor–fixed-resource interpretation, with labor receiving the revenue share $\eta_F = \varepsilon_F/(1+\varepsilon_F)$ and the remainder accruing as resource rent. Appendix D gives the complete unit-cost and factor-payment system.

Let κ_{ni}^j be the iceberg cost, $\tau_{ni,t}^j$ an import tariff, and $s_{i,t}^j$ an origin subsidy, which is zero outside renewable equipment. Define delivered unit cost by

$$d_{ni,t}^j = \kappa_{ni}^j (1+\tau_{ni,t}^j)(1-s_{i,t}^j) c_{i,t}^j.$$

The Eaton–Kortum price index and bilateral expenditure share are

$$P_{n,t}^j = \gamma^j \left[\sum_i A_{i,t}^j (d_{ni,t}^j)^{-\theta^j} \right]^{-1/\theta^j},$$
$$\pi_{ni,t}^j = \frac{A_{i,t}^j (d_{ni,t}^j)^{-\theta^j}}{\sum_m A_{m,t}^j (d_{nm,t}^j)^{-\theta^j}}, \qquad j \in \mathcal{J}^T, \tag{2}$$

where the first subscript of $\pi_{ni,t}^j$ denotes the purchasing destination and the second the producing origin, and γ^j is the usual Fréchet price constant. Appendix A.2 gives the within-sector aggregation. For computation, I write the within-period gravity block in exact hat algebra around the observed 2022 expenditure shares, following Dekle et al. (2007); Appendix D gives the full system. Let $X_{n,t}^j$ denote purchaser expenditure in destination n. Gross producer revenue at origin i, inclusive of the origin subsidy and available for factor payments, is

$$Y_{i,t}^j = \frac{1}{1-s_{i,t}^j} \sum_n \frac{\pi_{ni,t}^j X_{n,t}^j}{1+\tau_{ni,t}^j}, \tag{3}$$

Destination spending is allocated across origins by $\pi^{j}_{ni,t}$. Thus $Y^{j}_{i,t}$ is nominal producer revenue, not the real equipment-output measure that enters learning; equation (5) defines the latter for the equipment sectors. Cost minimization determines domestic factor payments, while tariff revenue accrues to the household.

3.3 Renewable Investment and Local Energy

Renewable investment links local deployment to demand in the world equipment market. A developer in country n purchases $I^{r}_{n,t}$ units of a renewable-investment bundle that combines traded equipment with balance-of-system inputs purchased in the destination. Let $C^{\text{proj},r}_{n,t}$ denote the bundle's unit cost. Its unit cost and conditional expenditure demands are

$$C^{\text{proj},r}_{n,t} = B_{nr}\left(\frac{P^{M^r}_{n,t}}{\omega^r}\right)^{\omega^r}\left(\frac{P^{O}_{n,t}}{1-\omega^r}\right)^{1-\omega^r},$$
$$X^{M^r}_{n,t} = \omega^r C^{\text{proj},r}_{n,t} I^{r}_{n,t}, \qquad X^{B,r}_{n,t} = (1-\omega^r) C^{\text{proj},r}_{n,t} I^{r}_{n,t}. \tag{4}$$

Here $I^{r}_{n,t}$ is real gross investment measured in units of this bundle, so $C^{\text{proj},r}_{n,t} I^{r}_{n,t}$ is total renewable-project expenditure. The parameter B_{nr} is a destination–technology project-cost shifter, ω^r is equipment's share of project cost, and $X^{B,r}_{n,t}$ is balance-of-system expenditure in destination n. Equipment demand is allocated across origins by $\pi^{M^r}_{ni,t}$. Thus $\omega^r \pi^{M^r}_{ni,t}$ is the share of renewable-project expenditure spent on equipment sourced from origin i. Tariff-exclusive buyer payments to origin i are obtained by additionally dividing by $1+\tau^{M^r}_{ni,t}$; gross producer revenue also includes the subsidy payment in equation (3).

To define experience accumulation, let $\theta^r \equiv \theta^{M^r}$, $a^{r}_{i,t} \equiv (A^{M^r}_{i,t})^{1/\theta^r}$, and $\tilde{c}^{M,r}_{i,t} \equiv c^{M^r}_{i,t}/a^{r}_{i,t}$. Because $A^{M^r}_{i,t}(c^{M^r}_{i,t})^{-\theta^r} = (c^{M^r}_{i,t}/a^{r}_{i,t})^{-\theta^r}$ in the Fréchet trade block, $a^{r}_{i,t}$ is the origin's unit-productivity shifter and $\tilde{c}^{M,r}_{i,t}$ is its productivity-adjusted factory-gate unit cost. Origin real equipment output is producer revenue deflated by this cost:

$$y^{M,r}_{i,t} = \frac{Y^{M^r}_{i,t}}{\tilde{c}^{M,r}_{i,t}}. \tag{5}$$

This index measures factory-gate real output within technology r, after deflating nominal producer revenue by the origin's productivity-adjusted unit cost.

Local energy. Installed capacity produces nontraded electricity,

$$G^{r}_{n,t} = z^{r}_{n} K^{r}_{n,t}, \qquad r \in \mathcal{R},$$
$$r^{r}_{n,t} = P^{G,r,p}_{n,t} z^{r}_{n} - m^{r}_{n} P^{O}_{n,t}. \tag{6}$$

Here z^{r}_{n} is generation per unit of installed capacity, and $r^{r}_{n,t}$ is the net operating return to one unit of

that capacity. The energy-services aggregator pays $P^{G,r,u}_{n,t}$, the generator receives $P^{G,r,p}_{n,t}$, and $m^r_n P^O_{n,t}$ is operating cost per unit of capacity. The two generation prices coincide absent a renewable-generation production subsidy. Generation has constant returns to installed capacity.

Solar and wind first form a renewable composite,

$$R_{n,t} = \left[(\eta^S_n)^{1/\rho}(G^S_{n,t})^{(\rho-1)/\rho} + (\eta^W_n)^{1/\rho}(G^W_{n,t})^{(\rho-1)/\rho}\right]^{\rho/(\rho-1)}, \qquad \eta^S_n + \eta^W_n = 1, \tag{7}$$

$$\mathcal{E}_{n,t} = \left[\eta_n^{1/\psi} R^{(\psi-1)/\psi}_{n,t} + (1-\eta_n)^{1/\psi}(F^d_{n,t})^{(\psi-1)/\psi}\right]^{\psi/(\psi-1)}. \tag{8}$$

Here $F^d_{n,t}$ is fossil absorption in destination n, and $P^E_{n,t}$ is the dual price of energy services. The η^r_n's and η_n are CES distribution parameters; ρ governs solar–wind substitution and ψ renewable–fossil substitution, with the usual Cobb–Douglas limits when either elasticity equals one. Aggregate energy services clear $\mathcal{E}_{n,t} = C^E_{n,t} + E^O_{n,t} + \sum_r E^{M^r}_{n,t}$, where $E^O_{n,t}$ and $E^{M^r}_{n,t}$ are the energy inputs of the non-energy and equipment sectors and all users face the same composite price $P^E_{n,t}$. Household demand assigns expenditure share μ^E_n to energy, and producer energy demand follows from the unit-cost functions above. Appendix A.3 gives the dual prices and expenditure shares. This annual formulation does not model dispatch, storage, or transmission; Appendix B.6 discusses the implications of this aggregation.

3.4 Income and Market Clearing

The household owns domestic firms, the fixed fossil resource, and capital through a national investment fund. The fund receives non-energy capital income, renewable operating surplus, fossil-resource rent, and tariff revenue, net of investment expenditure. Let $\mathrm{Div}_{n,t}$ denote this net payout. The consolidated household budget is

$$P^C_{n,t} C_{n,t} = w_{n,t} L_n + \mathrm{Div}_{n,t} + \mathcal{T}_{n,t} - \Lambda_{n,t}. \tag{9}$$

Here $\mathcal{T}_{n,t}$ is an international transfer and $\Lambda_{n,t}$ the lump-sum tax that finances policy outlays. The household and fund take prices, aggregate experience, and the tax that balances the government budget as given.

Countries do not trade an international asset. Let $\overline{nx}_n$ be 2022 net imports divided by world labor income, positive for a deficit country. The fixed transfer $\mathcal{T}_{n,t} = \overline{nx}_n \sum_m w_{m,t} L_m$, with $\sum_n \overline{nx}_n = 0$, preserves the benchmark trade imbalance as a share of world labor income; setting $\overline{nx}_n = 0$ imposes balanced trade.

Definition 1 (Conditional within-period equilibrium)**.** *Given the inherited state $\mathcal{S}_t$, current investment flows $\{I^k_{n,t}\}_{n,k}$, and the current wedges and transfers, a conditional within-period equilibrium consists of prices, allocations, and bilateral expenditure shares such that producers and developers*

minimize costs; sourcing and origin revenue satisfy equations (2) and (3); project demand and local energy satisfy equations (4) and (6)–(8); household absorption satisfies equation (9); and goods, factor, energy, and fossil markets clear. Appendix D states the complete dividend, external-balance, absorption, and market-clearing equations.

3.5 Investment and Learning Dynamics

The within-period equilibrium maps deployment into equipment production across origins. The state then evolves through two laws of motion: investment augments productive capacity, while origin production augments the experience stocks that govern future comparative advantage. For each $k \in \mathcal{K} \equiv \{S, W, D\}$, capital evolves according to

$$K^k_{n,t+1} = (1-\delta_k)K^k_{n,t} + \Phi^k_n(I^k_{n,t}, K^k_{n,t}), \qquad \Phi^k_n(I,K) = \chi^k_n I^{\alpha} K^{1-\alpha}. \tag{10}$$

For $r \in \mathcal{R}$, $K^r_{n,t}$ denotes installed capacity. The corresponding non-energy variables are gross investment $I^D_{n,t}$ and capital $K^D_{n,t}$. Here δ_k is depreciation, χ^k_n is installation efficiency, and inherited capital is a predetermined input into installation, so that $\alpha < 1$ creates a technological adjustment friction. This is the constant-elasticity adjustment form used by Ravikumar et al. (2019); it applies to all three capital stocks, including the solar and wind capacity whose investment bundles contain internationally traded equipment. The national fund chooses all three investment flows from 2023 onward; observed 2022 investment enters only through the inherited stocks.

Real origin output from equation (5) builds experience according to

$$Q^r_{i,t+1} = \mu_{Q,r} Q^r_{i,t} + y^{M,r}_{i,t}, \qquad Q^r_{\text{glob},t} \equiv \sum_i Q^r_{i,t}, \qquad \mu_{Q,r} \in (0,1]. \tag{11}$$

The retention parameter $\mu_{Q,r}$ allows experience to depreciate; $\mu_{Q,r} = 1$ gives cumulative production. Equipment productivity combines world and origin-specific experience,

$$A^{M^r}_{i,t} = A^{M^r}_{i,2022} \left[\left(\frac{Q^r_{\text{glob},t}}{Q^r_{\text{glob},2022}} \right)^{1-\lambda^r_L} \left(\frac{Q^r_{i,t}}{Q^r_{i,2022}} \right)^{\lambda^r_L} \right]^{\nu^r_{\text{tot}}}. \tag{12}$$

The stock $Q^r_{i,2022}$ accumulates origin i's production through 2021, and $Q^r_{\text{glob},2022} = \sum_i Q^r_{i,2022}$. The first policy state is $Q^r_{i,2023} = \mu_{Q,r} Q^r_{i,2022} + y^{M,r}_{i,2022}$, so 2022 production enters once.

Here $\nu^r_{\text{tot}} \geq 0$ is the elasticity of the Fréchet productivity parameter with respect to the composite experience index, $g_r = \nu^r_{\text{tot}}/\theta^r$ is the corresponding price–experience elasticity, and $\lambda^r_L \in [0,1]$ is the origin-specific learning weight. World experience shifts the frontier shared by all origins; origin experience tilts relative productivity toward the producing country. Setting $\lambda^r_L = 0$ nests the world-experience specification of Arkolakis and Walsh (2023); positive λ^r_L adds an origin-specific

component that can change comparative advantage. As in the competitive learning-by-doing tradition of Krugman (1987) and Young (1991), individual producers take aggregate experience as given. Equation (12) is therefore a country–sector Marshallian externality.

Let $s^{I,r}_{n,t}$ denote the deployment subsidy. Define the after-subsidy acquisition price of one unit of technology-r renewable investment as $\widetilde{P}^{\text{inv},r}_{n,t} = (1 - s^{I,r}_{n,t}) C^{\text{proj},r}_{n,t}$. Let $\Phi^{k}_{I,n,t}$ and $\Phi^{k}_{K,n,t}$ denote the derivatives of the installation technology with respect to investment and inherited capital, evaluated at $(I^{k}_{n,t}, K^{k}_{n,t})$. The shadow value of one additional unit of installed renewable capacity is $q^{r}_{n,t} = \widetilde{P}^{\text{inv},r}_{n,t} / \Phi^{r}_{I,n,t}$. The household's one-period nominal discount factor is

$$M_{n,t+1} \equiv \beta \frac{P^{C}_{n,t} C_{n,t}}{P^{C}_{n,t+1} C_{n,t+1}}.$$

It converts a nominal payoff at $t+1$ into its value at t. The fund's renewable-capital Euler equation is

$$q^{r}_{n,t} = M_{n,t+1} \left[r^{r}_{n,t+1} + q^{r}_{n,t+1} \left(1 - \delta_{r} + \Phi^{r}_{K,n,t+1} \right) \right], \qquad r \in \mathcal{R}. \tag{13}$$

Non-energy investment is chosen on the same margin, with shadow value $q^{D}_{n,t} = P^{O}_{n,t} / \Phi^{D}_{I,n,t}$ and Euler equation

$$q^{D}_{n,t} = M_{n,t+1} \left[r^{D}_{n,t+1} + q^{D}_{n,t+1} \left(1 - \delta_{D} + \Phi^{D}_{K,n,t+1} \right) \right]. \tag{14}$$

The renewable rental r^{r} is given by equation (6), while r^{D} is the rental on non-energy capital. Each Euler equation equates the current shadow value of installed capital to the discounted value of next period's rental income and continuation value. Solar, wind, and non-energy investment use the same household discount factor and the same installation functional form, with country- and capital-specific parameters. Because investment is forward looking, future learning affects current investment through anticipated project prices and capital values. No separate learning-externality term appears in equation (13) because the fund does not internalize how its own investment changes aggregate experience and future equipment productivity. The usual transversality condition closes the infinite-horizon capital problem; Appendix D describes how the calculation approximates capital values beyond 2100.

3.6 Policies and Transition Equilibrium

The two subsidies act at different points in investment and sourcing. An origin equipment subsidy $s^{M,r}_{i,t}$ lowers the delivered cost of equipment produced in i and can therefore change sourcing wherever that equipment is available. It can raise production and origin-specific experience at i while reducing both at competing origins; equations (2), (5), and (11) map these sourcing changes into real output and experience at the subsidized and competing origins. A deployment subsidy, $s^{I,r}_{n,t}$, implemented as an investment tax credit (ITC), instead lowers the after-subsidy acquisition price of

renewable investment in n and expands equipment demand across origins according to the endogenous sourcing shares.[7] Country n's investment fund pays the after-subsidy investment expenditure, $(1 - s^{I,r}_{n,t}) C^{\text{proj},r}_{n,t} I^r_{n,t}$, while the government of the same country finances the subsidy. The contemporaneous fiscal outlays are

$$\Lambda^{I,r}_{n,t} = s^{I,r}_{n,t} C^{\text{proj},r}_{n,t} I^r_{n,t},$$
$$\Lambda^{M,r}_{i,t} = \frac{s^{M,r}_{i,t}}{1 - s^{M,r}_{i,t}} \sum_n \frac{\pi^{M^r}_{ni,t} X^{M^r}_{n,t}}{1 + \tau^{M^r}_{ni,t}}, \qquad \Lambda_{n,t} = \sum_r (\Lambda^{I,r}_{n,t} + \Lambda^{M,r}_{n,t}).$$

The first expression equals the subsidy rate times the pre-subsidy project-expenditure base, $C^{\text{proj},r}_{n,t} I^r_{n,t}$. In the second, the sum is the tariff-exclusive expenditure of buyers worldwide on technology-r equipment from origin i. This expenditure equals the buyer-financed share $1 - s^{M,r}_{i,t}$ of gross producer revenue, so multiplying it by $s^{M,r}_{i,t} / (1 - s^{M,r}_{i,t})$ yields the government's subsidy outlay. Finally, $\Lambda_{n,t}$ sums country n's outlays on deployment and equipment subsidies across solar and wind. These outlays are financed by the contemporaneous lump-sum tax in equation (9).

The two instruments can therefore generate different cross-country emissions responses. A deployment subsidy enters the fund's after-subsidy acquisition price directly only in the installing country, whereas an equipment subsidy changes delivered equipment costs in every destination that buys from the subsidized origin. Foreign emissions can respond through renewable-equipment prices and production learning as well as through the world fossil market. Section 4 formalizes the investment-cost channel, and Section 6.2 quantifies the full response.

Definition 2 (Transition equilibrium)**.** *Given the inherited state $\mathcal{S}_1$, policy paths, and a trade-cost and transfer regime, a perfect-foresight transition equilibrium is a sequence of conditional within-period equilibria, investment flows $\{I^k_{n,t}\}_{n,k}$, shadow values $\{q^k_{n,t}\}_{n,k}$, and states $\{\mathcal{S}_{t+1}\}_{t \geq 1}$ such that capital evolves according to equation (10); real equipment output, experience, and equipment productivity satisfy equations (5), (11), and (12); the national investment funds satisfy equations (13) and (14); and the fiscal outlays defined above are financed by the lump-sum tax in equation (9). Equipment producers take aggregate experience as given, and the transversality and terminal conditions described in Appendix D hold.*

The subsidy paths eventually return to zero, while regime-specific trade policy, other trade costs, and transfers remain in place. Appendix D.4 describes the solution method, and Appendix C.9 reports learning-stock and horizon checks.

[7] A renewable-generation production tax credit (PTC) can also be represented in this framework. It raises generator revenue rather than reducing installation cost and subsidizes inframarginal generation as well as generation induced by new investment. Appendix B.1 maps a PTC into an equivalent marginal investment incentive, Appendix C.4 compares finite and permanent ITC and PTC windows, and Appendix D records their distinct fiscal outlays.

3.7 Welfare and Carbon Accounting

Fossil absorption generates territorial CO_2 emissions,

$$Z_t = \sum_n \epsilon_n F^d_{n,t}. \tag{15}$$

Here ϵ_n is country n's territorial emissions factor per unit of fossil absorption. Carbon damages do not enter private utility, and carbon is unpriced in the benchmark and subsidy counterfactuals. I value policy-induced emissions changes ex post when reporting carbon-inclusive welfare; Appendix C.8 separately introduces a fossil-use tax. For policy p, country n's consumption-welfare change is

$$\Delta U^{\text{cons}}_n = \sum_{t=1}^{\infty} \beta^{t-1} \left(\log C^p_{n,t} - \log C^0_{n,t} \right). \tag{16}$$

Let ℓ^{cons}_n denote country n's corresponding log consumption-equivalent change, and let $\overline{\mathcal{E}}_{n,T}$ be country n's present-value baseline expenditure under the common post-2100 continuation used across all policy counterfactuals. Let SCC_t denote the social cost of carbon, or global marginal damage per unit of CO_2. The superscript R denotes aggregation with reference-expenditure weights. The global consumption-welfare index and its carbon-inclusive counterpart are

$$\begin{aligned} \ell^R_W &= \sum_n \omega_n \ell^{\text{cons}}_n, & \omega_n &= \frac{\overline{\mathcal{E}}_{n,T}}{\sum_m \overline{\mathcal{E}}_{m,T}}, \\ \ell^{R,SCC}_W &= \ell^R_W - \frac{1}{\overline{\mathcal{E}}_{W,T}} \sum_{t=1}^{T} \beta^{t-1} SCC_t \left(Z^p_t - Z^0_t \right), & \overline{\mathcal{E}}_{W,T} &= \sum_n \overline{\mathcal{E}}_{n,T}. \end{aligned} \tag{17}$$

The index aggregates country consumption welfare using fixed reference-expenditure weights; the Pigouvian benchmark in Section 4 instead uses a cooperative planner objective. Changing the SCC path changes the global carbon-inclusive index; alternative allocations of global damages across countries affect only the corresponding national accounts. Neither changes the equilibrium paths. Appendix A.4 explains the post-2100 welfare adjustment, the conversion to percentage consumption-equivalent variation, and the allocation of global damages across countries.

Section 4 first isolates how equipment trade shapes each subsidy's effects on investment costs and production learning, and then embeds those effects in a welfare envelope. The direct formulas do not sign the general-equilibrium emissions responses reported in the quantitative analysis.

4 International Incidence and Policy Assignment

This section asks how deployment and equipment subsidies should be assigned when carbon pricing is unavailable. Throughout this section, the analytical reference transition has no pre-existing

policy wedges, and all marginal derivatives are evaluated at that reference unless stated otherwise. I first state the full-instrument benchmark and then characterize marginal second-best policy returns. Fiscal normalization puts the two subsidies on a common scale; sourcing shares determine where their investment-cost and learning effects accrue; and return differences determine assignment across countries. Technology and time indices are suppressed; Appendix A.1 gives the full derivations.

4.1 First-Best Benchmark and Marginal Second-Best Choice

First-best benchmark. Let $\mathcal{B}_{i,t}^{L,r}$ be the discounted marginal value of the future world and origin-specific experience created by one additional unit of current equipment output; Appendix equation (A.10) gives its expression in model primitives. Economically, it is a discounted sum of future productivity gains: it is larger when learning is steeper, experience persists longer, and future equipment output is large relative to the relevant experience stock, with λ_L^r determining the weight on origin-specific rather than world experience.

Lemma 1 (Pigouvian implementation)**.** *Suppose international lump-sum transfers support the planner allocation and learning and carbon are the only externalities. A Pigouvian implementation satisfies*

$$s_{i,t}^{M,r,*}\tilde{c}_{i,t}^{M,r} = \mathcal{B}_{i,t}^{L,r}, \qquad t_{n,t}^{F,*} = \epsilon_n SCC_t, \qquad s_{n,t}^{I,r,*} = 0, \tag{18}$$

where SCC_t is the date-t numeraire value of one unit of emissions and $t_{n,t}^F$ is a specific tax on all fossil use.

The fossil tax prices emissions, while the origin-indexed equipment subsidy prices the future productivity created by current output, as in the familiar carbon-tax-plus-clean-innovation-subsidy result of Acemoglu et al. (2012) and Hémous (2016).[8] The paper studies the constrained case in which the fossil tax is unavailable. Deployment and equipment subsidies then affect emissions through the equilibrium responses of investment, fossil use, and equipment production, and affect learning through the production they induce across origins. Neither subsidy generally reproduces both the origin-by-technology learning wedge and a tax on all fossil use. Their second-best ranking therefore depends on policy location, the pattern of equipment trade, and the objective used to value foreign effects.

Marginal second-best returns. Let I_n denote a deployment-subsidy policy in destination n, and let M_i denote an equipment-subsidy policy at origin i. Each policy specifies its technology coverage and time path; write $z \in \{I_n, M_i\}$. Let f_z be its present-value fiscal outlay in the model numeraire. Because the subsidy base is positive at the reference, the statutory rate can be locally inverted, so

[8] Implementing the learning wedge with the model's ad valorem subsidy additionally requires $0 \leq \mathcal{B}_{i,t}^{L,r}/\tilde{c}_{i,t}^{M,r} < 1$. Outside this range, $\mathcal{B}_{i,t}^{L,r}$ remains the required specific wedge but cannot be implemented with an ad valorem rate satisfying $1 - s_{i,t}^{M,r} > 0$.

f_z can parameterize policy changes around it. Fiscal normalization places policies with different statutory rates and subsidy bases on a common marginal scale.

Local variation in f_z scales the specified policy path while holding its timing and technology coverage fixed. For differentiation, f_z is extended as a signed local coordinate around zero; the subsidy problem below imposes $f_z \geq 0$.

I use three objectives. Let $\mathcal{W}^P$ denote the cooperative planner's carbon-inclusive objective underlying Lemma 1, normalized by its date-1 numeraire multiplier; let $\mathcal{W}^G = \ell_W^R$ and $\mathcal{W}^{N_n} = \ell_n^{\text{cons}}$, matching the log consumption-equivalent indices in Section 3.7. Unless stated otherwise, G and N_n exclude carbon; carbon-inclusive results append the damage account in equation (17).

The marginal return to policy z, per unit of fiscal outlay, can be written

$$R_z^o \equiv \left.\frac{\partial \mathcal{W}^o}{\partial f_z}\right|_0 = T_z^o + C_z^o + \sum_{r \in \mathcal{R}} \left[(1-\lambda_L^r) L_{z,W}^{o,r} + \lambda_L^r L_{z,O}^{o,r}\right]. \tag{19}$$

Equation (19) decomposes the return into four components. T_z^o is the dynamic terms-of-trade and transfer margin, while C_z^o values the induced emissions path. The two learning terms value future productivity created through world and origin-specific experience; their λ_L^r weights determine how strongly production location matters. All four terms are equilibrium directional derivatives.

At the reference, the marginal subsidy payment is offset by the lump-sum tax financing it, while private reoptimization vanishes by the envelope theorem. Hence R_z^o contains no separate -1 fiscal-cost term. International-price transfers cancel for the cooperative planner, so $T_z^P = 0$, but need not cancel under the global or national consumption objectives. The main consumption-welfare tables set $C_z^G = C_z^{N_n} = 0$ and report carbon separately; national learning components count only the productivity gains accruing to country n. Appendix A.1 provides the formal derivation.

Lemma 2 (Marginal second-best approximation)**.** *If $\mathcal{W}^o$ is twice differentiable in the fiscal coordinates, then*

$$\Delta \mathcal{W}^o(\mathbf{f}) = \mathbf{R}_o' \mathbf{f} - \frac{1}{2} \mathbf{f}' H_o \mathbf{f} + o(\|\mathbf{f}\|^2), \qquad H_o \equiv -\nabla_f^2 \mathcal{W}^o(\mathbf{0}). \tag{20}$$

If H_o is positive definite, the unconstrained maximizer of the quadratic approximation is $\mathbf{f}_{o,\text{quad}}^* = H_o^{-1}\mathbf{R}_o$. With nonnegative subsidies, its constrained maximizer satisfies

$$\mathbf{f}_{o,\text{quad}}^* \geq 0, \qquad H_o \mathbf{f}_{o,\text{quad}}^* - \mathbf{R}_o \geq 0,$$
$$f_{z,o,\text{quad}}^* (H_o \mathbf{f}_{o,\text{quad}}^* - \mathbf{R}_o)_z = 0 \quad \textit{for every } z.$$

If only policy z is available and $h_z^o = (H_o)_{zz} > 0$,

$$f_{z,o,\text{quad}}^* = \max\left\{0, \frac{R_z^o}{h_z^o}\right\}, \qquad \Delta \mathcal{W}_{z,o,\text{quad}}^* = \frac{[\max\{R_z^o, 0\}]^2}{2h_z^o}. \tag{21}$$

The lemma separates ranking from scale. For sufficiently small common-budget changes, R_z^o

ranks policies; H_o determines preferred scale and interactions. This maps directly to the quantitative exercises: common-budget comparisons fix scale, unilateral profiles vary it, and joint-policy exercises allow interactions. All finite policies are solved from the nonlinear transition.

4.2 How Equipment Trade Routes the Two Subsidies

I first derive how marginal changes in the two statutory subsidy rates—and, after fiscal normalization, the associated marginal fiscal outlays—change renewable-project costs and production experience across countries. For one technology and date,[9] let $\widetilde{P}_d^{\text{inv}} = (1 - s_d^I)C_d^{\text{proj}}$ be the acquisition price paid by investors, ω equipment's project-cost share, and π_{di} destination d's sourcing share from origin i. Holding current origin costs, productivities, and trade costs fixed, the direct price effects are

$$-\left.\frac{\partial \log \widetilde{P}_d^{\text{inv}}}{\partial s_n^I}\right|_0 = \mathbf{1}\{d = n\}, \qquad -\left.\frac{\partial \log \widetilde{P}_d^{\text{inv}}}{\partial s_i^M}\right|_0 = \omega \pi_{di}. \tag{22}$$

Let $V_{di}^M \equiv \pi_{di} X_d^M$ denote destination d's expenditure on equipment from origin i. Let J_n be total project spending in destination n, and let $S_i = \sum_d V_{di}^M$ be origin i's total equipment sales. Since $\partial \Lambda_n^I / \partial s_n^I|_0 = J_n$ and $\partial \Lambda_i^M / \partial s_i^M|_0 = S_i$, the same effects per marginal fiscal dollar are

$$-\left.\frac{\partial \log \widetilde{P}_d^{\text{inv}}}{\partial \Lambda_n^I}\right|_0 = \frac{\mathbf{1}\{d = n\}}{J_n}, \qquad -\left.\frac{\partial \log \widetilde{P}_d^{\text{inv}}}{\partial \Lambda_i^M}\right|_0 = \frac{\omega \pi_{di}}{S_i}. \tag{23}$$

The deployment subsidy therefore changes acquisition cost directly only in the subsidizing destination. The equipment subsidy affects every destination that buys from the subsidized origin. Sourcing shares govern the magnitude and geographic distribution of foreign price pass-through.

For the remainder of this section, $k \in \{W, O\}$ indexes the experience component, with W denoting world and O origin-specific experience. Let $v_i^{k,o}$ be the discounted learning value generated by one dollar of equipment sales at origin i. Appendix A.1 gives its expression in model primitives using the continuation value per unit of real output. Its magnitude depends on the objective: the planner, the reference-weighted global index, and a national government value the same productivity path differently. Suppress o below, and let $\varsigma_{ni} = V_{ni}^M / S_i$ denote destination n's share in origin i's sales.

The two direct learning measures below use different denominators. Define the deployment equipment-spending multiplier—the increase in own-destination equipment spending per dollar of deployment-subsidy outlay—as

$$\xi_n^I \equiv \left.\frac{\partial X_n^M / \partial s_n^I}{\partial \Lambda_n^I / \partial s_n^I}\right|_0.$$

[9] Appendix A.1 marks reference-path levels with an overbar; the bars are suppressed here. The full policy return in equation (19) discounts and aggregates the date- and technology-specific formulas.

Multiplying the deployment routing value by τ_n^I expresses its own-destination contribution per fiscal dollar. The multiplier ξ_n^I captures the equipment-spending response in the subsidizing destination; the sourcing shares π_{ni} then allocate those purchases across origins.

Proposition 1 (Direct learning effects of the two subsidies)**.** *For $k \in \{W, O\}$, the deployment routing value per additional dollar of equipment purchases in destination n is*

$$\mathcal{V}_n^{I,k} = \sum_i \pi_{ni} v_i^k. \tag{24}$$

Its own-destination contribution per marginal dollar of deployment-subsidy outlay is $\xi_n^I \mathcal{V}_n^{I,k}$. By contrast, the equipment measure below is already normalized per marginal dollar of equipment-subsidy outlay. At fixed current destination spending, origin costs, and productivities, it is

$$\mathcal{V}_i^{M,k} = \underbrace{v_i^k}_{\text{additional origin output}} + \underbrace{\theta \sum_n \varsigma_{ni} \sum_{h \neq i} \pi_{nh} (v_i^k - v_h^k)}_{\text{production reallocation}}. \tag{25}$$

Combining the world and origin-specific components gives

$$\mathcal{V}_n^I = (1 - \lambda_L)\mathcal{V}_n^{I,W} + \lambda_L \mathcal{V}_n^{I,O}, \qquad \mathcal{V}_i^M = (1 - \lambda_L)\mathcal{V}_i^{M,W} + \lambda_L \mathcal{V}_i^{M,O}.$$

The deployment subsidy inherits the destination's sourcing mix: its learning value is a sourcing-weighted average of supplier learning values. The equipment subsidy combines additional production at the subsidized origin with the production it displaces elsewhere. Reallocation raises the direct learning value when the displaced suppliers have lower learning values than the subsidized origin, and lowers it in the reverse case. Because the equipment measure is normalized by fiscal outlay, a large sales base alone does not imply a high learning return. Neither export scale nor hub status alone therefore signs this direct effect.

These are direct effects, not general-equilibrium welfare derivatives: they hold current spending, origin costs, and productivities fixed. After aggregation over dates and technologies, these measures supply the routing components of the L terms in equation (19); the full L terms also include indirect equilibrium responses in spending, costs, and productivity. Appendix A.1 embeds them in the dynamic welfare envelope and gives the technology- and date-specific sourcing formulas.

Corollary 1 (Direct investment-cost pass-through and abatement)**.** *Suppose a lower renewable-investment price weakly raises discounted abatement through the direct investment-cost channel. Let $\mathcal{A}_d^o \geq 0$ be the welfare value under objective o of the abatement generated by a one-unit decline in destination d's log renewable-investment price, holding the other general-equilibrium channels fixed. With J_n and S_i the project-spending and equipment-sales totals defined for equation (23), the*

corresponding direct carbon components per marginal fiscal dollar are

$$C_{I_n}^{o,\mathrm{dir}} = \frac{\mathcal{A}_n^o}{J_n}, \qquad C_{M_i}^{o,\mathrm{dir}} = \frac{\omega}{S_i}\sum_d \pi_{di}\mathcal{A}_d^o. \tag{26}$$

Thus the deployment subsidy generates direct abatement only in the subsidizing destination. The equipment subsidy can transmit direct abatement to every destination that sources from the subsidized origin. These expressions sign the direct renewable-investment-cost channel, not the general-equilibrium response of territorial emissions.

Section 6 solves the additional spending, sourcing, supplier-cost, fossil-price, income, and scale responses in the nonlinear model.[10]

Equipment sourcing governs where production and learning occur; because renewable electricity is local, the destination of investment governs where new capacity directly displaces fossil generation. Electricity trade would add a second network linking generation to emissions, without changing the equipment-sourcing formulas.

4.3 Global and Origin-Specific Learning, Equipment Trade, and Policy Assignment

Under world experience, the direct learning value depends on additional total output. Under origin-specific experience, it also depends on where that output is produced. At a fixed reference transition, shifting weight toward origin-specific experience raises a policy's direct learning value exactly when its origin component exceeds its world component. This comparison holds continuation values fixed; recalibrating λ_L and resolving the equilibrium is a separate counterfactual.

Corollary 2 (Learning scope and equipment trade)**.** *Suppose first that $\lambda_L = 0$ and productivity-adjusted unit costs are common across active origins. Then $v_i^W = v^W$, and the two displayed learning measures coincide numerically:*

$$\mathcal{V}_n^{I,W} = \mathcal{V}_i^{M,W} = v^W \qquad \text{for every } n \text{ with } J_n > 0 \text{ and } i \text{ with } S_i > 0. \tag{27}$$

Production reallocation then has no direct effect on learning through the composition of output across origins; only total real output matters. The fiscal-normalized own-destination deployment contribution is $\xi_n^I v^W$, so equation (27) does not equate full policy returns. If instead equipment is not traded,

$$\pi_{ni} = \mathbf{1}\{n = i\} \quad \Longrightarrow \quad \mathcal{V}_n^{I,k} = v_n^k, \qquad \mathcal{V}_i^{M,k} = v_i^k, \qquad k \in \{W, O\}, \tag{28}$$

[10] The direct-versus-total distinction is related to the reverse-leakage mechanism in Kotchen and Maggi (2025). Their green subsidy supports the output of a traded green-energy input. Here deployment and equipment production are separate margins, and equipment trade and production learning make foreign investment costs and emissions depend on the instrument and location.

and an equipment subsidy has no direct foreign investment-cost pass-through.

With origin-specific experience, equal world output need not imply equal future productivity. At a fixed reference path, the origin-specific learning value v_i^O rises with discounted future sales and falls with inherited experience. Equipment trade then adds two distinct margins: deployment-induced demand can create experience abroad, and an origin-indexed equipment subsidy can displace production at competing suppliers while lowering project costs in foreign destinations. Both cross-border equipment-trade margins disappear in the no-trade limit, but world-experience spillovers and fossil-fuel trade remain. Even under pure world experience, equipment trade can matter through delivered equipment costs, carbon, and international-price incidence.

Which location should receive deployment support, and which should receive equipment support? Let U denote an importing market and H a production hub, and fix an objective o. The returns $R^o_{I_n}$ and $R^o_{M_n}$ are defined per common marginal fiscal dollar in equation (19). Define location n's instrument return gap, measured as deployment minus equipment, by

$$\begin{aligned}\Delta^o_n &\equiv R^o_{I_n} - R^o_{M_n} \\ &= \Delta T^o_n + \Delta C^o_n + \sum_{r\in\mathcal{R}} \left[(1-\lambda^r_L)\Delta L^{o,r}_{n,W} + \lambda^r_L \Delta L^{o,r}_{n,O}\right],\end{aligned} \tag{29}$$

where $\Delta T^o_n = T^o_{I_n} - T^o_{M_n}$, $\Delta C^o_n = C^o_{I_n} - C^o_{M_n}$, and $\Delta L^{o,r}_{n,k} = L^{o,r}_{I_n,k} - L^{o,r}_{M_n,k}$. The cooperative carbon-inclusive assignment uses $o = P$; the noncarbon global tables below use $o = G$. Proposition 2 uses these gaps to state the assignment condition.

Proposition 2 (Policy assignment by return gaps)**.** *Suppose there are two locations, U and H, and the same infinitesimal fiscal outlay is implemented at each, with the instrument chosen between deployment and equipment subsidies. Among the two assignments that use one subsidy of each type, assigning the deployment subsidy to U and the equipment subsidy to H has the larger first-order welfare effect if and only if*

$$R^o_{I_U} + R^o_{M_H} > R^o_{M_U} + R^o_{I_H}, \qquad \textit{equivalently} \qquad \Delta^o_U > \Delta^o_H. \tag{30}$$

The stronger conditions $\Delta^o_U > 0 > \Delta^o_H$ are necessary and sufficient for (I_U, M_H) to dominate each of the other three assignments, (M_U, I_H), (I_U, I_H), and (M_U, M_H), at first order. If zero subsidy is also feasible, $R^o_{I_U} > 0$ and $R^o_{M_H} > 0$ are additionally required for the selected pair to dominate allocations that omit either policy.[11]

What determines the return gaps? The assignment condition is comparative: deployment need not dominate equipment at the importer in levels; its deployment advantage need only exceed the hub's.

[11] The assignment rule extends to multiple locations; Appendix A.1 gives the corresponding sorting result.

Together with the deployment multiplier ξ_n^I, equations (24), (25), and (26) map the routing primitives into the direct learning and carbon components of Δ_n^o. Deployment's own-destination learning contribution per fiscal dollar, $\xi_n^I \sum_i \pi_{ni} v_i^k$, is larger when the subsidy induces more equipment spending and when sourcing is tilted toward suppliers with high learning values. For equipment support, ς_{di} weights the destinations in which rival production is displaced. Conditional on v^k, π, and ς, θ scales this reallocation, whose sign depends on whether the subsidized origin displaces lower- or higher-value suppliers. Direct price exposure in destination d is governed by $\omega \pi_{di}/S_i$, while the direct abatement term additionally weights that exposure by $\mathcal{A}_d^o$.

At a fixed reference transition and fixed component values, the origin-specific weight changes the learning portion of the return gap according to

$$\frac{\partial}{\partial \lambda_L^r}\left[(1-\lambda_L^r)\Delta L_{n,W}^{o,r} + \lambda_L^r \Delta L_{n,O}^{o,r}\right] = \Delta L_{n,O}^{o,r} - \Delta L_{n,W}^{o,r}.$$

Thus greater origin specificity has no uniform effect on the instrument ranking. The full gap also includes international-price incidence and endogenous equilibrium responses, so neither importer status nor a large production base is sufficient by itself.

National versus global objectives. Evaluating the same return gaps under national objectives can reverse the globally preferred mixed assignment. Corollary 3 gives the corresponding conditions.

Corollary 3 (National–global misalignment)**.** *Consider the two-location assignment of Proposition 2 and suppose each location's instrument is chosen to maximize that location's own national return. Because first-order returns are additive across locations, each choice is independent of the instrument assigned at the other location. If, at each location, the national return favors the instrument opposite to that in the globally preferred mixed assignment,*

$$R_{M_U}^{N_U} > R_{I_U}^{N_U}, \qquad R_{I_H}^{N_H} > R_{M_H}^{N_H}, \qquad \text{and} \qquad R_{I_U}^{G} - R_{M_U}^{G} > R_{I_H}^{G} - R_{M_H}^{G}, \tag{31}$$

then own-return maximization selects (M_U, I_H)*, the assignment that condition* (30) *ranks below the reverse pair* (I_U, M_H) *under the global objective at first order.*

These results turn policy assignment into a comparison of location-specific return gaps. Equipment trade shapes those gaps by routing direct price, learning, and abatement incidence across borders; because national objectives count only the incidence captured domestically, their preferred mixed assignment can reverse the global ordering of the two mixed assignments. Trade position alone does not sign the gaps. Section 6 evaluates these mechanisms and assignment comparisons in the nonlinear model; Table 5 and Appendix Table C.4 provide numerical counterparts to the pass-through and marginal-return formulas.

5 Calibration

I calibrate the objects identified in Section 4 around a 2022 benchmark and retain historical tariff data for the applications with observed policy wedges. Input–output accounts, capacities, generation, labor, and emissions supply observed anchors; balanced equipment-expenditure matrices and bilateral tariff equivalents are constructed from source data. Factor and expenditure shares, installation efficiencies, and initial states are mapped internally; structural elasticities are externally sourced, constructed, or imposed as summarized in Table 1; and two solar-learning moments are estimated.

5.1 Countries, Sectors, and Benchmark Trade

The model contains 15 countries and regions: the United States, China, Canada, Mexico, Japan, Korea, India, and Russia, together with Europe including the U.K., Australia–New Zealand, Southeast Asia, South America, Africa, the Middle East, and the rest of the world. Appendix B.2 gives the aggregation. The four traded composites are solar equipment, wind equipment, fossil fuel, and a pooled non-energy good; solar and wind generation are local.

The 2022 OECD inter-country input–output accounts determine bilateral fossil and non-energy flows, sectoral output and absorption, value-added shares, and investment. Separate matrices determine solar- and wind-equipment expenditures. The solar matrix balances Comtrade cell and module flows to independent production and absorption margins. Wind sourcing is imputed by combining trade in generating sets and selected parts with installation-based absorption margins. The resulting equipment expenditures are $139.6 billion for solar and $58.0 billion for wind; China supplies 80.1 percent of solar equipment, while the imputed wind matrix assigns China 59.8 percent of wind supply. The solar matrix is also used as the sourcing proxy for the traded solar components included in the equipment share of installation cost. Appendix B.3 gives the construction and product-boundary assumptions.

The historical data include observed trade policy. Sourcing shares embed the full trade-cost wedge, while the tariff matrix records its policy component separately. The central China-to-U.S. solar wedge is a 70.4 percent tariff equivalent; other material bilateral rates and their construction are listed in Appendix B.3. The central assignment exercise sets this policy component to zero while preserving calibrated productivity and other trade-cost fundamentals; sourcing then adjusts endogenously. The IRA and source-restriction applications restore the observed tariff matrix. Whenever tariffs are active, purchaser expenditure includes tariffs, tariff revenue accrues to the importer, and origin factor revenue is tariff-exclusive. Fixed transfers preserve the 2022 external positions and sum to zero.

Trade elasticities govern supplier reallocation within each composite. I set $(\theta^S, \theta^W, \theta^F, \theta^O) = (5.38, 4.38, 3.16, 6.43)$. The first three values are estimates from Fontagné et al. (2022) mapped to the corresponding ICIO sectors; θ^O is the baseline-expenditure-weighted aggregation of the component

non-energy sectors.

5.2 Equipment Learning

Two moments parameterize learning for each technology. A world price–experience curve disciplines $g_r = \nu^r_{\text{tot}}/\theta^r$, while an exporter-capability regression disciplines $b_r = \lambda^r_L \nu^r_{\text{tot}}$. I estimate both moments for solar. For wind, the shorter price–experience series provides descriptive evidence on total learning, while the benchmark total-learning elasticity and origin-specific learning weight are calibrated from external evidence. The mapping into the model is

$$\nu^r_{\text{tot}} = \theta^r g_r, \qquad \lambda^r_L = \frac{b_r}{\theta^r g_r}.$$

For solar, I estimate

$$\log P^{S,\text{real}}_{W,t} = a - g_S \log Q^{S,\text{beg}}_{\text{glob},t} + \beta_P \log P^{\text{poly,real}}_t + u_t,$$

where $P^{S,\text{real}}_{W,t}$ is the real world module price, $Q^{S,\text{beg}}_{\text{glob},t}$ is the retained world experience stock at the beginning of the year, and $P^{\text{poly,real}}_t$ is the real polysilicon price. On the benchmark $\mu_Q = 0.99$ stock, the 2000–24 regression gives $\hat{g}_S = 0.465$ (Newey–West s.e. 0.023), implying a 27.5 percent price decline per doubling.

To estimate the origin-specific component, I first recover exporter-year capability $\hat{S}^S_{n,t}$ from a PPML gravity equation with importer-year and country-pair fixed effects. This capability is a gravity-based index of relative exporter competitiveness: in the model, it rises with origin productivity but may also reflect time-varying origin costs and policy wedges. I then estimate

$$\log \hat{S}^S_{n,t} = b_S \log Q^{S,\text{beg}}_{n,t} + \delta_n + \delta_t + \varepsilon_{n,t},$$

where δ_n and δ_t are origin and year fixed effects. Using module production for 22 origins over 2005–23 gives $\hat{b}_S = 0.308$ (country-clustered s.e. 0.092). At $\theta^S = 5.38$, the two estimates map into $(\nu^S_{\text{tot}}, \lambda^S_L) = (2.50, 0.12)$.

These regressions discipline the calibration rather than identify a causal learning effect: cumulative production also covaries with R&D, scale, input prices, and other experience-linked innovation. Still, a corroborating specification that instruments cumulative output with exporters' predetermined exposure to foreign feed-in-tariff expansions gives a coefficient of 0.313 (s.e. 0.119), close to the benchmark estimate of 0.308, with a first-stage F-statistic of 19.8. Appendix B.5 reports the estimates, data construction, sensitivity to stock timing, initial experience, sample windows, and the policy-based instrument; Appendix Table B.3 summarizes the mapped learning parameters. Appendix C.2 reports joint low- and high-learning-parameter scenarios, counterfactuals that shut down

individual learning mechanisms, and technology-specific marginal policy returns.

The descriptive wind fit in Figure 2 implies a 21 percent price decline per doubling of cumulative installed capacity. The price history is short, however, and installed capacity is an imperfect proxy for manufacturing experience, so the data do not support the same two-moment decomposition with comparable precision. The analogous manufacturer-location regression gives an own-country coefficient near 0.6 but uses only six producer clusters and reverses sign when China is removed. Covert and Sweeney (2024) find that experience at other wind-turbine firms receives only 0.1–0.6 percent of the marginal weight on own-model experience, although these firm-level estimates do not map directly into a country-level weight. I therefore calibrate wind learning rather than estimate it. Conditional on $\nu^W_{\text{tot}} = 1$, I conservatively round the noisy own-country coefficient near 0.6 to $\lambda^W_L = 0.50$ and bracket it with 0.25 and 0.75. The resulting benchmark implies a 14.6 percent price decline per doubling, within the empirical range (Wiser et al., 2024).

I initialize country experience stocks from historical equipment production—module output from the annual IEA PVPS *Trends* reports for solar (IEA PVPS, 2024) and manufacturer-location data for wind—and scale them to match the model's 2022 equipment-output units. Following Arkolakis and Walsh (2023), the benchmark assumes one percent annual experience obsolescence, $\mu_Q = 0.99$, and constructs beginning-of-2023 stocks so that 2022 output enters once. The cumulative-output case $\mu_Q = 1$ re-estimates the solar moments and reconstructs the stocks; no-learning and no-global-learning exercises leave the 2022 trade and capital anchors unchanged.

5.3 Renewable Investment and the Energy Block

IRENA capacities for 2021 and 2022 determine beginning- and end-of-2022 solar and wind stocks. Dividing Ember's 2022 generation by end-of-2022 capacity gives the country–technology coefficient z^r_n in equation (6). The benchmark assumes constant returns to installed capacity, $G^r_n = z^r_n K^r_n$. This specification uses observed country–technology generation per unit of capacity to capture average utilization differences and leaves within-country resource congestion outside the benchmark.

The inner solar–wind aggregator is Cobb–Douglas, $\rho = 1$.[12] The outer renewable–fossil elasticity is $\psi = 3$, close to the estimate in Jo (2025) and within the ranges in Papageorgiou et al. (2017) and Wiskich (2023); Appendix C.3 considers $\psi \in \{2, 4\}$.

Total benchmark energy expenditure is the sum of fossil expenditure $X^F_{n,2022}$ and renewable-electricity expenditure $E^R_{n,2022}$; it is divided between households and production. Household energy expenditure is the residual

$$X^{E,H}_{n,2022} = X^F_{n,2022} + E^R_{n,2022} - \zeta^{E,O}_n Y^O_{n,2022} - \sum_{r \in \mathcal{R}} \zeta^{E,r}_n Y^{M^r}_{n,2022}.$$

[12] Appendix B.4 reports $\hat{\rho} = 0.645$ (s.e. 0.554). The estimate is too imprecise to distinguish it from the Cobb–Douglas benchmark and is not used as a structural estimate.

Here $\zeta_n^{E,O}$ and $\zeta_n^{E,r}$ are the energy-input expenditure shares of non-energy and equipment production. The household energy weight and the renewable share of the common energy composite are then

$$\mu_n^E = \frac{X_{n,2022}^{E,H}}{P_{n,2022}^C C_{n,2022}}, \qquad \omega_n^R = \frac{E_{n,2022}^R}{X_{n,2022}^F + E_{n,2022}^R}.$$

The production energy shares are constructed from intermediate-input expenditures in the OECD input–output accounts. For equipment, I use the pooled electrical-equipment and machinery parent industries as a proxy for solar- and wind-equipment producers. The energy shares count direct fossil inputs and the renewable share of domestic electricity purchases; other purchased electricity remains in the pooled non-energy composite. Households, non-energy producers, and equipment producers purchase the same local energy composite, which is why ω_n^R is calibrated from total energy expenditure. Across regions, the household energy weight averages 2.4 percent and ranges from 0.4 to 5.4 percent. Solar's share of renewable expenditure is based on generator revenue rather than the raw TWh split. Generator revenue is constructed to cover the user cost implied by the 2021–22 capital transition and annual fixed O&M of $23 per kW for solar and $43 per kW for wind (National Renewable Energy Laboratory, 2024). With benchmark energy prices normalized to one, these expenditure shares map directly into the CES distribution parameters. Appendix A.3 gives the mapping.

Traded equipment accounts for $\omega^S = 0.40$ of solar installation cost and $\omega^W = 0.55$ of wind installation cost. The solar share reflects modules, inverters, and electrical balance-of-system in the NREL utility-scale benchmark (Ramasamy et al., 2025); the wind share is the ex-works turbine share in the land-based cost benchmark of Stehly et al. (2024). The remaining installation spending is local non-energy input.

5.4 Initial States, Transition Parameters, and Carbon Reporting

Renewable installation efficiency χ_n^r is chosen so the capital law maps observed 2021 capacity and matrix-implied 2022 renewable-investment spending into observed end-of-2022 capacity. Appendix B.7 describes the regularization used for three region–wind observations with near-zero matrix-implied spending. Observed 2022 non-energy investment is gross fixed capital formation net of renewable-investment spending; its initial stock and installation scale make the 2022 investment rate stationary under the same capital law. I set $\alpha = 0.76$, following Ravikumar et al. (2019), and $(\delta_S, \delta_W, \delta_D) = (0.03, 0.03, 0.07)$, following Arkolakis and Walsh (2023). The non-energy capital share in value added is 0.30 (Greenwood et al., 1997).

Motivated by the 0.5–1 oil-field extraction-cost calibration in Kortum and Weisbach (2021), I set the fossil-supply elasticity to $\varepsilon_F = 1$ and use the competitive mapping $\eta_F = \varepsilon_F/(1+\varepsilon_F) = 0.5$ for the labor-payment share, as required by the exact envelope derivation in Appendix A.1. Appendix C.3 varies ε_F over $\{0.5, 2\}$ while holding $\eta_F = 0.5$, and hence the benchmark income split,

fixed; those reduced-form cases do not use the exact envelope decomposition. The supply elasticity governs aggregate extraction, while the fossil trade elasticity governs sourcing conditional on origin prices.

End-of-2022 capital and beginning-of-2023 experience are inherited states, and 2023 is the first endogenous investment year. The explicit transition runs through 2100; Appendix C.9 reports the learning-stock and horizon exercises.

The 2022 emissions account is anchored to 37.1 $GtCO_2$ of fossil and industry emissions from the Global Carbon Project (Friedlingstein et al., 2026). I aggregate those emissions to the model regions and set ϵ_n equal to regional emissions divided by benchmark fossil absorption. I value counterfactual changes at a flat $210 per metric ton in 2022 dollars. This benchmark converts the $185 central mean SC–$CO_2$ estimate in 2020 dollars from Rennert et al. (2022), computed at a 2 percent near-term risk-free discount rate, to 2022 dollars and rounds it. Using an SC–CO_2 value matches the model's physical emissions measure rather than folding other greenhouse gases into a CO_2-equivalent price. Appendices B.8 and B.9 describe the physical account and its ex-post valuation; Appendix C.8 separately imposes the same value as a specific fossil-use tax, rebates its revenue, and re-solves the equilibrium.

Table 1: Key Calibration Objects

Object	Benchmark value	Status	Source or identifying moment
Trade elasticities θ^j	$(5.38, 4.38, 3.16, 6.43)$ for solar equipment, wind equipment, fossil, and non-energy	External; constructed	Fontagné et al. (2022) estimates mapped to ICIO sectors; expenditure-weighted aggregation for non-energy
Observed tariff matrix $\tau^{M^r}_{ni,2022}$	Bilateral; China–U.S. solar 70.4 percent	Constructed benchmark input	Source-specific policy rates; zero tariffs define a counterfactual regime, not a recalibration
Solar learning (g_S, b_S)	$(0.465, 0.308)$	Estimated	World price curve and exporter capability under $\mu_Q = .99$ retention
Solar $(\nu^S_{\text{tot}}, \lambda^S_L)$	$(2.501, 0.123)$	Mapped	$\nu = \theta g$, $\lambda = b/\nu$
Wind $(\nu^W_{\text{tot}}, \lambda^W_L)$	$(1.00, 0.50)$	External scenario	Literature range and manufacturer-location evidence
Experience retention μ_Q	0.99; 1 in robustness	External	Arkolakis and Walsh (2023); one-percent annual obsolescence with matched estimation and state construction
Energy elasticities (ρ, ψ)	$(1, 3)$	Imposed/external	Cobb–Douglas inner nest; $\psi = 3$ follows Jo (2025); sensitivity reported for $\psi \in \{2, 4\}$
Generation returns	$G^r_n = z^r_n K^r_n$	Maintained specification	Constant returns in the benchmark
Equipment shares in installation cost (ω^S, ω^W)	$(0.40, 0.55)$	External	NREL solar and CoWER wind cost decompositions
Installation curvature α	0.76	External; maintained	Ravikumar et al. (2019)
Depreciation $(\delta_S, \delta_W, \delta_D)$	$(0.03, 0.03, 0.07)$	External	Arkolakis and Walsh (2023)
Fossil supply (ε_F, η_F)	$(1, 0.5)$	Imposed/derived	ε_F imposed and η_F implied by the competitive benchmark; sensitivity holds η_F fixed
Fixed O&M (m_S, m_W)	$(23, 43)$ dollars/kW-year	Externally set	NREL Annual Technology Baseline (National Renewable Energy Laboratory, 2024)
Factor-cost and model energy-expenditure shares	Region-sector specific	Calibrated	2022 input–output accounts and energy-demand closure
Generation z^r_n, installation χ^r_n	Region-technology specific	Mapped; calibrated	2022 TWh/capacity and the 2021–22 capacity boundary
Non-energy (K^D_n, χ^D_n)	Region specific	Calibrated closure	Common adjustment technology; $K^D_{n,0} = I^D_{n,0}/(\alpha\delta_D)$ and $\chi^D_n = \delta_D(K^D_{n,0}/I^D_{n,0})^\alpha$
Discount factor β	0.96 annually	Imposed	Annual model convention
Carbon value and tax	\$210/tCO$_2$; zero tax in reference and subsidy paths	External; policy use	Flat SC–CO$_2$ from Rennert et al. (2022); imposed as a tax only in Appendix C.8
Emissions intensity ϵ_n	Region specific; world total 37.1 GtCO$_2$	Mapped	Global Carbon Project fossil and industry emissions (Friedlingstein et al., 2026) divided by benchmark fossil absorption

Table 2: Calibration Fit and Non-Targeted Outcomes

Measure	Year	Model	Data
Panel A. Balanced calibration matrices versus source data			
China share of world solar-equipment output (percent)	2022	80.1	80.1
China share of world wind-equipment output (percent)	2022	59.8	59.8
Southeast Asia and Korea share of U.S. solar-equipment imports (percent)	2022	94.2	93.4
China share of U.S. solar-equipment imports (percent)	2022	0.8	0.8
Panel B. Non-targeted reference-transition outcomes			
China share of world solar-equipment output (percent)	2023	77.1	83.3
	2024	78.6	86.4
Solar share of solar-plus-wind additions (percent)	2023	62.1	75.2
	2024	61.4	79.8

Notes: Panel A reports benchmark objects; U.S. origin shares use equipment imports as the denominator. The wind output share in Panel A is implied by the imputed wind matrix, whereas the solar output share is anchored to IEA PVPS production data. Panel B uses the reference transition with observed 2022 tariffs and no post-2022 deployment or equipment subsidy. Model solar output combines cells and modules, whereas the observed China series covers modules.

5.5 Calibration Fit and Non-Targeted Outcomes

Panel A of Table 2 compares selected shares from the balanced equipment matrices with their source-data counterparts. Balancing to production and absorption margins modestly changes the raw U.S. import composition. The remaining 2022 bilateral expenditures, input–output margins, capacities, household energy shares, tariff receipts, and transfers also enter the benchmark accounts.

Panel B compares non-targeted outcomes from the model's reference transition with observed supplier concentration and technology composition in 2023–24. The reference transition retains observed 2022 tariffs but imposes no deployment or equipment subsidy after 2022. It preserves China's dominance in solar equipment and reproduces the increase in its share from 2023 to 2024, although the model share is lower than the observed module share. It also predicts that solar remains the majority of combined solar-and-wind additions. The lower model solar share may partly reflect technology-specific policy take-up in the data: realized PTC and ITC support can differ between solar and wind, whereas the reference transition does not reproduce those technology-specific effective subsidies. These comparisons assess supplier concentration and technology composition, two margins that matter directly for policy incidence.

6 Policy Assignment in Global Equipment Markets

Using the calibrated model, I ask which instrument and location deliver the larger welfare effect at a common fiscal cost and how much the ranking depends on equipment trade. The benchmark compares deployment and equipment subsidies in an importing market (the United States) and a production hub (China). It removes equipment tariffs while retaining calibrated technologies and nonpolicy trade costs and allowing sourcing to adjust.[13] This benchmark isolates the roles of equipment trade and learning from second-best interactions with existing tariffs. I then use welfare accounting and domestic–foreign quantity responses to relate the equilibrium outcomes to the model's mechanisms.

The deployment instrument is an investment tax credit (ITC) against renewable-investment cost, whereas the equipment instrument is an ad valorem production subsidy that lowers domestic equipment producers' effective unit cost. Policies cover solar and wind and, unless stated otherwise, operate during 2023–2032. Country effects are exact consumption-equivalent variations (CEVs). The global welfare index is the baseline-expenditure-weighted average of country log consumption equivalents defined in equation (17). Carbon damages do not enter the equilibrium allocation, so consumption effects and discounted physical CO_2 are presented separately.

6.1 Instruments and Locations at a Common Fiscal Cost

Common subsidy rates are not economically comparable because the four tax bases differ. I define the common budget within the tariff-free benchmark as the $48.35 billion fiscal present value of a 30 percent U.S. equipment subsidy. I then solve separately for the rate of each other instrument–location pair that attains this budget. Table 3 reports the matched rates.

Panel A of Table 3 shows a location-dependent ranking at fixed spending.[14] In the United States, deployment support produces a positive global effect and a smaller domestic consumption loss than equipment support, whose global effect is slightly negative. In China, deployment support produces the larger national gain, whereas equipment support produces the larger global gain. Thus, at this common fiscal scale, global welfare favors deployment support in the importing market and equipment support at the production hub, as in Proposition 2. This ordering is a calibrated outcome of sourcing, international incidence, and learning values rather than a consequence of the importer and hub labels alone.

All four policies reduce discounted emissions. The carbon-benefit column is an ex-post damage account rather than a benefit–cost calculation: the CEV columns report the equilibrium consumption consequences of financing and policy distortions but exclude climate damages. The consumption effects are small in level; the substantive result is their ordering across instruments and locations.

The benchmark equipment sectors are small: model-covered solar and wind expenditures were

[13] The IRA and source-restriction applications restore observed 2022 equipment tariffs.

[14] Tariff removal leaves calibrated nonpolicy trade costs in place; it does not imply frictionless sourcing.

Table 3: Policy Instruments and Locations With and Without Equipment Trade

Policy	Rate (percent)	World log CEV	U.S. CEV (percent)	China CEV	CO_2 ($MtCO_2$)	Carbon benefit (2022 $bn)
Panel A: Equipment trade						
U.S. deployment	12.86	0.0006	−0.0010	0.0033	−361.7	75.9
U.S. equipment	30.00	−0.0001	−0.0024	0.0022	−296.8	62.3
China deployment	3.12	0.0010	0.0004	0.0049	−503.2	105.7
China equipment	4.70	0.0017	0.0008	0.0038	−542.1	113.8
Panel B: Equipment trade shut down						
U.S. deployment	14.38	0.0006	−0.0004	0.0033	−366.1	76.9
U.S. equipment	26.89	0.0009	−0.0005	0.0045	−403.7	84.8
China deployment	3.12	0.0009	0.0004	0.0050	−496.8	104.3
China equipment	6.57	0.0013	0.0006	0.0066	−553.6	116.3

Notes: Policies operate during 2023–2032. Within each trade regime, every rate is solved to match a $48.35 billion fiscal present value. Panel A is the tariff-free benchmark. Panel B shuts down international trade in solar and wind equipment while trade in fossil fuel and other goods and world-experience learning remain active. "Deployment" denotes the deployment ITC and "equipment" the production subsidy. CO_2 entries are changes in territorial fossil combustion over 2023–2100, discounted at the annual factor $\beta = 0.96$. The carbon-benefit column values those discounted changes at $210 per ton and is not included in the CEV columns.

$139.6 billion and $58.0 billion in 2022. Panel B asks whether trade in these sectors can nevertheless change the policy ranking and reports results from the model without equipment trade. The shutdown eliminates two direct equipment-market margins: U.S. equipment support can no longer expand domestic production partly by displacing foreign suppliers, while Chinese equipment support can no longer transmit its purchaser-price reduction through China's export network. Direct equipment-market incidence is therefore concentrated in the policy country, although world-experience learning, fossil trade, and other equilibrium responses remain.[15]

The ranking consequences differ by country and objective. For U.S. policies, the global-welfare ranking reverses from deployment to equipment support: the global effect of equipment support rises from a 0.0001 percent loss to a 0.0009 percent gain. U.S. welfare, however, continues to rank deployment support slightly higher, although its advantage nearly disappears. For Chinese policies, equipment support remains globally preferred, whereas China's national ranking reverses from deployment under trade to equipment without trade. Consistent with the loss of the export-network channel, the global-welfare advantage of Chinese over U.S. equipment support narrows from 0.0018 to 0.0005 percentage points. Equipment trade therefore affects both instrument choice in an importing market and the value of locating production support at an export hub.

[15] Corollary 2 does not sign these finite changes: Panel B re-solves the full transition at fiscally matched rates, so the removed margins interpret rather than decompose the welfare responses.

Table 4: Domestic and Foreign Responses at a Common Fiscal Cost

	ΔCO_2 ($MtCO_2$)		Δ equipment output (present value, 2022 $bn)	
Policy	Policy country	Rest of world	Policy country	Rest of world
U.S. deployment subsidy	−461.7	+100.0	+20.4	+84.1
U.S. equipment subsidy	−308.2	+11.4	+171.5	−45.4
China deployment subsidy	−620.7	+117.5	+101.8	+9.3
China equipment subsidy	−482.2	−59.9	+194.2	−11.3

Notes: These are the tariff-free equal-budget experiments in Table 3. CO_2 entries are discounted full-horizon changes in territorial fossil combustion. Equipment entries are discounted changes in solar-plus-wind real output by producing region, valued at constant 2022 unit costs.

6.2 International Transmission: Production, Emissions, and Investment Costs

Because electricity remains local while equipment and fossil fuels are traded, the two subsidies transmit production and emissions differently across countries. Table 4 separates responses in the policy country and abroad. Deployment support concentrates abatement in the policy country while routing equipment demand through its supplier network. Under U.S. deployment support, 80.5 percent of the increase in world equipment output occurs abroad. Foreign emissions nevertheless rise by 100.0 $MtCO_2$ under U.S. deployment support and by 117.5 $MtCO_2$ under Chinese deployment support. Deployment support can generate leakage through both fossil prices and a short-run equipment-price channel: the additional equipment demand raises delivered costs and reduces renewable investment abroad. Production learning lowers those costs over time.

Equipment support moves the equipment-price margin in the opposite direction. It lowers the delivered cost of the subsidized origin's equipment in destinations that source from it while concentrating production in that origin. Both equipment policies expand world output while reducing output abroad, so aggregate expansion and geographic reallocation occur together. The net foreign-emissions response differs by policy origin: foreign emissions rise by 11.4 $MtCO_2$ under U.S. support and fall by 59.9 $MtCO_2$ under Chinese support. This contrast is consistent with direct pass-through: Chinese support delivers substantially more equipment-price relief to foreign buyers, whereas U.S. support shifts production toward U.S. suppliers while providing much less direct relief abroad.

Table 5 decomposes 2023–32 average project-cost changes into direct-policy and learning components in the policy country and abroad. Because the table reports policy-window averages, its positive foreign deployment entries show that learning-induced cost declines eventually more than offset the initial equipment-price pressure described above.

In Panel A, the foreign Direct column is zero by construction: the statutory deployment wedge applies only to projects in the granting country. The average foreign cost declines therefore arise almost entirely through learning. Equipment support, by contrast, lowers the purchaser price of subsidized-origin equipment in every destination that sources from that origin. The magnitude of this

Table 5: Domestic and Cross-Border Transmission to Renewable-Investment Costs

		Policy country			Foreign destinations		
Country	Technology	Total	Direct	Learning	Total	Direct	Learning
Panel A: Deployment support							
United States	Solar	14.35	13.77	0.58	0.36	0.00	0.36
United States	Wind	14.22	13.77	0.45	0.18	0.00	0.18
China	Solar	3.59	3.17	0.43	0.40	0.00	0.40
China	Wind	3.45	3.17	0.29	0.16	0.00	0.15
Panel B: Equipment support							
United States	Solar	4.44	3.69	0.75	0.36	0.09	0.26
United States	Wind	15.73	14.25	1.49	1.76	1.58	0.17
China	Solar	2.57	1.88	0.69	1.67	1.05	0.62
China	Wind	3.02	2.62	0.40	0.88	0.70	0.17

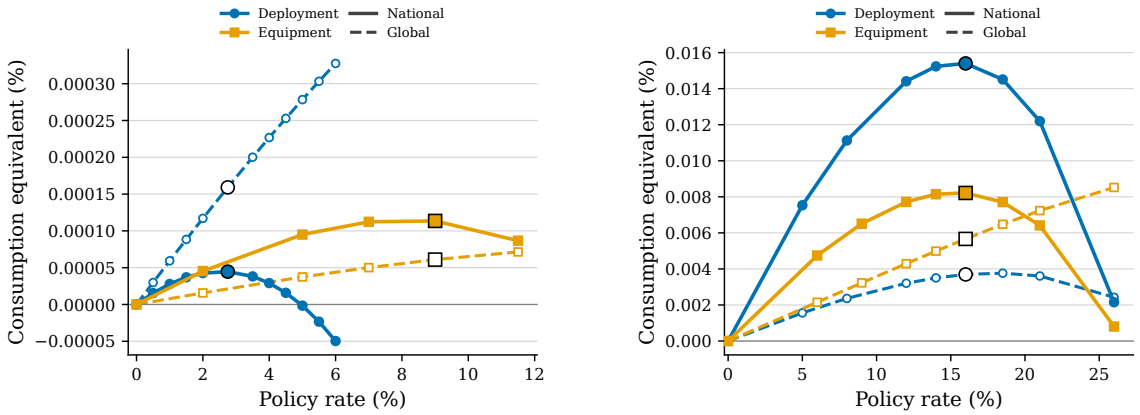
Notes: Entries are 2023–2032 average log project-cost declines, multiplied by 100; positive values denote reductions, and foreign destinations are unweighted means over nonpolicy regions. Direct is the deployment wedge or equipment purchaser-price effect; learning combines world and origin-specific experience. Total includes supplier-cost and balance-of-system terms, jointly below 0.01 point.

pass-through depends on equipment's share of project cost and the destination's sourcing exposure. Under U.S. equipment support, learning accounts for most of the foreign solar cost decline, whereas direct pass-through accounts for about 90 percent of the foreign wind decline. Chinese equipment support produces sizable direct cost reductions abroad in both technologies. The learning column combines world and origin-specific experience.

Appendix Table C.1 reports the corresponding welfare decomposition. Terms-of-trade incidence measured from pre-subsidy origin costs is substantial nationally but nearly cancels globally, while learning contributes positively under all four policies. Within the learning component, world experience accounts for roughly four-fifths of the global contribution under all four policies. Chinese equipment support generates about 60 percent more global learning value than Chinese deployment support, consistent with its global ranking. The two U.S. instruments generate nearly identical global learning contributions, so learning does not account for their ranking.

6.3 Policy Scale and National–World Incidence

The common-budget exercise fixes policy scale. I now allow scale to differ and ask which instrument each country prefers when acting alone. Figure 3 traces national and global welfare over unilateral policy rates, with every other green policy set to zero. U.S. national welfare is higher under equipment support, which peaks at about 9 percent with a consumption gain of 0.000114 percent, than under deployment support, which peaks at about 3 percent with a gain of 0.000045 percent. China's national ranking is the reverse: deployment support generates a gain of 0.0154 percent at its approximately 16 percent peak, compared with 0.0082 percent under equipment support at about the same



(a) United States

(b) China

Figure 3: National and Global Welfare Rank the Instruments Differently

Notes: Markers are solved unilateral grid points; lines connect them for readability. Solid lines report the policy country's consumption equivalent and dashed lines the global consumption equivalent. Large black-outlined markers evaluate both objectives at each instrument's approximate national-welfare peak; they do not identify global optima. All other green policies are zero, and climate damages are excluded. Appendix Table C.7 reports the broader grids.

rate.[16]

Evaluating global welfare at those same instrument-specific national peaks reverses the ordering in both countries. Global welfare is 0.000159 percent under U.S. deployment support and 0.000061 percent under U.S. equipment support; for China, it is 0.00566 percent under equipment support and 0.00369 percent under deployment support. Thus unilateral national incentives favor U.S. equipment and Chinese deployment, whereas global welfare favors U.S. deployment and Chinese equipment. This is the finite-policy counterpart to Corollary 3.

This reversal is consistent with the importer–hub incidence characterized in Proposition 2. In the importing United States, equipment support directs production toward domestic suppliers, whereas deployment support routes demand and learning through the supplier network. At the Chinese hub, deployment support concentrates its direct price effect on domestic investors, whereas equipment support reaches both domestic and foreign buyers. Its price reduction for foreign buyers is a terms-of-trade-like transfer financed by China; in the equal-budget decomposition, it also generates more global learning than Chinese deployment support. National welfare counts the international-price, fiscal, and learning incidence accruing to domestic residents; global welfare also counts foreign buyer and learning gains.

[16]These are separate national optima rather than an equal-budget comparison. Appendix Table C.7 reports the corresponding fiscal present values and broader evaluated grids.

Table 6: Joint U.S.–China Policy Assignments

U.S. instrument	China instrument	World log CEV	U.S. CEV	China CEV	CO_2	U.S. fiscal PV	China fiscal PV
		(percent)			($MtCO_2$)	(2022 $bn)	
Joint policies at equal-budget rates							
Deployment (12.86%)	Deployment (3.12%)	0.0016	−0.0005	0.0082	−864	48.5	48.4
Deployment (12.86%)	Equipment (4.70%)	0.0022	−0.0003	0.0071	−905	48.8	48.9
Equipment (30.00%)	Deployment (3.12%)	0.0009	−0.0020	0.0071	−799	48.4	48.4
Equipment (30.00%)	Equipment (4.70%)	0.0016	−0.0016	0.0060	−827	47.5	47.6
Joint policies at unilateral preferred rates							
Deployment (2.75%)	Deployment (16%)	0.0038	0.0028	0.0159	−3,164	7.6	363.5
Deployment (2.75%)	Equipment (16%)	0.0058	0.0032	0.0087	−2,206	7.9	217.9
Equipment (9%)	Deployment (16%)	0.0037	0.0029	0.0155	−3,142	5.3	363.4
Equipment (9%)	Equipment (16%)	0.0057	0.0033	0.0085	−2,172	4.4	216.3

Notes: Policies operate during 2023–2032 in the tariff-free benchmark. World log CEV uses baseline-expenditure weights and excludes climate damages. The first block uses the unilateral equal-budget rates; the second uses each country's approximate unilateral preferred rate. Fiscal present values are realized joint-policy bills. Negative CO_2 denotes discounted emissions reductions.

6.4 Joint U.S.–China Policy Assignments

The unilateral profiles imply opposing national and global assignments; I next ask whether these rankings persist when both countries intervene.

Table 6 compares all four U.S.–China instrument assignments at two rate sets.[17] The first block uses each instrument's unilateral equal-budget rate. The second uses the rate that maximizes the corresponding country's welfare when that instrument is used alone, so policy scale differs across instruments.

At both rate sets, U.S. deployment combined with Chinese equipment support produces the largest global gain among the four assignments, while the reverse assignment produces the smallest. The ranking matches the calibrated importer–hub comparative advantage in policy use: U.S. deployment support routes additional demand through the importing market's supplier network, whereas Chinese equipment support accumulates production experience at the established hub and lowers equipment costs in the markets it serves. At both evaluated rate sets, the pairwise comparisons preserve the ordering associated with the stronger assignment conditions in Proposition 2.

At the unilateral preferred rates, U.S. deployment plus Chinese equipment support raises welfare in both policy countries and globally. Each country gains when the partner's policy is added to its own. Yet the instruments each country prefers in isolation—U.S. equipment and Chinese deployment—form the assignment with the smallest global gain. National one-instrument incentives therefore point away from the globally leading assignment, as in Corollary 3.

[17] These are fixed-rate assignments, not a strategic equilibrium.

Table 7: A Stylized IRA-Scale Package and Its Components

Experiment	Fiscal present value (2022 $bn)	World log CEV	U.S. CEV (percent)	China CEV	Discounted CO_2 ($MtCO_2$)
Deployment component, 30%	196.1	≈ 0	−0.0091	0.0096	−1,088
Equipment component, 30%	55.4	≈ 0	−0.0032	0.0038	−384
Combined 30/30 package	371.6	−0.0028	−0.0271	0.0189	−1,851

Notes: Policies apply to solar and wind during 2023–2032 with observed 2022 equipment tariffs in place. Consumption effects exclude carbon damages; negative CO_2 entries denote emissions reductions. The package is solved jointly and is not the sum of the component rows.

6.5 A Stylized IRA-Scale Package

To connect the benchmark experiments to current U.S. policy, I restore observed 2022 equipment tariffs and combine a 30 percent deployment ITC with a 30 percent subsidy to domestic solar- and wind-equipment production during 2023–2032.[18] Unlike the tariff-free common-budget comparison in Table 3, the equipment component here operates with observed tariffs and has an endogenous fiscal present value of $55.4 billion rather than $48.35 billion. Table 7 reports the two components and the jointly solved package.

The package shifts consumption welfare away from the United States and toward China while reducing world emissions. It lowers the U.S. consumption equivalent by 0.0271 percent, raises China's by 0.0189 percent, and lowers global consumption welfare excluding climate damages by 0.0028 percent. Discounted world emissions fall by 1.85 $GtCO_2$. Valuing that reduction at the flat $210 benchmark gives about $390 billion in avoided damages and raises carbon-inclusive global welfare by 0.0156 percent.

The joint bill exceeds the sum of the component-only bills because each policy expands the other's fiscal base; this does not imply welfare complementarity.

6.6 Trade Restrictions and Policy Assignment

I restore observed 2022 equipment tariffs and impose a permanent restriction under which U.S. buyers cannot source solar equipment from China or Southeast Asia.[19] The restriction creates no tariff revenue or quota rents; wind equipment and other trade links remain unchanged.

The restriction and the equipment subsidy act on different margins. The subsidy lowers one origin's cost in every destination, whereas the restriction removes selected origin–destination links

[18] The deployment rate is the statutory ITC rate under Sections 48 and 48E. Section 45X instead provides component-specific per-unit credits; the 30 percent equipment wedge places solar and wind on a common policy scale. Appendix B.1 gives the detailed mapping and scope. This is a stylized rate experiment rather than a legislative budget score.

[19] The experiment is broader than current U.S. rules, which restrict Xinjiang-linked imports and address circumvention through Southeast Asian assembly rather than imposing a blanket regional exclusion (U.S. Customs and Border Protection, 2025; U.S. Department of Commerce, 2023). Appendix C.6 studies narrower and temporary alternatives.

Table 8: Trade Restrictions and Policy Assignment

Trade regime	U.S. policy	Rate (percent)	World log CEV	U.S. CEV (percent)	China CEV	Discounted CO_2 ($MtCO_2$)
Panel A: Total effect relative to the factual no-policy transition						
China and Southeast Asia excluded	None	—	−0.0022	−0.0065	−0.0004	140
Panel B: Policy effect relative to no policy under the same trade regime						
Factual sourcing	Deployment subsidy	14.58	0.0009	−0.0007	0.0048	−398
Factual sourcing	Equipment subsidy	30.00	0.0002	−0.0033	0.0055	−387
Restricted sourcing	Deployment subsidy	15.21	0.0008	−0.0009	0.0052	−410
Restricted sourcing	Equipment subsidy	28.33	0.0008	−0.0020	0.0064	−436

Notes: Observed 2022 equipment tariffs are active throughout. Panel A compares the restriction-only transition with factual no policy. Panel B compares each subsidy with no policy under the same trade regime; policy rates are solved to attain the same fiscal target in each trade regime. The restricted no-policy path emits 140 $MtCO_2$ more than factual no policy. The restriction generates no revenue or rents.

only from the U.S. market. It therefore changes both the sourcing composition of deployment-induced demand and the foreign reach of equipment support. Panel A of Table 8 compares the restriction-only transition with factual no policy. Panel B compares each subsidy with no policy under the same trade regime. Policy rates are solved to attain the same fiscal target in each trade regime.

The restriction is costly on its own and shifts relative performance toward equipment support, but it does not make either matched subsidy beneficial for U.S. consumption. Without either subsidy, it lowers U.S. consumption by 0.0065 percent and global consumption welfare by 0.0022 percent, and raises discounted emissions by 140 $MtCO_2$. Because China supplies only 0.6 percent of U.S. solar-equipment purchases under observed tariffs, the restriction primarily removes Southeast Asian rather than Chinese supply.

The two subsidies respond differently to the restricted sourcing environment. The smaller U.S. investment base raises the deployment rate required to match the fiscal target from 14.58 to 15.21 percent; the U.S. loss widens slightly and the global gain falls. For equipment support, the larger protected domestic production base is consistent with a lower matched rate, a smaller U.S. loss, and a larger global gain. China's gains rise slightly in both comparisons, a pattern consistent with its limited direct exposure to the restricted U.S. market and its continued access to world-experience spillovers.

The larger incremental emissions reductions in the restricted-sourcing rows reflect the dirtier restricted no-policy baseline, not cleaner U.S. production. Relative to factual no policy, combining the restriction with either subsidy reduces emissions by less than the corresponding subsidy under factual sourcing. Appendix C.6 reports China-only and temporary restrictions, finite tariffs, and alternative allocations of tariff revenue.

6.7 Carbon Pricing and Green Industrial Policy

The preceding experiments work through renewable-investment costs, equipment demand, and production learning; they do not price the broader carbon externality. This subsection asks whether carbon pricing makes green industrial policy redundant and what distinct role remains for each instrument. Appendix C.8 solves carbon-tax counterfactuals in which fossil use is taxed and the revenue is rebated lump sum in the taxing region. The carbon-tax paths operate through 2100, whereas the subsidy windows close after 2032. A worldwide $210-per-ton tax reduces discounted emissions by 576 $GtCO_2$. Its noncarbon consumption cost is 2.6 percent of world consumption, but avoided emissions valued at the same rate produce a net global gain of about 4.1 percent. This is not a like-for-like return comparison: the tax operates for longer and reaches fossil use throughout the economy, whereas the subsidies act through investment and learning in a small traded-input sector.

Carbon pricing does not make learning support redundant. Adding to the worldwide tax an origin- and technology-specific equipment-subsidy schedule derived from Lemma 1 and iterated to a fixed point on the taxed path raises the carbon-inclusive global gain from 4.09 to 4.20 percent. The schedule imposes an additional noncarbon consumption cost of 0.19 percentage points but generates an additional carbon benefit of 0.30 percentage points, so targeted learning support can remain valuable after emissions are priced. This mirrors the lemma's two-wedge division of labor: carbon pricing acts directly on fossil use, while origin-indexed equipment support targets production experience; in that full-instrument benchmark, deployment support has no first-best role. The combined counterfactual illustrates this division of labor; it is not a solved first best.

Unilateral carbon pricing also has international incidence. Appendix C.8 evaluates U.S.-only taxes of $50, $100, $150, and $210 per ton, each operating through 2100. The $210 tax reduces discounted world emissions by 40.4 $GtCO_2$, lowers U.S. consumption welfare by 1.35 percent before assigning climate benefits, and produces a carbon-inclusive global gain of 0.12 percent. Among the four rates, that gain is largest at $100 per ton, at 0.18 percent; the ranking values avoided damages worldwide, and the exercise does not solve for the tax that would maximize U.S. welfare. The discounted foreign-emissions offset remains between 47 and 48 percent as lower fossil demand in the United States depresses the world fossil price. The tax's direct international incidence begins in the fossil market, whereas the subsidies enter through renewable investment and traded equipment; both markets adjust in general equilibrium. Carbon pricing is therefore the primary instrument for emissions, while origin-indexed equipment support addresses the production-learning externality that the tax leaves unpriced.

7 Conclusion

This paper studies how the instrument and location of green subsidies shape incidence when renewable equipment is traded but electricity remains local. At a common fiscal cost, global welfare favors

U.S. deployment and Chinese equipment support; deployment also minimizes the U.S. loss, while China itself favors deployment. Shutting down equipment trade reverses the U.S. global ranking and removes most of China's location advantage. World experience provides about four-fifths of global learning value, and U.S. deployment with Chinese equipment support yields the largest gain among the four fixed-rate assignments.

Equipment trade and the scope of learning drive these rankings: trade patterns determine where investment-cost relief arrives and where production expands, while world and origin-specific experience determine future costs. At their unilateral national-welfare peaks, the United States favors equipment and China deployment—the reverse of the global ranking. The IRA application shifts consumption welfare away from the United States but cuts discounted world emissions by 1.85 $GtCO_2$; at \$210 per ton, its carbon-inclusive global effect is positive. Carbon pricing targets the broader fossil-use margin, while equipment support corrects production-learning externalities.

References

Abuin, Constanza, "Power Decarbonization in a Global Energy Market: The Climate Effect of U.S. LNG Exports," Working Paper, Johns Hopkins University 2026.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The Environment and Directed Technical Change," *American Economic Review*, February 2012, *102* (1), 131–166.

_ , Ufuk Akcigit, Douglas Hanley, and William Kerr, "The Transition to Clean Technology," *Journal of Political Economy*, 2016, *124* (1), 52–104.

Aghion, Philippe, Lint Barrage, Eric Donald, David Hémous, and Ernest Liu, "Transition to Green Technology along the Supply Chain," NBER Working Paper 33934, National Bureau of Economic Research 2025.

Arkolakis, Costas and Conor Walsh, "Clean Growth," *Working Paper*, 2023.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

Bai, Yan, Keyu Jin, Dan Lu, and Hanxi Wang, "Optimal Trade Policy with International Technology Diffusion," NBER Working Paper 32097, National Bureau of Economic Research 2024.

Baker, Erin, Meredith Fowlie, Derek Lemoine, and Stanley S. Reynolds, "The Economics of Solar Electricity," *Annual Review of Resource Economics*, 2013, *5*, 387–426.

Baldwin, Robert E., "The Case against Infant Industry Protection," *Journal of Political Economy*, 1969, *77* (3), 295–305.

Bardhan, Pranab K., "On Optimum Subsidy to a Learning Industry: An Aspect of the Theory of Infant-Industry Protection," *International Economic Review*, 1971, *12* (1), 54–70.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, and Nahim Zahur, "Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry," NBER Working Paper 33378, National Bureau of Economic Research January 2025.

_ , Myrto Kalouptsidi, and Nahim Bin Zahur, "Industrial Policy Implementation: Empirical Evidence from China's Shipbuilding Industry," *Review of Economic Studies*, November 2025, *92* (6), 3611–3648.

Bistline, John, Neil Mehrotra, and Catherine Wolfram, "Economic Implications of the Climate Provisions of the Inflation Reduction Act," *Brookings Papers on Economic Activity*, 2023, *2023* (1), 77–182.

Bolinger, Mark, Ryan Wiser, and Eric O'Shaughnessy, "Levelized cost-based learning analysis of utility-scale wind and solar in the United States," *iScience*, 2022, *25* (6).

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 2015, *82* (1), 1–44.

_ , Marcelo Dolabella, Mauricio Moreira, Matthew Murillo, and Fernando Parro, "Voluntary Emis-

sion Restraints in Developing Economies: The Role of Trade Policy," *Working Paper*, 2024.

_ , Maximiliano Dvorkin, and Fernando Parro, "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 2019, *87* (3), 741–835.

Casey, Gregory, Woongchan Jeon, and Christian Traeger, "The Macroeconomics of Clean Energy Subsidies," *Journal of Environmental Economics and Management*, 2026, *138*, 103311.

Copeland, Brian R., Joseph S. Shapiro, and M. Scott Taylor, "Globalization and the environment," *Handbook of International Economics*, 2021, *5*, 61–146.

Covert, Thomas R. and Richard L. Sweeney, "Winds of Change: Estimating Learning by Doing without Cost or Input Data," *Working Paper*, 2024.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced Trade," *American Economic Review Papers and Proceedings*, 2007, *97* (2), 351–355.

Donald, Eric, "Spillovers and the Direction of Innovation," *Working Paper*, 2024.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

_ , _ , Brent Neiman, and John Romalis, "Trade and the Global Recession," *American Economic Review*, 2016, *106* (11), 3401–3438.

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

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

Fried, Stephie, "Climate Policy and Innovation: A Quantitative Macroeconomic Analysis," *American Economic Journal: Macroeconomics*, January 2018, *10* (1), 90–118.

Friedlingstein, Pierre et al., "Global Carbon Budget 2025," *Earth System Science Data*, 2026, *18*, 3211–3288.

Gentile, Claudia, "Relying on Intermittency: Clean Energy, Storage, and Innovation in a Macro Climate Model," *Working Paper*, 2024.

Gerarden, Todd D., Bryan K. Bollinger, Kenneth Gillingham, and Daniel Xu, "Strategic Avoidance and the Welfare Impacts of U.S. Solar Panel Tariffs," Working Paper 34401, National Bureau of Economic Research October 2025.

Goldberg, Pinelopi K., Réka Juhász, Nathan J. Lane, Giulia Lo Forte, and Jeff Thurk, "Industrial Policy in the Global Semiconductor Sector," *NBER Working Paper*, 2024, *32651*.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 1997, *87* (3), 342–362.

Hahn, Robert W., Nathaniel Hendren, Robert D. Metcalfe, and Ben Sprung-Keyser, "A Welfare Analysis of Policies Impacting Climate Change," *American Economic Review*, July 2026, *116* (7), 2368–2421.

Helm, Carsten and Mathias Mier, "Steering the energy transition in a world of intermittent electric-

ity supply: Optimal subsidies and taxes for renewables and storage," *Journal of Environmental Economics and Management*, 2021, *109*, 102497.

Hémous, David, "The Dynamic Impact of Unilateral Environmental Policies," *Journal of International Economics*, November 2016, *103*, 80–95.

Hirth, Lion, "The Market Value of Variable Renewables: The Effect of Solar Wind Power Variability on Their Relative Price," *Energy Economics*, 2013, *38*, 218–236.

Hsiao, Allan, "Coordination and Commitment in International Climate Action: Evidence from Palm Oil," *Working Paper*, 2024.

IEA PVPS, "Trends in Photovoltaic Applications 2024," Technical Report, International Energy Agency Photovoltaic Power Systems Programme 2024.

Internal Revenue Service, "Advanced Manufacturing Production Credit," IRS.gov 2026. Page last reviewed or updated February 14, 2026.

Irwin, Douglas A. and Peter J. Klenow, "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy*, December 1994, *102* (6), 1200–1227.

Jo, Ara, "Substitution between Clean and Dirty Energy with Biased Technical Change," *International Economic Review*, May 2025, *66* (2), 883–902.

Joskow, Paul L., "Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies," *American Economic Review*, May 2011, *101* (3), 238–241.

Juhász, Réka, "Temporary protection and technology adoption: Evidence from the napoleonic blockade," *American Economic Review*, 2018, *108* (11), 3339–76.

Juhász, Réka, Nathan Lane, and Dani Rodrik, "The New Economics of Industrial Policy," *Annual Review of Economics*, August 2024, *16*, 213–242.

Kleinman, Benny, Ernest Liu, and Stephen J Redding, "Dynamic Spatial General Equilibrium," *Econometrica*, 2023, *91* (2), 385–424.

Kortum, Samuel E. and David A. Weisbach, "Optimal Unilateral Carbon Policy," *Working Paper*, 2021.

Kotchen, Matthew J. and Giovanni Maggi, "Carbon Taxes and Green Subsidies in a World Economy," NBER Working Paper 34080, National Bureau of Economic Research July 2025.

Krugman, Paul, "The Narrow Moving Band, the Dutch Disease, and the Competitive Consequences of Mrs. Thatcher," *Journal of Development Economics*, 1987, *27*, 41–55.

Lane, Nathan, "Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea," *Quarterly Journal of Economics*, August 2025, *140* (3), 1683–1741.

Liu, Ernest, "Industrial Policies in Production Networks," *Quarterly Journal of Economics*, 2019, *134* (4), 1883–1948.

Melitz, Marc J., "When and How Should Infant Industries Be Protected?," *Journal of International Economics*, May 2005, *66* (1), 177–196.

Mutreja, Piyusha, Michael Sposi, and B. Ravikumar, "Capital Goods Trade, Relative Prices and

Economic Development," *Review of Economic Dynamics*, January 2018, *27*, 101–122.

National Renewable Energy Laboratory, "2024 Annual Technology Baseline," National Renewable Energy Laboratory 2024.

Papageorgiou, Chris, Marianne Saam, and Patrick Schulte, "Substitution Between Clean And Dirty Energy Inputs," *The Review of Economics and Statistics*, May 2017, *99* (2), 281–290.

Ramasamy, Vignesh, Jarett Zuboy, David Feldman, Sridhar Narayanaswami, Michael Woodhouse, and Robert Margolis, "Documenting 15 Years of Reductions in U.S. Solar Photovoltaic System Costs," Technical Report NREL/TP-7A40-92536, National Renewable Energy Laboratory 2025.

Ravikumar, B., Ana Maria Santacreu, and Michael Sposi, "Capital Accumulation and Dynamic Gains From Trade," *Journal of International Economics*, 2019, *119*, 93–110.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, David Smith, Kevin Cromar, Delavane Diaz, Frances C. Moore, Ulrich K. Müller, Richard J. Plevin, Adrian E. Raftery, Hana Ševčíková, Hannah Sheets, James H. Stock, Tammy Tan, Mark Watson, Tony E. Wong, and David Anthoff, "Comprehensive Evidence Implies a Higher Social Cost of CO2," *Nature*, 2022, *610*, 687–692.

Rodrik, Dani, "Green Industrial Policy," *Oxford Review of Economic Policy*, 2014, *30* (3), 469–491.

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

Shin, David D., "Climate Policies under Dynamic Factor Adjustment," *Working Paper*, 2026.

Stehly, Tyler, Patrick Duffy, and Daniel Mulas Hernando, "Cost of Wind Energy Review: 2024 Edition," Technical Report NREL/PR-5000-91775, National Renewable Energy Laboratory 2024.

Thornton, Rebecca Achee and Peter Thompson, "Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding," *American Economic Review*, December 2001, *91* (5), 1350–1368.

U.S. Customs and Border Protection, "What Is the Uyghur Forced Labor Prevention Act?," U.S. Customs and Border Protection 2025.

U.S. Department of Commerce, "Department of Commerce Issues Final Determination of Circumvention Inquiries of Solar Cells and Modules from China," U.S. Department of Commerce 2023.

Way, Rupert, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer, "Empirically Grounded Technology Forecasts and the Energy Transition," *Joule*, September 2022, *6* (9), 2057–2082.

Wiser, Ryan, Mark Bolinger, Ben Hoen, Dev Millstein, Joseph Rand, Galen Barbose, Naim Darghouth, Will Gorman, Seongeun Jeong, Andrew Mills, and Bentham Paulos, "Land-Based Wind Market Report: 2024 Edition," Annual Report, Lawrence Berkeley National Laboratory 2024.

Wiskich, Anthony, "Perfect So Far? Substitutability Between Wind & Solar and Dirty Electricity Generation," *Working Paper*, 2023.

World Meteorological Organization and World Trade Organization, "Supporting the Renewable Elec-

tricity Transition through Trade: Unlocking Re-Globalization Opportunities via Interconnection," Technical Report, World Meteorological Organization and World Trade Organization, Geneva 2024.

Xiang, Wei, "Clean Growth and Environmental Policies in the Global Economy," *Working Paper*, 2023.

Young, Alwyn, "Learning by Doing and the Dynamic Effects of International Trade," *The Quarterly Journal of Economics*, 1991, *106* (2), 369–405.

Online Appendix:
Green Industrial Policy in a Globalized Economy

Sungwan Hong
University of Pittsburgh

Appendix A Derivations **1**
A.1 Proofs for International Incidence and Policy Assignment 1
A.2 Composite Intermediate Goods (Dixit–Stiglitz) 8
A.3 Energy Aggregation and Expenditure Shares 8
A.4 Consumption Welfare and the Ex-post Carbon Account 8

Appendix B Calibration: Data and Estimation **10**
B.1 Renewable Energy Provisions in the IRA 10
B.2 Geographic and Time Coverage 11
B.3 Equipment Trade and Benchmark Sourcing 12
B.4 Evidence on the Inner-Nest Elasticity 13
B.5 Learning Calibration and Supporting Estimates 14
B.6 Intermittency, Value Factors, and Annual Aggregation 18
B.7 Measuring the Initial Productive Stocks 19
B.8 Emissions Accounting 20
B.9 Social Cost of Carbon 21

Appendix C Mechanisms and Robustness **21**
C.1 Welfare Decomposition 21
C.2 Learning Robustness and Technology-Specific Returns 23
C.3 Structural Sensitivity 24
C.4 Policy Duration and the PTC Mapping 25
C.5 Unilateral Policy-Rate Profiles 26
C.6 Source Restrictions and Tariff Alternatives 26
C.7 A Local Pigouvian Benchmark 28

C.8 Carbon Pricing in Equilibrium 29
C.9 Learning Stock, Horizon, and Accounting Revaluations 30

Appendix D Exact Hat Algebra in the Trade Block and Computational Implementation 30
D.1 Within-period trade block 30
D.2 Resource and Policy Accounting 32
D.3 Post-2100 Values and Numerical Solution 33
D.4 Solution Algorithm 34

Appendix A Derivations

A.1 Proofs for International Incidence and Policy Assignment

Trade-price envelope. For the extension with pre-existing tariffs, let $V^j_{ni,t} = \pi^j_{ni,t} X^j_{n,t}/(1+\tau^j_{ni,t})$ be the tariff-exclusive value of purchases by destination n from origin i. With $\ell^j_{i,t} = \log[(1-s^j_{i,t})c^j_{i,t}]$, the Eaton–Kortum price envelope implies

$$X^j_{n,t}\, d\log P^j_{n,t} = \sum_i (1+\tau^j_{ni,t}) V^j_{ni,t} \left[d\log(1+\tau^j_{ni,t}) + d\ell^j_{i,t} - \frac{1}{\theta^j} d\log A^j_{i,t} \right]. \tag{A.1}$$

Tariff revenue in the same sector is $TR^j_{n,t} = \sum_i \tau^j_{ni,t} V^j_{ni,t}$, so

$$dTR^j_{n,t} = \sum_i \left[V^j_{ni,t}\, d\tau^j_{ni,t} + \tau^j_{ni,t}\, dV^j_{ni,t} \right].$$

At fixed iceberg costs, define source-specific real purchases by $d\log \mathsf{m}^j_{ni,t} = d\log V^j_{ni,t} - d\ell^j_{i,t} + (1/\theta^j) d\log A^j_{i,t}$. Combining the household's purchaser-price change with the tariff-revenue rebate cancels the direct $d\tau$ term. The remaining contribution from link (n,i,j) is

$$\underbrace{-V^j_{ni,t} d\ell^j_{i,t}}_{\text{import-price incidence}} + \underbrace{\frac{V^j_{ni,t}}{\theta^j} d\log A^j_{i,t}}_{\text{productivity gain}} + \underbrace{\tau^j_{ni,t} V^j_{ni,t} d\log \mathsf{m}^j_{ni,t}}_{\text{tariff interaction}}.$$

The export side of the first term enters domestic factor income through origin-market clearing.

For the exact dynamic envelope, let a scalar p scale a complete policy sequence around the reference transition; overbars denote levels on that transition, and dots denote total derivatives with

respect to p. At a tariff-free reference, define

$$\mathcal{I}^p_{n,t} = \sum_j \sum_d \bar{V}^j_{dn,t} \dot{\ell}^j_{n,t} - \sum_j \sum_i \bar{V}^j_{ni,t} \dot{\ell}^j_{i,t} + \dot{\mathcal{T}}_{n,t},$$
$$\mathcal{L}^p_{n,t} = \sum_{r \in \mathcal{R}} \sum_i \frac{\bar{V}^{M^r}_{ni,t}}{\theta^r} \frac{d \log A^{M^r}_{i,t}}{dp}. \tag{A.2}$$

The first line is international-price and transfer incidence; the second values the induced equipment-productivity path. With pre-existing tariffs, $\mathcal{H}^p_{n,t}$ collects the tariff-interaction terms in the last underbrace above. Let

$$\overline{\mathcal{E}}_n \equiv \sum_{t=1}^{\infty} \beta^{t-1} \bar{E}_{n,t}, \qquad \varpi_n \equiv \frac{\overline{\mathcal{E}}_n}{\sum_m \overline{\mathcal{E}}_m}, \qquad \mathcal{B}_\beta \equiv \sum_{t=1}^{\infty} \beta^{t-1}.$$

Dynamic envelope. Start from $E_{n,t} = P^C_{n,t} C_{n,t}$ and define the reference nominal welfare weight

$$\mu_{n,t} \equiv \frac{\beta^{t-1}}{\bar{E}_{n,t}}.$$

The contribution of date t to the derivative of lifetime utility is therefore

$$\beta^{t-1} \frac{d \log C_{n,t}}{dp} = \mu_{n,t} \left[\dot{E}_{n,t} - \bar{E}_{n,t} \frac{d \log P^C_{n,t}}{dp} \right]. \tag{A.3}$$

Before differentiating, consolidate the fiscal accounts:

$$(1 - s^{I,r}_{n,t}) C^{\text{proj},r}_{n,t} I^r_{n,t} + \Lambda^{I,r}_{n,t} = C^{\text{proj},r}_{n,t} I^r_{n,t}, \qquad \text{deployment subsidy},$$
$$Y^{M^r}_{i,t} - \Lambda^{M,r}_{i,t} = \sum_n V^{M^r}_{ni,t}, \qquad \text{equipment subsidy}.$$

The subsidy payment and the lump-sum tax financing it thus cancel as a direct transfer at a zero-rate reference.

Differentiate the consolidated budget and substitute firms' zero-profit and factor-demand conditions, goods-market clearing, and the trade-price envelope above. Intermediate purchases and factor payments cancel between users and suppliers. The remaining date-t money metric is

$$\dot{E}_{n,t} - \bar{E}_{n,t} \frac{d \log P^C_{n,t}}{dp} = \mathcal{I}^p_{n,t} + \mathcal{L}^p_{n,t} + \mathcal{H}^p_{n,t} + \sum_{k \in \mathcal{K}} \left(\bar{r}^k_{n,t} \dot{K}^k_{n,t} - \bar{P}^{\text{inv},k}_{n,t} \dot{I}^k_{n,t} \right), \tag{A.4}$$

where $\mathcal{H}^p_{n,t}$ is the tariff-interaction term and is zero in the tariff-free reference. The sum runs over all three capital stocks, because the fund chooses non-energy investment alongside the two renewable margins. For the renewable technologies the acquisition price is the renewable investment bundle,

$\bar{P}^{\text{inv},r}_{n,t} = \bar{C}^{\text{proj},r}_{n,t}$; non-energy investment is purchased out of the non-energy composite, so $\bar{P}^{\text{inv},D}_{n,t} = \bar{P}^{O}_{n,t}$.

It remains to eliminate the privately chosen capital variations. Let $a^k_{n,t} \equiv 1 - \delta_k + \bar{\Phi}^k_{K,n,t}$. Differentiating the capital law and using the definition of the installed-capital value at the zero-subsidy reference gives, for every $k \in \mathcal{K}$,

$$\begin{aligned} \dot{K}^k_{n,t+1} &= a^k_{n,t}\dot{K}^k_{n,t} + \bar{\Phi}^k_{I,n,t}\dot{I}^k_{n,t}, & \bar{P}^{\text{inv},k}_{n,t} &= \bar{q}^k_{n,t}\bar{\Phi}^k_{I,n,t}, \\ \bar{r}^k_{n,t}\dot{K}^k_{n,t} - \bar{P}^{\text{inv},k}_{n,t}\dot{I}^k_{n,t} &= -\bar{q}^k_{n,t}\dot{K}^k_{n,t+1} + \left(\bar{r}^k_{n,t} + \bar{q}^k_{n,t}a^k_{n,t}\right)\dot{K}^k_{n,t}. \end{aligned} \tag{A.5}$$

Non-energy capital enters this construction as a third stock with the same adjustment technology. Equations (13) and (14), together with $\mu_{n,t-1}\bar{M}_{n,t} = \mu_{n,t}$, imply, for $t \geq 2$ and every $k \in \mathcal{K}$,

$$\mu_{n,t}\left(\bar{r}^k_{n,t}\dot{K}^k_{n,t} - \bar{P}^{\text{inv},k}_{n,t}\dot{I}^k_{n,t}\right) = \mu_{n,t-1}\bar{q}^k_{n,t-1}\dot{K}^k_{n,t} - \mu_{n,t}\bar{q}^k_{n,t}\dot{K}^k_{n,t+1}.$$

All three Euler equations share the discount factor $\bar{M}_{n,t}$ and the same adjustment technology, so the three blocks telescope under one argument. Summing over dates therefore leaves only boundary terms in each block:

$$\sum_{t=1}^{T}\mu_{n,t}\left(\bar{r}^k_{n,t}\dot{K}^k_{n,t} - \bar{P}^{\text{inv},k}_{n,t}\dot{I}^k_{n,t}\right) = \mu_{n,1}\left(\bar{r}^k_{n,1} + \bar{q}^k_{n,1}a^k_{n,1}\right)\dot{K}^k_{n,1} - \mu_{n,T}\bar{q}^k_{n,T}\dot{K}^k_{n,T+1}.$$

The initial state is fixed, so $\dot{K}^k_{n,1} = 0$ for every $k \in \mathcal{K}$. End-of-2022 renewable capacities and the inherited non-energy stock are data, and observed 2022 non-energy investment enters that stock before the first policy choice rather than as a date-2023 decision, so the first-date boundary term vanishes in all three blocks. The transversality condition removes each limiting terminal term, including $\lim_{T\to\infty}\mu_{n,T}\bar{q}^D_{n,T}\dot{K}^D_{n,T+1}$. Combining equations (A.3)–(A.5) yields

$$\dot{\ell}^{\text{cons}}_{n} = \frac{1}{\mathcal{B}_\beta}\sum_{t=1}^{\infty}\frac{\beta^{t-1}}{\bar{E}_{n,t}}\left(\mathcal{I}^p_{n,t} + \mathcal{L}^p_{n,t} + \mathcal{H}^p_{n,t}\right). \tag{A.6}$$

Appending the separately defined carbon account, multiplying by ϖ_n, and summing over countries gives

$$\begin{aligned} \dot{\ell}^{R,SCC}_{W} = &\frac{1}{\mathcal{B}_\beta}\sum_n \varpi_n \sum_{t=1}^{\infty}\frac{\beta^{t-1}}{\bar{E}_{n,t}}\left(\mathcal{I}^p_{n,t} + \mathcal{L}^p_{n,t} + \mathcal{H}^p_{n,t}\right) \\ &- \frac{1}{\sum_n \bar{\mathcal{E}}_n}\sum_{t=1}^{\infty}\beta^{t-1}SCC_t\dot{Z}_t. \end{aligned} \tag{A.7}$$

At the tariff-free reference, $\mathcal{H}^p_{n,t} = 0$. Equations (A.6) and (A.7) are the full-model envelope underlying the fiscal-normalized return decomposition in equation (19).

Reindexing destinations and origins gives the accounting identity

$$\sum_n \left[\sum_j \sum_d \bar{V}^j_{dn,t} \dot{\ell}^j_{n,t} - \sum_j \sum_i \bar{V}^j_{ni,t} \dot{\ell}^j_{i,t} \right] = 0.$$

Because $\sum_n \dot{\mathcal{T}}_{n,t} = 0$, $\sum_n \mathcal{I}^p_{n,t} = 0$ in dollars. In welfare, however, each term is multiplied by $\varpi_n / \bar{E}_{n,t}$. A common marginal welfare weight per dollar across countries at date t therefore guarantees that international incidence cancels for every admissible perturbation. With heterogeneous weights, incidence may still cancel for a particular policy perturbation, but it need not do so. This establishes a sufficient transfer-neutral special case without imposing it on the quantitative baseline.

Fiscal normalization and marginal choice. Let a_z scale the statutory rates of policy z, and let $\mathcal{F}_z(\mathbf{a})$ be its present-value fiscal bill. At the zero-policy reference, $\partial \mathcal{F}_z / \partial a_z > 0$, while $\partial \mathcal{F}_z / \partial a_{z'} = 0$ for $z' \neq z$: a policy with a zero own rate has a zero bill even if another policy changes its prospective base. The inverse-function theorem therefore permits the local change of coordinates $\mathbf{a} \mapsto \mathbf{f}$, where $f_z = \mathcal{F}_z(\mathbf{a})$. For any scalar perturbation p that scales policy z, the fiscal-normalized return is explicitly

$$R^o_z = \left. \frac{d\mathcal{W}^o / dp}{df_z / dp} \right|_{p=0}. \tag{A.8}$$

Dividing each component of the policy envelope by the same df_z/dp defines T^o_z, C^o_z, $L^{o,r}_{z,W}$, and $L^{o,r}_{z,O}$. A second-order Taylor expansion of objective $\mathcal{W}^o$ in $\mathbf{f}$ then gives equation (20). The first-order conditions with $\mathbf{f} \geq 0$ are the complementary-slackness conditions in Lemma 2; when all constraints are slack, they reduce to $H_o \mathbf{f} = \mathbf{R}_o$. Equation (21) follows by completing the square in one dimension.

The learning term already includes the entire induced productivity path and must not be added again as an output co-state. To obtain the stripped components used in equation (19), differentiate equations (11) and (12). Under the reference-weighted global objective, one unit of technology-r output at date t has world-experience value

$$\begin{aligned}
\Omega^{G,r}_{W,t} &= \frac{1}{\mathcal{B}_\beta} \sum_{u>t} \beta^{u-1} \frac{\nu^r_{\text{tot}}}{\theta^r} \underbrace{(\mu_{Q,r})^{u-t-1}}_{\text{retention to } u} \underbrace{\frac{\sum_{n,h} (\varpi_n / \bar{E}_{n,u}) \bar{V}^{M^r}_{nh,u}}{\bar{Q}^r_{\text{glob},u}}}_{\text{world-experience value}}, \\
\Omega^{G,r}_{i,O,t} &= \frac{1}{\mathcal{B}_\beta} \sum_{u>t} \beta^{u-1} \frac{\nu^r_{\text{tot}}}{\theta^r} \underbrace{(\mu_{Q,r})^{u-t-1}}_{\text{retention to } u} \underbrace{\frac{\sum_{n} (\varpi_n / \bar{E}_{n,u}) \bar{V}^{M^r}_{ni,u}}{\bar{Q}^r_{i,u}}}_{\text{origin-specific value}}, \\
\Omega^{G,r}_{i,t} &= (1 - \lambda^r_L) \Omega^{G,r}_{W,t} + \lambda^r_L \Omega^{G,r}_{i,O,t}.
\end{aligned} \tag{A.9}$$

Thus the direct learning contribution to reference-weighted global welfare is $\sum_{r,i,t} \Omega^{G,r}_{i,t} \dot{y}^{M,r}_{i,t}$. The planner and national components replace the reference weights in these expressions with their re-

spective marginal values of income. International-price incidence remains in $\mathcal{I}^p$, and any interaction with a pre-existing tariff remains in $\mathcal{H}^p$.

Pigouvian implementation. Let λ_t^W denote the planner's common shadow value of the date-t numeraire and let $m_{t,u}^W \equiv \lambda_u^W / \lambda_t^W$. With $g_r \equiv \nu_{\text{tot}}^r / \theta^r$, the marginal learning value used in Lemma 1 is

$$\mathcal{B}_{i,t}^{L,r} = \sum_{u>t} m_{t,u}^W (\mu_{Q,r})^{u-t-1} g_r \left[(1-\lambda_L^r) \frac{\sum_h Y_{h,u}^{M^r}}{Q_{\text{glob},u}^r} + \lambda_L^r \frac{Y_{i,u}^{M^r}}{Q_{i,u}^r} \right]. \tag{A.10}$$

The first term inside the brackets values the common experience stock; the second values the origin-specific stock.

Differentiating the experience law and equation (12) gives, for $u > t$,

$$\frac{\partial \log A_{h,u}^{M^r}}{\partial y_{i,t}^{M,r}} = (\mu_{Q,r})^{u-t-1} \nu_{\text{tot}}^r \left[\frac{1-\lambda_L^r}{Q_{\text{glob},u}^r} + \frac{\lambda_L^r}{Q_{i,u}^r} \mathbf{1}\{h=i\} \right].$$

At a tariff-free planner allocation, the Eaton–Kortum envelope values a log productivity increase at origin h by $Y_{h,u}^{M^r}/\theta^r$ units of the date-u numeraire. Multiplying by the derivative above, summing over future dates and origins, and applying $m_{t,u}^W = \lambda_u^W / \lambda_t^W$ gives $\mathcal{B}_{i,t}^{L,r}$. Equations (5) and the equipment-subsidy budget imply $\Lambda_{i,t}^{M,r} = s_{i,t}^{M,r} Y_{i,t}^{M^r} = s_{i,t}^{M,r} \tilde{c}_{i,t}^{M,r} y_{i,t}^{M,r}$, so equating the subsidy per unit of real output to its marginal learning value gives the first equality in equation (18). When $0 \le \mathcal{B}_{i,t}^{L,r} / \tilde{c}_{i,t}^{M,r} < 1$, the specific learning wedge is implemented by the admissible ad valorem rate $s_{i,t}^{M,r,*} = \mathcal{B}_{i,t}^{L,r} / \tilde{c}_{i,t}^{M,r}$. Otherwise the same specific wedge remains the Pigouvian object, but the ad valorem parameterization with $1 - s_{i,t}^{M,r} > 0$ cannot implement it.

One unit of fossil absorption in country n, whether purchased through household energy services or as the fossil component of a producer's energy input, creates ϵ_n units of emissions. A specific tax $\epsilon_n SCC_t$ levied at every use margin therefore prices its marginal global damage; were the tax confined to the household margin, producers' fossil use would remain an unpriced emitting margin and the two instruments would no longer separate. With learning and carbon priced directly, all remaining margins are competitive and the planner and private fund face the same capital laws, including the non-energy accumulation technology. Their renewable- and non-energy-investment Euler equations consequently coincide without an acquisition or generation subsidy; the transversality conditions remove the limiting capital and experience terms. This proves Lemma 1.

The lemma identifies the required wedges, not a unique statutory instrument. An alternative instrument mix is equivalent only if it reproduces the origin-by-technology production wedge and the fossil-use wedge along the same allocation; matching the capitalized investment incentive alone does not equalize fiscal outlays, especially for a PTC paid on inframarginal generation.

Proof of Proposition 1 and Corollary 1. For equipment technology r, differentiation of the gravity

share at fixed X, c, A gives

$$\frac{\partial \log \pi_{nh,t}^{M^r}}{\partial s_{i,t}^{M,r}} = \frac{\theta^r}{1 - s_{i,t}^{M,r}} \left[\mathbf{1}\{h = i\} - \pi_{ni,t}^{M^r} \right].$$

Conditional on current destination spending, unit costs, and productivities held fixed in this derivative, equations (3) and (5) imply, in a tariff-free neighborhood of the reference,

$$y_{h,t}^{M,r} = \frac{1}{(1 - s_{h,t}^{M,r}) \bar{\tilde{c}}_{h,t}^{M,r}} \sum_n V_{nh,t}^{M^r}.$$

At the zero-subsidy reference, this identity reduces to

$$\bar{y}_{h,t}^{M,r} = \frac{1}{\bar{\tilde{c}}_{h,t}^{M,r}} \sum_n \bar{V}_{nh,t}^{M^r}.$$

Differentiating the neighborhood identity, rather than only its zero-policy restriction, and substituting the share derivative gives

$$\left. \frac{\partial y_{h,t}^{M,r}}{\partial s_{i,t}^{M,r}} \right|_{X,c,A} = \frac{1}{\bar{\tilde{c}}_{h,t}^{M,r}} \sum_n \bar{V}_{nh,t}^{M^r} \left\{ \theta^r \left[\mathbf{1}\{h = i\} - \bar{\pi}_{ni,t}^{M^r} \right] + \mathbf{1}\{h = i\} \right\}.$$

The supported origin gains both the sourcing response and the subsidy top-up; every competing origin loses sales wherever the two origins overlap. Define $\bar{S}_{i,t}^r \equiv \sum_n \bar{V}_{ni,t}^{M^r}$ and $\varsigma_{ni,t}^r \equiv \bar{V}_{ni,t}^{M^r} / \bar{S}_{i,t}^r$ for every active origin. For objective o and learning component k, define

$$v_{h,t}^{o,r,k} \equiv \frac{\Omega_{h,t}^{o,r,k}}{\bar{\tilde{c}}_{h,t}^{M,r}}, \qquad k \in \{W, O\},$$

where the world component has $\Omega_{h,t}^{o,r,W} = \Omega_{W,t}^{o,r}$. Multiplying the output derivative by $\Omega_{h,t}^{o,r,k}$, summing over origins, and using $\bar{V}_{ni,t}^{M^r} = \bar{\pi}_{ni,t}^{M^r} \bar{X}_{n,t}^{M^r}$ gives, for either component,

$$\sum_h \Omega_{h,t}^{o,r,k} \left. \frac{\partial y_{h,t}^{M,r}}{\partial s_{i,t}^{M,r}} \right|_{X,c,A} = \bar{S}_{i,t}^r \left[v_{i,t}^{o,r,k} + \theta^r \sum_n \varsigma_{ni,t}^r \sum_{h \neq i} \bar{\pi}_{nh,t}^{M^r} \left(v_{i,t}^{o,r,k} - v_{h,t}^{o,r,k} \right) \right].$$

Because $\partial \Lambda_{i,t}^{M,r} / \partial s_{i,t}^{M,r}|_0 = \bar{S}_{i,t}^r$, the bracket is the direct learning return per marginal equipment-subsidy dollar in equation (25). The nested sum remains well defined as trade vanishes.

For deployment, one additional dollar of purchaser expenditure in destination n creates $\bar{\pi}_{ni,t}^{M^r}$ dollars of receipts at origin i in the tariff-free benchmark. Dividing by $\bar{\tilde{c}}_{i,t}^{M,r}$ and weighting by $\Omega_{i,t}^{o,r,k}$ gives $\mathcal{V}_{n,t}^{I,r,k} = \sum_i \bar{\pi}_{ni,t}^{M^r} v_{i,t}^{o,r,k}$, proving equation (24). Dividing the induced destination-spending response

by $\partial \Lambda_n^{I,r} / \partial s_n^{I,r}|_0$ gives the equipment-spending multiplier ξ_n^I in the main text. With a pre-existing tariff, the corresponding origin receipt is $\bar{\pi}_{ni,t}^{M^r} / (1 + \bar{\tau}_{ni,t}^{M^r})$.

The direct price derivatives follow from the after-subsidy acquisition price and the Eaton–Kortum price-index envelope. More generally, for any differentiable policy perturbation p,

$$\begin{aligned} \frac{d \log \widetilde{P}_{n,t}^{\text{inv},r}}{dp} = & -\frac{1}{1 - s_{n,t}^{I,r}} \frac{ds_{n,t}^{I,r}}{dp} - \omega^r \sum_i \bar{\pi}_{ni,t}^{M^r} \frac{1}{1 - s_{i,t}^{M,r}} \frac{ds_{i,t}^{M,r}}{dp} \\ & + \omega^r \sum_i \bar{\pi}_{ni,t}^{M^r} \frac{d \log c_{i,t}^{M^r}}{dp} + (1 - \omega^r) \frac{d \log P_{n,t}^{O}}{dp} \\ & - \frac{\omega^r}{\theta^r} \sum_i \bar{\pi}_{ni,t}^{M^r} \frac{d \log A_{i,t}^{M^r}}{dp}. \end{aligned} \tag{A.11}$$

No derivative of the sourcing share appears because the expenditure shares are envelope weights in the price-index derivative. Holding the last two lines fixed and evaluating at zero gives equation (22). Dividing by the two marginal fiscal bases gives equation (23); applying the destination abatement values gives equation (26). This proves Proposition 1 and Corollary 1.

Limiting cases and assignment. If $\lambda_L^r = 0$, $\Omega_{W,t}^{o,r}$ is common across origins before division by real unit cost. Common $\bar{\tilde{c}}_{i,t}^{M,r}$ therefore implies common $v_i^{W,o}$; substituting that value into equations (24) and (25) makes every production-reallocation difference zero and proves equation (27). Under no equipment trade, substituting $\pi_{ni} = \mathbf{1}\{n = i\}$ makes the nested reallocation sum zero and yields equation (28). The same substitution in equation (22) eliminates direct foreign pass-through. These substitutions do not remove world-experience spillovers or fossil trade.

Finally, under the constrained assignment in Proposition 2, first-order returns are additive across locations, so the return of (I_U, M_H) exceeds that of (M_U, I_H) exactly when $(R_{I_U}^o + R_{M_H}^o) - (R_{M_U}^o + R_{I_H}^o) > 0$, which proves equation (30). The gain of (I_U, M_H) over (I_U, I_H) is $R_{M_H}^o - R_{I_H}^o$, over (M_U, M_H) it is $R_{I_U}^o - R_{M_U}^o$, and over (M_U, I_H) it is the sum of these two gaps, so the two stronger inequalities establish dominance over all three alternatives, and positivity of the two selected returns extends dominance to assignments that omit either policy.

The rule extends to multiple locations. Suppose all returns are evaluated at the same zero-policy reference and objective, every location receives the same infinitesimal fiscal outlay, and exactly K locations must use deployment support. Up to the common positive fiscal-outlay factor, total first-order welfare equals the return from assigning equipment support everywhere plus the sum of Δ_n^o over the K deployment locations. Selecting the K largest gaps is therefore optimal, up to ties. This result is conditional on K; it does not determine how many locations should use either instrument or whether any subsidy should be implemented. At finite common budgets, subtracting the corresponding quadratic terms in equation (20) adds own- and cross-curvature; independently optimized policies instead depend on the return-to-curvature ratios in equation (21).

A.2 Composite Intermediate Goods (Dixit–Stiglitz)

Composite producers in sector j aggregate a unit continuum of varieties $\omega^j \in [0,1]$ through a standard Dixit–Stiglitz technology with elasticity of substitution $\sigma^j > 1$ across varieties. Cost minimization delivers the usual isoelastic variety demands and the CES sectoral price index $P^j_{n,t}$ that appears in the main text; because the aggregator is entirely standard, the familiar expressions are omitted.

A.3 Energy Aggregation and Expenditure Shares

The dual price indices corresponding to the quantity nests in Section 3.3 are

$$P^R_{n,t} = \left[\sum_{r\in\mathcal{R}} \eta^r_n (P^{G,r,u}_{n,t})^{1-\rho}\right]^{1/(1-\rho)},$$
$$P^E_{n,t} = \left[\eta_n (P^R_{n,t})^{1-\psi} + (1-\eta_n)(P^F_{n,t})^{1-\psi}\right]^{1/(1-\psi)}. \tag{A.12}$$

Here $P^{G,r,u}_{n,t}$ is the user price of technology-r electricity paid by the energy aggregator; it differs from the producer price $P^{G,r,p}_{n,t}$ received by the generator only when a generation subsidy is present. The Cobb–Douglas limits apply at unit elasticities. Total spending on energy services is equation (D.4), and its fossil share is

$$\gamma^F_{n,t} = \left[1 + \frac{\eta_n}{1-\eta_n}\left(\frac{P^F_{n,t}}{P^R_{n,t}}\right)^{\psi-1}\right]^{-1}. \tag{A.13}$$

Renewable spending is the complement, and technology-r's share of renewable spending is

$$\frac{\eta^r_n (P^{G,r,u}_{n,t})^{1-\rho}}{\sum_{h\in\mathcal{R}} \eta^h_n (P^{G,h,u}_{n,t})^{1-\rho}}.$$

Thus fossil, solar, and wind expenditures sum exactly to $X^E_{n,t}$. At $\rho = 1$ this share equals η^r_n; only the outer renewable–fossil share responds to relative prices. Equations (A.12) and (D.5) give the corresponding zero-profit and market-clearing conditions. Because the composite shares are properties of the CES aggregator, the same share algebra applies to total energy expenditure $X^E_{n,t}$ regardless of its split across household demand and the energy-input demand of non-energy and equipment producers, all of whom face the common price $P^E_{n,t}$.

A.4 Consumption Welfare and the Ex-post Carbon Account

This subsection defines two distinct objects: the utility-based consumption equivalent and an ex-post present-value expenditure account for carbon damages. Both are dimensionless shares of reference consumption, so their sum is a useful carbon-inclusive reporting measure. In general, however, that

sum is not the derivative of the private utility in equation (1). The finite implementation follows consumption through year T and holds the terminal policy–reference log-consumption gap fixed in the post-2100 welfare tail. Let $d_{n,t} = \log C^p_{n,t} - \log C^0_{n,t}$, $H_T = \sum_{t=1}^{T} \beta^{t-1} + \beta^T/(1-\beta)$, and define

$$\ell_n^{\text{cons}} = \frac{1}{H_T}\left[\underbrace{\sum_{t=1}^{T} \beta^{t-1} d_{n,t}}_{\text{modeled transition through } T} + \underbrace{\frac{\beta^T}{1-\beta} d_{n,T}}_{\text{post-}T\text{ welfare tail}}\right]. \tag{A.14}$$

The first term is the discounted consumption effect along the explicitly solved transition. The second carries the terminal proportional consumption gap beyond T. Dividing by H_T converts their discounted sum into the constant proportional change in consumption that would deliver the same lifetime utility. This is the implemented approximation to the infinite-horizon consumption equivalent in equation (16); it does not assert that the stationary tail equals the exact value implied by every possible state in 2100.

The corresponding level consumption equivalent is

$$\text{CEV}_n^{\text{cons}} = \exp(\ell_n^{\text{cons}}) - 1. \tag{A.15}$$

Country CEV columns report $100\,\text{CEV}_n^{\text{cons}}$. Columns labeled "World log CEV" instead report $100\,\ell_W^R$, defined below.

Define present-value reference expenditure under the same tail convention by

$$\overline{\mathcal{E}}_{n,T} = \underbrace{\sum_{t=1}^{T} \beta^{t-1} \bar{E}_{n,t}}_{\text{modeled reference expenditure}} + \underbrace{\frac{\beta^T}{1-\beta} \bar{E}_{n,T}}_{\text{post-}T\text{ expenditure tail}},$$
$$\overline{\mathcal{E}}_{W,T} = \sum_n \overline{\mathcal{E}}_{n,T}.$$

The denominator is therefore the present value of reference consumption expenditure under the same horizon convention. The world carbon account values damages along the explicitly solved 2023–2100 emissions path and expresses them as a share of this present value. Unlike the consumption calculation, it does not extrapolate the policy-induced emissions gap beyond year T:

$$\ell_W^Z = -\frac{1}{\overline{\mathcal{E}}_{W,T}} \sum_{t=1}^{T} \beta^{t-1} \underbrace{SCC_t (Z_t^p - Z_t^0)}_{\text{change in global dollar damages}}. \tag{A.16}$$

The leading minus sign means that an emissions reduction, $Z_t^p - Z_t^0 < 0$, raises carbon-inclusive welfare. If damage allocation s assigns country n the date-t share $q_{n,t}^{\mathrm{dam},(s)}$ of global dollar damages, with $\sum_n q_{n,t}^{\mathrm{dam},(s)} = 1$, the corresponding national money-metric component is

$$\ell_n^{Z,(s)} = -\frac{1}{\bar{\mathcal{E}}_{n,T}} \sum_{t=1}^{T} \beta^{t-1} q_{n,t}^{\mathrm{dam},(s)} SCC_t \Delta Z_t. \tag{A.17}$$

This is an ex-post allocation of a dollar damage account, not a change in private utility. For national incidence, $\ell_n^{Z,(s)}$ is added to ℓ_n^{cons} as a wealth-equivalent reporting adjustment. It would coincide with the log-consumption equivalent of contemporaneous damages only under additional restrictions—for example, if reference expenditure were constant over time or if the assigned damage bill were proportional to reference expenditure in every period. The reference-expenditure-weighted average of the national accounts reproduces the global account exactly because the damage shares sum to one in every period. The noncarbon global consumption-welfare measure is

$$\ell_W^R = \sum_n \varpi_n \ell_n^{\mathrm{cons}}, \qquad \varpi_n = \frac{\sum_{t=1}^{T} \beta^{t-1} \bar{E}_{n,t} + \frac{\beta^T}{1-\beta} \bar{E}_{n,T}}{\sum_m \left[\sum_{t=1}^{T} \beta^{t-1} \bar{E}_{m,t} + \frac{\beta^T}{1-\beta} \bar{E}_{m,T} \right]}.$$

The carbon-inclusive world measure is $\ell_W^{R,SCC} = \ell_W^R + \ell_W^Z$. The baseline-expenditure weights are held fixed across policies. They provide a transparent incidence summary but do not generally equalize weighted marginal utilities of income at every date. Consequently ℓ_W^R can value pure international redistribution and is not a planner objective based on planner first-order-condition weights. Its local counterpart is equation (A.7).

Appendix B Calibration: Data and Estimation

B.1 Renewable Energy Provisions in the IRA

The deployment instrument is an investment tax credit (ITC) applied to gross renewable-investment cost. It operates on the same investment margin as the traded equipment input. To compare it with a renewable-generation production tax credit (PTC), let $M_{n,t,u}$ be the date-t value of one nominal dollar at date u, and let $R_{n,t,u}^r$ be the marginal date-u generation revenue produced by capacity installed at t. Let $q_{n,t}^r = P_{n,t}^{\mathrm{proj},r} / \Phi_{I,n,t}^r$ denote the pre-credit installed-capital cost. The ITC rate that produces the same marginal investment incentive as the PTC is

$$s_{n,t}^{\mathrm{eq},r} = \frac{\sum_{u>t} M_{n,t,u} \tau_{n,u}^{D,r} R_{n,t,u}^r}{q_{n,t}^r}. \tag{B.1}$$

The numerator is the capitalized PTC payment on one additional unit of installed capacity; $R^r_{n,t,u}$ includes the propagation of that marginal capacity to date u along the reference path. The mapping does not impose fiscal equivalence: a PTC pays on generation, including output from inframarginal capacity, so the PTC experiment records the actual generation-based fiscal bill.

The equipment instrument is an ad valorem subsidy to domestic solar- and wind-equipment output. Counterfactual rates are imposed directly rather than inferred from statutory component-specific schedules. This keeps the comparison between destination deployment support and origin equipment support on a common ad valorem basis. Along an integrated solar supply chain, the Section 45X credits sum to approximately \$0.14 per watt. Relative to the 2022 U.S. module price of about \$0.42 per watt in the IEA PVPS data, this is of the same order as a 30 percent ad valorem subsidy (Internal Revenue Service, 2026; IEA PVPS, 2024). I apply the same rate to wind to keep the two equipment sectors on a common policy scale, not because every wind component receives a statutory 30 percent credit.

The policy application in Section 6.5 combines a 30 percent deployment ITC with a 30 percent equipment subsidy during 2023–2032. I call it a *stylized IRA-scale package*: the 30 percent ITC matches the statutory investment-credit rate, while the equipment rate summarizes Section 45X's component-specific per-unit credits as an ad valorem production wedge. The experiment does not reproduce every eligibility rule, phaseout, bonus, or technology covered by the law; it measures the incidence of policy wedges rather than producing a legislative budget score of the IRA, and its fiscal present value is the endogenous general-equilibrium outlay on the two modeled tax bases.

The central instrument-and-location comparison holds fiscal present value fixed rather than imposing a common rate. Its \$48.35 billion target is the model-implied cost of a 30 percent U.S. equipment subsidy in the tariff-free benchmark. For each instrument–location pair, I solve for the subsidy rate that matches this target in equilibrium. The resulting bills differ from the target by at most \$0.003 billion. This procedure compares allocations generated by a common amount of public spending; it does not linearly rescale policy outcomes and it does not equalize the fiscal burden as a share of national income.

B.2 Geographic and Time Coverage

List of countries and regions. I calibrate the model to 15 countries and regions, aggregated from the OECD Inter-Country Input–Output (ICIO) tables. Eight economies are modeled individually: the United States, China (including Hong Kong and Taiwan), Canada, Mexico, Japan, Korea, India, and Russia. The remaining economies form seven aggregates. Europe comprises the 27 European Union member states plus Norway, Switzerland, and the United Kingdom; Australia–New Zealand comprises Australia and New Zealand; Southeast Asia comprises Bangladesh, Brunei, Cambodia, Indonesia, Laos, Malaysia, Myanmar, the Philippines, Singapore, Thailand, and Vietnam; and South America comprises Argentina, Brazil, Chile, Colombia, Costa Rica, and Peru. Africa comprises

Angola, Cameroon, Côte d'Ivoire, the Democratic Republic of the Congo, Egypt, Morocco, Nigeria, São Tomé and Príncipe, Senegal, South Africa, and Tunisia, and the Middle East comprises Israel, Jordan, Saudi Arabia, Türkiye, and the United Arab Emirates. A synthetic rest of the world collects all remaining economies in the ICIO tables (Belarus, Iceland, Kazakhstan, Pakistan, Ukraine, and the ICIO rest-of-world residual), so that the 15 regions sum exactly to the world.

Sector aggregation. The model has four traded composites: solar equipment, wind equipment, fossil fuel, and a pooled non-energy good. Four exhaustive rules map the source ICIO industries into these composites: B05_06, mining and quarrying of energy-producing products, identifies the fossil-extraction good; C26, computer, electronic and optical equipment, is the solar-equipment input–output proxy; C28, machinery and equipment n.e.c., is the wind-equipment proxy; and every remaining ICIO industry, including background electricity (D), enters the pooled non-energy composite. The replication package contains the full code-level crosswalk. Solar and wind generation are local goods whose 2022 scale is calibrated from observed generation and capacity rather than ICIO industries, and the equipment sectors' bilateral trade matrices are calibrated separately from Comtrade and industry production data (Appendix B.3).

Time horizon. The year 2022 is the historical boundary supplied by the input–output accounts, equipment trade matrices, and renewable-capacity data. The first endogenous choice is made in 2023. The modeled 2023–2100 transition therefore contains 78 choice years, with the treatment of values beyond 2100 described in Section 3.6; some results focus on the years through 2050.

B.3 Equipment Trade and Benchmark Sourcing

Production and deployment shares (Figure 1). Bilateral trade data come from UN Comtrade at the 6-digit HS level. For the 2022 solar matrix I use HS `854142` (PV cells not assembled in modules) and `854143` (PV cells assembled in modules or panels). For wind equipment I combine HS `850231` (wind-powered generating sets) with a destination-specific allocation of HS `850300` (parts of electric motors and generators) and exclude HS `841290`. This is a broad product-boundary proxy: HS `850231` omits important non-generator components, while HS `850300` also includes motor and generator parts used outside wind. The destination-specific share is applied uniformly across origins within a destination; this is an identification assumption rather than an observed origin-specific wind-content measure. I reconcile mirror flows, scale destination totals to installation-based absorption, and impute domestic shipments as the residual. Solar manufacturing shares are 2022 production-origin shares from the IEA PVPS *Trends* data (IEA PVPS, 2024) and provide the solar row-margin targets. Benchmark wind origin shares are the PPML-imputed row sums of this proxy and should not be read as observed turbine-market shares; they place China at about 60% and Europe at about 22%. The separate manufacturer-location panel enters the wind experience stock, not the trade matrix. Deployment shares for both technologies are based on 2022 regional capacity additions.

The solar trade matrix covers cells and modules, whereas the equipment share of installation cost

also includes inverters and electrical balance-of-system. In the absence of comparable bilateral data for those components, I use cell-and-module sourcing shares as a proxy for the origin composition of the full traded solar bundle. This assumption allocates equipment spending across origins; it does not determine equipment's total share of installation cost. The experience stock uses module output in MW rather than the sum of cells and modules, avoiding double counting of cells embodied in modules.

Observed equipment tariffs. Observed bilateral expenditures are duty inclusive, while the model records the policy component of the trade wedge separately. The calibrated observed-policy China–U.S. solar rate is a representative 2022 tariff equivalent of 70.41 percent: 25 percent under Section 301, 14.75 percent under the Section 201 safeguard, a 14.79 percent antidumping separate rate, and a 15.87 percent countervailing-duty non-selected rate. These duties are additive on entered value. The remaining solar rates are 14.75 percent from Japan to the United States, 15 percent from Korea and Southeast Asia to the United States, and 40 percent on imports into India. The Indian model good is module dominant, so I use the module rather than cell rate. For wind, the aggregate-equipment proxies are 25 percent from China and 5 percent from Korea to the United States. Tariff revenue equals $\tau/(1+\tau)$ of duty-inclusive expenditure and accrues to the importer.[1]

U.S. import penetration. For solar, I use the U.S. Energy Information Administration's *Monthly Solar Photovoltaic Module Shipments* to proxy domestic supply and compute import penetration as the share of imports in total U.S. supply. For wind, I use the U.S. Department of Energy's *Land-Based Wind Market Report*, which reports domestic-content shares by major component (nacelle assembly, towers, blades/hubs). I aggregate these component shares using component value-cost weights to obtain a plant-level domestic-content measure and the implied import share.

Figure 2 (equipment prices). The solar module-price series splices PVPS average-price observations through 2012 with the BNEF/ITRPV global c-Si ASP series for 2013–23, cross-checked against PVPS spot/lowest-price observations, and Fraunhofer/PVPS evidence for 2024. For wind equipment, I use IRENA's global onshore wind cost index as a proxy for turbine/equipment price dynamics.

B.4 Evidence on the Inner-Nest Elasticity

Energy production in the model is a two-level CES. As auxiliary evidence on the inner-nest elasticity between solar and wind, ρ, I report the log–log relationship

$$\log\left(\frac{y^{s}_{n,t}}{y^{w}_{n,t}}\right) = -\rho \log\left(\frac{p^{s}_{n,t}}{p^{w}_{n,t}}\right) + FE_n + FE_t + \varepsilon_{n,t}, \tag{B.2}$$

where FE_n and FE_t denote country and year fixed effects. Country fixed effects absorb time-invariant cross-country heterogeneity (e.g., average site quality), while year fixed effects capture

[1] The U.S. solar stack uses Proclamation 10339 (87 FR 7357), the Section 301 List 2 schedule, and the U.S. Department of Commerce's 2019–20 antidumping and 2019 countervailing-duty administrative reviews. India's module duty follows the Ministry of New and Renewable Energy notice effective April 1, 2022. See also Gerarden et al. (2025).

Table B.1: Estimates for Substitution Elasticity (Inner Nest)

	Dependent Variable: $\ln(y^s_{n,t}/y^w_{n,t})$			
	(1)	(2)	(3)	(4)
$\ln(p^s_{n,t}/p^w_{n,t})$	-1.067	-1.173	-0.332	-0.645
	(0.211)	(0.356)	(0.413)	(0.554)
Model F	25.56	56.19	54.00	
1st-stage F				52.46
N	144	144	144	144
Country FE		✓	✓	✓
Year FE			✓	✓
IV				✓

Note: The unit of observation is country-year. Standard errors are clustered at the country level and reported in parentheses.

common time-varying factors such as global technological progress and input-price movements. Equation (B.2) is the demand system dual to the inner CES aggregator in equation (7).

Relative prices and quantities are jointly determined. I therefore also instrument $\log(p^s_{n,t}/p^w_{n,t})$ with wind subsidies that vary across country–year. Because these policies may affect technology choice through channels other than measured relative cost, I treat the IV estimate as suggestive reduced-form evidence rather than structural identification of ρ.

Data. Solar and wind output (y^s, y^w) come from the International Renewable Energy Agency (IRENA). For prices, I construct an unbalanced 15-country sample by combining the levelized cost of electricity for solar and wind (IRENA) with wind deployment subsidies (OECD) for 2010–2021.

Estimates. Table B.1 reports the results. OLS suggests an elastic response to relative prices; adding FE_n and FE_t attenuates the slope. The IV specification with country and year fixed effects gives $\hat{\rho} = 0.645$ (s.e. 0.554). Within the model, ρ governs how the renewable composite rebalances in response to relative cost changes. The estimates support treating solar and wind as imperfect rather than perfect substitutes, but are too imprecise to distinguish the point estimate from one. I therefore adopt the parsimonious Cobb–Douglas case ($\rho = 1$) in the quantitative analysis.

B.5 Learning Calibration and Supporting Estimates

The model requires the total response of equipment costs to world experience and the additional loading on an origin's own experience. I discipline these objects with two complementary moments rather than interpreting either regression as a causal estimate of learning-by-doing. Define

$$g_r \equiv -\frac{\partial \log P^r_t}{\partial \log Q^r_{\text{glob},t}}, \qquad b_r \equiv \frac{\partial \log \hat{S}^r_{n,t}}{\partial \log Q^r_{n,t}}.$$

The structural mapping into equation (12) is

$$\nu^r_{\text{tot}} = \theta^r g_r, \qquad \lambda^r_L = \frac{b_r}{\theta^r g_r} = \frac{b_r}{\nu^r_{\text{tot}}}.$$

Thus the world curve disciplines total price progress, while the exporter panel disciplines the origin-specific component. Because cumulative production can also proxy for R&D, scale economies, and other experience-linked improvements, the resulting estimates discipline the calibration rather than identify a causal learning parameter.

Common timing and units. All experience stocks are beginning-of-year objects. The baseline sets the retention rate to $\mu_Q = 0.99$, so $Q_{t+1} = 0.99Q_t + y_t$; $\mu_Q = 1$ provides the cumulative-output robustness case. The world solar stock used in estimation starts with PVPS module production in 1993 and a conservative 110 MW beginning-1993 seed, equal to cumulative PV capacity in IEA PVPS reporting countries at the end of 1992. The construction of initial transition states begins its historical flow panel in 2000, which changes the benchmark world solar stock-to-flow ratio by about 0.07 percent. Country experience shares are floored at 0.1 percent and renormalized to avoid zero local stocks. Finally, historical 2022 equipment flow is valued at purchaser prices when the initial experience state is built, whereas endogenous transition output is valued at factory-gate resource cost; removing tariffs from that historical flow lowers the world 2022 solar and wind flows by 2.54 and 0.25 percent. Module and polysilicon prices are converted to 2024 dollars using the U.S. CPI.

Country experience is measured by module output. Adding cell and module output would double-count watts embodied in modules; alternative product boundaries provide sensitivity checks.

World price–experience moment. For solar I estimate

$$\log P^{S,\text{real}}_{W,t} = a - g_S \log Q^{S,\text{beg}}_{\text{glob},t} + \beta_P \log P^{\text{poly,real}}_t + u_t,$$

over 2000–24 and use Newey–West standard errors with two lags. Under the benchmark retention specification, the estimate is $\hat{g}_S = 0.4649$ (s.e. 0.023), which implies $\hat{\nu}^S_{\text{tot}} = 2.501$ at $\theta^S = 5.38$ and a 27.5 percent price decline per doubling. Table B.2 reports sensitivity to price units, the pre-sample stock, stock timing, retention, and the sample period. The variation across sample windows, rather than the HAC standard error alone, defines the empirical range used in the learning robustness exercises.

Exporter-capability moment. I recover an exporter-by-year capability $\hat{S}^S_{n,t}$ from a PPML gravity equation with importer-by-year and country-pair fixed effects. The importer-year effects absorb the common equipment frontier and destination demand; the exporter-year term is the model's origin capability, although time-varying origin costs and policy wedges may also enter it. Under equation (2), this term combines $\log A^{M^r}_{n,t} - \theta^r \log c^{M^r}_{n,t}$ with origin policy wedges. After the second-stage year effects absorb common world experience, its structural own-experience loading is $\lambda^r_L \nu^r_{\text{tot}}$; the empirical coefficient equals that loading only if the remaining origin-cost and policy terms are conditionally

Table B.2: Sensitivity of the Solar Price–Experience Estimate to Timing and Price Measurement

Specification	ν^S_{tot}	LR^S
Real price, beginning-year stock, $\mu_Q = 0.99$	2.501	27.5%
Nominal prices	2.256	25.2%
Zero beginning-1993 seed	2.460	27.2%
Cumulative-output case $\mu_Q = 1$	2.504	27.6%
End sample in 2022	2.457	27.1%
Add post-2013 source-splice indicator	2.470	27.3%

Notes: World solar modules, 2000–24. Regressions control for the polysilicon price and use beginning-of-year experience unless stated otherwise. Real prices are in 2024 dollars. $\nu^S_{\text{tot}} = 5.38 g_S$ and $LR^S = 1 - 2^{-g_S}$.

unrelated to experience. I then estimate

$$\log \hat{S}^S_{n,t} = b_S \log Q^{S,\text{beg}}_{n,t} + \delta_n + \delta_t + \varepsilon_{n,t}.$$

Using module production over 2005–23 under the benchmark retention specification gives $\hat{b}_S = 0.308$ (country-clustered s.e. 0.092). Dropping China gives 0.293. Omitting the observed pre-2000 origin history lowers the estimate to 0.256, while using the maximum of cell and module output gives 0.306. The estimate is therefore not sensitive to this alternative product boundary.

Figure B.1 plots the exporter-year capability estimates for the major producing origins. Because the United States is normalized to zero in every year, the paths show changes in relative exporter capability rather than movements in the common world frontier.

In the Comtrade matrices, rows are importing destinations and columns are exporting origins. Let $\omega_{nd,0}$ be origin n's average annual share of exports sold to destination d during 2000–04, normalized so that $\sum_d \omega_{nd,0} = 1$, and let $F_{d,t}$ indicate a pre-specified generous solar feed-in-tariff or demand-support window in destination d. Define

$$f^{\text{FIT}}_{n,t} = \sum_d \omega_{nd,0} F_{d,t}, \qquad \mathcal{F}^{\text{FIT}}_{n,t} = \sum_{u=2000}^{t} f^{\text{FIT}}_{n,u}, \qquad z^{\text{FIT}}_{n,t} = \log \mathcal{F}^{\text{FIT}}_{n,t}. \tag{B.3}$$

I use $z^{\text{FIT}}_{n,t}$ to instrument module-only beginning-of-year cumulative output in the 2005–23 capability regression; this corroborating specification uses $\mu_Q = 1$. The policy window is used instead of realized installations, which are equilibrium outcomes, and is coded from REN21, the IEA policy database, and country policy histories. The IV estimate is 0.313 (s.e. 0.119), with a first-stage $F = 19.8$ and a 95 percent Anderson–Rubin acceptance interval of $[0.070, 0.589]$. Dropping China gives 0.319. Leave-one-destination-out coefficients range from 0.236 to 0.352, although omitting Japan lowers the first-stage F to 4.9. The design provides corroborating rather than stand-alone causal evidence.

China's 2018 policy reversal (descriptive). China's May 2018 "531" reform sharply reduced do-

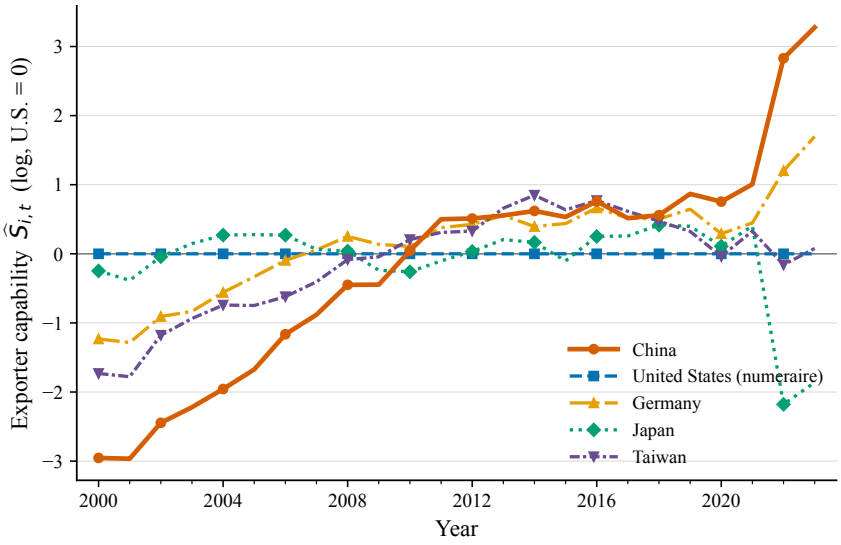


Figure B.1: Solar Exporter Capability

Notes: Exporter-year capability $\hat{S}^S_{n,t}$ is recovered from the PPML gravity equation with importer-year and country-pair fixed effects. The United States is normalized to zero within each year. The estimates therefore measure relative exporter capability and may also incorporate time-varying origin costs and policy wedges.

mestic solar deployment support. China's relative exporter capability did not decline after the reform: it rose from 0.56 in 2018 to 1.01 in 2021, while China's share of world solar exports rose from 32 to 42 percent. This before–after pattern is consistent with persistence in accumulated production capability after domestic demand support is withdrawn. It is not used for identification or calibration, however, because the single-country episode has no control group and coincided with other changes in global solar markets.

Combining the two unrounded moments gives

$$\hat{\lambda}^S_L = \frac{0.3076}{2.5011} = 0.1230 \qquad \text{(delta-method s.e. } 0.038\text{)}.$$

This parameter is an own-stock loading in addition to a common frontier; it is not literally the fraction of all social learning benefits captured by the origin.

Wind. The wind data do not support the same decomposition with comparable precision. The world price series is short and its slope varies with the price concept; the country-production panel maps manufacturer market shares to headquarters and is therefore a noisy proxy for production location. The own-country coefficient is about 0.6 but has only six producer clusters and reverses sign when China is removed. The calibration therefore sets $(\nu^W_{\text{tot}}, \lambda^W_L) = (1.00, 0.50)$. Appendix Table C.3 reports a joint lower envelope $(0.50, 0.25)$ and upper envelope $(1.20, 0.75)$, paired with the solar ranges described above. These values are calibrated rather than estimated and are not comparable to

Table B.3: Learning calibration and empirical moments

	ν^r_{tot}	b_r	λ^r_L	LR^r
Solar, estimated/mapped	2.501	0.308	0.123	27.5%
	(0.123)	(0.092)	(0.038)	
Wind, calibrated value	1.00	0.50	0.50	14.6%

Note: $b_r = \lambda^r_L \nu^r_{\text{tot}}$. Solar standard errors are Newey–West, country-clustered, and delta-method, respectively. For wind, $(\nu^W_{\text{tot}}, \lambda^W_L)$ are calibration choices; b_W and LR^W are implied values rather than separately estimated moments.

the solar moments.

Quantitative parameters. The model sets $(\nu^S_{\text{tot}}, \nu^W_{\text{tot}}) = (2.501, 1.00)$, $(\lambda^S_L, \lambda^W_L) = (0.123, 0.50)$, and $\mu_Q = 0.99$. Beginning-of-2022 stocks accumulate physical production through 2021 under the corresponding retention specification and satisfy $Q^r_{\text{glob},2022} = \sum_n Q^r_{n,2022}$. A trade-elasticity sensitivity must preserve the empirical sufficient statistics: changing θ^r requires recomputing both $\nu^r_{\text{tot}} = \theta^r g_r$ and $\lambda^r_L = b_r / (\theta^r g_r)$, followed by a fresh equilibrium solve. The no-learning and origin-specific-learning-only cases isolate the mechanism, $\mu_Q = 1$ provides the stock-retention sensitivity, and Table C.3 reports the joint parameter envelope.

Data construction. Country module production comes from the annual *Trends in Photovoltaic Applications* reports of the International Energy Agency Photovoltaic Power Systems Programme (IEA PVPS, 2024). The world module-price series uses PVPS average-price observations through 2012, the BNEF/ITRPV global c-Si ASP series for 2013–23 with PVPS cross-checks, and Fraunhofer/PVPS evidence in 2024. The world series covers 2000–24; the production pre-sample covers 1993–99, and the 1992 cumulative-capacity observation supplies a conservative earlier stock. U.S. CPI series CPIAUCSL converts module and polysilicon prices to 2024 dollars. Bilateral equipment trade comes from UN Comtrade; the gravity step uses the full bilateral matrix and the FIT exposure uses the 2000–04 origin-column-normalized destination shares in equation (B.3). The country experience flow is module MW, not the sum of cell and module MW.

For wind, global installations and price concepts come from IRENA and LBNL. Country manufacturing is a proxy constructed from manufacturer market shares and headquarters locations, scaled to global output. This proxy is distinct from deployment but remains too concentrated in a small number of producers to support a precise local-share estimate; it therefore provides only directional evidence.

B.6 Intermittency, Value Factors, and Annual Aggregation

The model aggregates generation and energy use annually and therefore omits storage operation, curtailment, hourly transmission constraints, and within-day and seasonal dispatch. The renewable–fossil substitution elasticity summarizes these integration costs in reduced form. A fixed annual elasticity cannot steepen endogenously as renewable penetration grows, so the model may understate

late-transition integration costs if value factors decline more sharply than the annual CES implies, or overstate them if storage becomes sufficiently cheap. Models with explicit dispatch and storage include Helm and Mier (2021) and Gentile (2024).

The demand-side content of intermittency at annual frequency is a declining value factor. Because solar and wind produce when the resource is available rather than when power is most valuable, the market value of a megawatt-hour of renewable output falls relative to the average power price as renewable penetration rises; this is the central lesson of the literature comparing intermittent and dispatchable generation (Joskow, 2011) and documenting the declining market value of variable renewables (Hirth, 2013). In the model this penalty is carried by the outer nest of the energy aggregator (8). With a finite elasticity ψ between the renewable composite and fossil energy, absorbing a larger renewable share requires the relative price of renewable output to fall, so the marginal value of renewable generation declines with penetration. The range $\psi \in \{2, 4\}$ spans stronger and weaker versions of this penalty.

Within the renewable bundle, the baseline calibration sets the inner solar–wind nest to Cobb–Douglas ($\rho = 1$), which treats the complementarity between the two technologies in reduced form. Renewable expenditure divides between solar and wind in fixed proportions, so the model imposes neither the additional penalty on an unbalanced mix that gross complements ($\rho < 1$) would generate nor the near-perfect substitutability that would let a single technology serve the entire renewable share at no cost. The counterfactual analysis holds $\rho = 1$. Observed 2022 generation per unit of capacity embeds average capacity factors in the country–technology coefficients z_n^r.

B.7 Measuring the Initial Productive Stocks

The quantitative model has three capital states: solar capacity, wind capacity, and non-energy capital. Fossil extraction instead uses the within-period supply curve in equation (D.2); there is no fossil-capital stock.

Solar and wind. I aggregate IRENA installed capacity in 2021 and 2022 to the model regions. End-of-2022 capacity is the renewable state inherited in 2023. For each technology, the model-consistent 2022 gross capacity flow is end-of-2022 capacity less undepreciated end-of-2021 capacity; it includes replacement and is not the observed net addition. Baseline renewable-investment expenditure is inferred from equipment expenditure and the calibrated equipment share, equation (4). I choose χ_n^r so that the capital law in equation (10) reproduces the observed 2021–22 capacity boundary. For wind in Africa, Australia–New Zealand, and the rest of the world, I use the world-average installed cost when matrix-implied spending is near zero. Thus the historical boundary fixes both the initial stock and the units in which renewable investment augments capacity; it does not impose a 2022 Euler equation.

Non-energy capital. I measure 2022 non-energy investment as total gross fixed capital formation less solar and wind renewable-investment expenditure. Its baseline rental income is the capital share

of non-energy value added, $\overline{\mathcal{R}}_n^D = \gamma_K \gamma_n^V \bar{Y}_n^O$. Because non-energy capital uses the same adjustment technology as renewables, I initialize this block by combining a unit marginal installed-capital price with a stationary 2022 law of motion:

$$K_{n,0}^D = \frac{I_{n,0}^D}{\alpha \delta_D}, \qquad \chi_n^D = \delta_D \left(\frac{K_{n,0}^D}{I_{n,0}^D} \right)^{\alpha}.$$

These two equations imply $\Phi_I^D = 1$ and $q_{n,0}^D = 1$ at the reference point, and $K_{n,t+1}^D = K_{n,t}^D$ whenever real I^D remains at its 2022 value. Rental income determines $r_{n,0}^D = \overline{\mathcal{R}}_n^D / K_{n,0}^D$. This initialization does not impose an Euler condition on observed 2022 investment; from 2023 onward, non-energy investment is chosen endogenously under equation (14).

Fossil supply. The 2022 origin value of fossil production fixes the level of the reduced-form supply curve. I impose the supply elasticity $\varepsilon_F = 1$ and derive the labor-payment share $\eta_F = \varepsilon_F / (1 + \varepsilon_F) = 0.5$ under the competitive labor–fixed-resource interpretation; the remaining share accrues to the domestic fossil resource. The code stores the two parameters separately, which permits the sensitivity exercises to vary the supply response while holding the benchmark income split fixed.

B.8 Emissions Accounting

Flow emissions in the model are linear in fossil-energy use, $Z_t = \sum_n \epsilon_n F_{n,t}^d$ (equation (15)). Physical emissions and country intensities are constructed as follows.

Physical anchor. The Global Carbon Project's fossil-carbon accounts give world emissions of 37.1 $GtCO_2$ in 2022 (Friedlingstein et al., 2026). I aggregate national combustion and cement emissions to the model regions and assign the world-total residual, including international bunkers, to the rest of the world. Non-CO_2 gases are outside the physical measure.

Country intensities. The intensity for region n is the ratio of territorial emissions to model fossil use in the base year,

$$\epsilon_n = \frac{E_{n,2022}^{\text{GCP}}}{\bar{F}_{n,2022}^d},$$

where $\bar{F}_{n,2022}^d$ is baseline real fossil absorption—destination expenditure on the fossil-energy sector deflated by its user price—measured at the burning site and inclusive of non-electricity fossil demand. Emissions are thus attributed to the region where the fossil energy is used, matching the territorial basis of the anchor, rather than to the region extracting the fuel.

Time variation. The equilibrium calculation determines percentage changes in fossil use and hence in emissions. For welfare reporting, the baseline normalization fixes world emissions and the world-income denominator at their 2022 levels throughout the horizon: $\bar{Z}_t = \bar{Z}_{2022}$ and $\bar{Y}_t^W = \bar{Y}_{2022}^W$. Policy-induced emissions hats vary over time, so the physical deviation used in the account is $Z_t' - \bar{Z}_t = \bar{Z}_{2022}(\widehat{Z}_t - 1)$. This normalization converts the model's percentage deviations into tons, but

it is not a projected baseline emissions or GDP path. The results in Section 6 that exclude carbon damages do not depend on this physical scaling.

B.9 Social Cost of Carbon

In the reference and green-subsidy experiments, carbon damages do not enter private objectives or equilibrium allocations. The tables that exclude carbon damages are paired with the ex-post account in equations (A.16) and (A.17), under which any SC–CO_2 path can revalue the solved emissions deviations. This revaluation changes welfare accounting but not prices, quantities, investment, or emissions. Appendix C.8 instead imposes a fossil-use tax and re-solves the equilibrium. World noncarbon incidence uses the present-value baseline-expenditure weights described in Appendix A.4.

The global damage price is the flat \$210 per tCO_2 benchmark described in Section 5. It applies to the territorial fossil CO_2 measure above; non-CO_2 gases are not folded into either the emissions quantity or the price. I do not report country-specific SCCs. Such an exercise requires a separate and contestable allocation of a fixed global damage account across recipients, changes no equilibrium object, and does not affect the global carbon benefit or the paper's policy-assignment comparisons. The generic national allocation in Appendix A.4 is retained only so that any chosen recipient shares can be applied transparently.

Appendix C Mechanisms and Robustness

This appendix distinguishes the tariff-free exercises that accompany the central policy-assignment comparison from observed-tariff extensions used for the IRA application and parameter sensitivities. Every table states the relevant regime. Equilibrium robustness exercises re-solve the full transition; ex-post carbon revaluations instead hold solved consumption and emissions paths fixed. Appendix C.8 is different: it introduces a fossil-use tax and re-solves the equilibrium. Unless noted otherwise, policies operate during 2023–2032, the learning specification sets $\mu_Q = 0.99$, and consumption equivalents exclude climate damages.

C.1 Welfare Decomposition

Table C.1 decomposes the equal-budget welfare changes in Table 3, separating world-experience and origin-specific learning. For each policy, I evaluate the marginal welfare components at intermediate subsidy rates and integrate them as the subsidy rises from zero to the reported rate. The international-price term values changes in export receipts and import costs using pre-subsidy origin costs, and the two learning terms value the productivity effects of world and origin-specific experience. The policy term is instrument specific: the deployment row reports the investment-margin ITC wedge, whereas the equipment rows report purchaser-price benefits net of marginal fiscal outlay. Carbon is excluded

Table C.1: Welfare Decomposition in the Tariff-Free Benchmark

Policy	Intl. prices	World exp.	Origin exp.	ITC wedge	Equip. incidence	Other	Total
Panel A: Global consumption welfare							
U.S. deployment subsidy	−0.08	7.46	1.91	−4.23	—	0.70	5.75
U.S. equipment subsidy	0.09	7.71	1.48	—	−10.94	0.56	−1.10
China deployment subsidy	−0.07	8.16	2.18	−0.93	—	0.67	10.01
China equipment subsidy	−0.13	13.23	3.29	—	−0.45	0.89	16.82
Panel B: Policy-country consumption welfare							
U.S. deployment subsidy	−2.01	2.86	3.39	−15.33	—	1.43	−9.65
U.S. equipment subsidy	9.85	2.94	7.50	—	−45.19	1.01	−23.89
China deployment subsidy	22.62	20.92	8.14	−5.36	—	2.57	48.89
China equipment subsidy	30.78	33.54	13.53	—	−42.05	2.60	38.40

Notes: Entries are in 10^{-4} percentage points of consumption-equivalent welfare. Rates are 12.86 percent for U.S. deployment, 30 percent for U.S. equipment, 3.12 percent for China deployment, and 4.70 percent for China equipment. Channels are 12-point Gauss–Legendre path integrals. "International prices" use pre-subsidy origin costs; for equipment subsidies, the statutory $(1 - s^M)$ purchaser-price wedge is recorded instead, net of marginal fiscal outlay, under "Equip. incidence" and includes foreign buyer gains. The ITC wedge $-sP^{\text{proj}}\,dI$ is not comparable to this equipment-policy term. "Other" is the signed unallocated residual. Its largest absolute value is 6.1 percent of the sum of absolute channel contributions; quadrature error is below 10^{-11}. Climate damages are excluded.

and reported separately in the main text.

Learning is positive in every row. For U.S. deployment support, the world-experience component accounts for about four-fifths of the total learning contribution, and total learning exceeds the ITC wedge. The two U.S. policies generate nearly identical world-experience contributions at the displayed precision, but the difference is smaller than either row's unallocated residual. Chinese equipment support produces a world-experience contribution about 60 percent larger than Chinese deployment support. International-price incidence is close to zero globally but substantial in the policy-country rows. The table therefore supports the importance of learning and national incidence without treating the two instrument-specific policy columns as a common measure of efficiency cost.

Destination allocation of buyer-price benefits. The equipment-policy term contains a direct buyer-price benefit. Let $a \in [0,1]$ scale the reported equipment-subsidy path at origin i, so that $s_{i,t}^{M,r}(a) = a s_{i,t}^{M,r}(1)$ and all other equipment-subsidy rates remain zero. The discounted benefit delivered to destination n is

$$\mathcal{P}_{n\leftarrow i}^{\text{buy}} \equiv -\int_0^1 \sum_{t=1}^{T} \beta^{t-1} \sum_{r\in\mathcal{R}} V_{ni,t}^{M^r}(a) \frac{\partial \log\left[1 - s_{i,t}^{M,r}(a)\right]}{\partial a}\, da. \tag{C.1}$$

Here $V_{ni,t}^{M^r}(a)$ is tariff-exclusive bilateral equipment expenditure at the intermediate equilibrium. Its variation with a incorporates endogenous sourcing and quantities; changes in origin costs and productivity remain in the international-price and learning channels. Define $\mathcal{P}_i^{\text{dom}} \equiv \mathcal{P}_{i\leftarrow i}^{\text{buy}}$ and $\mathcal{P}_i^{\text{for}} \equiv \sum_{n\neq i} \mathcal{P}_{n\leftarrow i}^{\text{buy}}$.

Table C.2: Destination Allocation of Equipment-Subsidy Buyer-Price Benefits

	Fiscal PV ($bn)	Buyer-price benefit (PV $bn)			Foreign incidence	
Policy		Total	Policy country	Foreign	Share (%)	Cents/$
U.S. equipment subsidy	48.350	25.394	22.137	3.258	12.83	6.74
China equipment subsidy	48.348	45.810	31.373	14.437	31.51	29.86

Notes: Dollar entries are present values in 2022 dollars, discounted at $\beta = 0.96$. Equation (C.1) is integrated over the same 12-point policy path as Table C.1. "Foreign share" divides foreign by total buyer-price benefit; "cents/$" divides foreign benefit by fiscal outlay. These are direct effects of the policy wedge, not welfare effects or estimated pass-through. The difference between fiscal outlay and measured buyer benefit is not decomposed here.

Table C.3: Equipment-Learning Robustness with Observed Tariffs

	Lower learning	Observed-tariff benchmark	Upper learning
Panel A: Fiscal-matched policy rate (percent)			
U.S. deployment subsidy	14.66	14.58	14.55
U.S. equipment subsidy	30.46	30.00	29.39
China deployment subsidy	4.71	4.68	4.69
China equipment subsidy	6.81	6.75	6.72
Panel B: Global consumption CEV (percent)			
U.S. deployment subsidy	0.0004	0.0008	0.0019
U.S. equipment subsidy	−0.0004	0.0002	0.0015
China deployment subsidy	0.0008	0.0013	0.0027
China equipment subsidy	0.0015	0.0023	0.0045
Panel C: Policy-country consumption CEV (percent)			
U.S. deployment subsidy	−0.0012	−0.0008	−0.0001
U.S. equipment subsidy	−0.0043	−0.0035	−0.0021
China deployment subsidy	0.0062	0.0086	0.0136
China equipment subsidy	0.0026	0.0064	0.0150

Notes: Observed 2022 equipment tariffs are retained. The three columns use solar pairs $(\nu^S_{\text{tot}}, b_S)$ of $(2.04, .221)$, $(2.50, .308)$, and $(3.52, .408)$ and wind pairs $(\nu^W_{\text{tot}}, \lambda^W_L)$ of $(.50, .25)$, $(1, .50)$, and $(1.20, .75)$. Rates are fiscal-matched; climate damages are excluded.

C.2 Learning Robustness and Technology-Specific Returns

Table C.3 recomputes the four observed-tariff equal-budget policies under joint lower and upper learning-parameter scenarios. Within each scenario, the policy rate is solved again to attain the same fiscal present value rather than held fixed. Chinese equipment support has the largest world effect throughout, while U.S. equipment support is close to zero and becomes negative under lower learning. Both U.S. policies reduce U.S. consumption welfare throughout the range; in the upper-learning scenario, China's domestic ranking switches from deployment to equipment support.

Separate observed-tariff counterfactuals for U.S. deployment support, U.S. equipment support, and Chinese equipment support shut down individual learning mechanisms while holding the parameters fixed. Shutting down learning makes the global effect negative for these three policies. Re-

Table C.4: Marginal Global Welfare per Fiscal Dollar: Solar and Wind

Policy	Solar	Wind	Solar/wind
U.S. deployment subsidy	2.783	1.859	1.50
U.S. equipment subsidy	3.513	1.763	1.99
China deployment subsidy	2.664	1.830	1.46
China equipment subsidy	4.711	2.385	1.97

Notes: The solar and wind columns are in 10^{-5} percentage points of global consumption-equivalent welfare per billion 2022 dollars of fiscal cost, obtained from Richardson-extrapolated small policy changes in the tariff-free benchmark. The underlying calculations use 1 and 2 percent rate changes and show no detectable curvature at this scale. Each row varies one technology-specific instrument while all other green policies remain zero. These are marginal returns, not finite-policy welfare effects.

taining only origin-specific learning leaves only Chinese equipment support slightly positive. These counterfactuals provide a finite observed-tariff counterpart to the learning-scope distinction in Corollary 2: when common experience is removed, the calibrated value of production location is sufficient to preserve a small gain only for Chinese equipment support among these three policies.

Table C.4 compares the technologies in marginal fiscal experiments. Because it divides an infinitesimal welfare change by its fiscal cost, it is the closest numerical counterpart to the return R_z^G in equation (19). For all four combinations of country and instrument, the global return per fiscal dollar is higher for solar than for wind. The ratio is about 1.5 for deployment subsidies and about 2 for equipment subsidies. Because solar and wind differ in learning, trade, initial production, deployment, and energy productivity, this comparison summarizes the full calibration rather than isolating any one parameter. It also concerns marginal policy changes and need not reproduce the ranking of finite equal-budget policies.

C.3 Structural Sensitivity

Table C.5 varies the outer energy elasticity and the two distinct fossil elasticities. Greater substitution between renewable and fossil energy amplifies both policies without changing their ranking. The fossil trade elasticity has little effect on global welfare. The fossil supply elasticity is more important because it directly governs the fossil-price response: at $\varepsilon_F = 2$, the already small effect of the U.S. deployment subsidy becomes slightly negative, whereas the effect of Chinese equipment support remains positive throughout the evaluated range.

The small U.S. aggregate effects are parameter-sensitive. In policy-country results not displayed in Table C.5, U.S. consumption falls by 0.0002 percent at $\psi = 2$ and by 0.0001 percent at $\psi = 3$, but rises by 0.0002 percent at $\psi = 4$. Changes in ε_F likewise reverse both the U.S. and world signs for that experiment. I therefore interpret U.S. deployment support as approximately neutral in consumption terms and emphasize its cross-country incidence. By contrast, the world and Chinese effects of Chinese equipment support remain positive over the evaluated ψ and ε_F ranges.

Table C.5: Structural Sensitivity with Observed Tariffs

Parameter setting	U.S. deployment subsidy	China equipment subsidy
Panel A: Renewable–fossil elasticity		
$\psi = 2$	0.0003	0.0016
$\psi = 3$	0.0007	0.0035
$\psi = 4$	0.0012	0.0060
Panel B: Fossil supply elasticity		
$\varepsilon_F = 0.5$	0.0013	0.0057
$\varepsilon_F = 1$	0.0007	0.0035
$\varepsilon_F = 2$	−0.0001	0.0011
Panel C: Fossil trade elasticity		
$\theta_F = 3.16$	0.0007	0.0035
$\theta_F = 8$	0.0006	0.0035
$\theta_F = 15$	0.0006	0.0035

Notes: Entries are world log CEVs for a 10 percent policy with observed 2022 equipment tariffs retained; one parameter block varies at a time.

Table C.6: Policy Duration and PTC Mapping with Observed Tariffs

Experiment	Policy horizon	World log CEV (percent)	Fiscal present value ($bn)	Discounted emissions ($MtCO_2$)
Panel A: Duration of a U.S. investment credit, $\mu_Q = 1$				
ITC	5 years	0.0004	17.3	−160
ITC	10 years	0.0006	30.7	−262
ITC	20 years	0.0008	49.2	−379
ITC	Permanent	0.0009	77.8	−501
Panel B: ITC and PTC comparison, $\mu_Q = 0.99$				
ITC	10 years	0.0007	30.8	−264
Per-technology PTC	10 years	0.0003	31.7	−83
ITC	Permanent	0.0009	77.4	−503
Per-technology PTC	Permanent	0.0009	140.1	−623

Notes: Observed 2022 equipment tariffs are retained. Panel A uses $\mu_Q = 1$; Panel B uses $\mu_Q = .99$. PTC rates use the capitalized marginal return, $q = P^{\text{proj}}/\Phi$. Inframarginal PTC payments make its fiscal cost differ from the ITC.

C.4 Policy Duration and the PTC Mapping

Panel A of Table C.6 compares policy durations under the cumulative-experience case $\mu_Q = 1$; relative to the benchmark, it is therefore a joint duration-and-retention exercise rather than a duration-only sensitivity. Panel B returns to the benchmark retention rate $\mu_Q = .99$ to compare ITC and PTC incentives. As the window expands from five years to a permanent policy in Panel A, the global CEV rises from 0.0004 to 0.0009 percent, while the present-value fiscal bill rises from $17.3 billion to $77.8 billion. Welfare therefore rises much less than the fiscal bill over the same experiments. Appendix C.9 shows that the material policy effects are insensitive to the change in retention.

For permanent policies, the appropriately capitalized PTC and ITC generate nearly identical global-welfare effects, although the PTC's fiscal bill is substantially larger. The equivalence breaks for a ten-year calendar window. The finite-window PTC produces less than half the effect of the ITC with a similar fiscal present value because installations completed late in the window receive few subsidized operating years and existing capacity receives inframarginal payments. The PTC and ITC are therefore not interchangeable temporary policies.

C.5 Unilateral Policy-Rate Profiles

Table C.7 reports selected points from the tariff-free unilateral grids behind Section 6.3. U.S. welfare peaks at about 3 percent for deployment support and about 9 percent for equipment support. Chinese welfare peaks near 16 percent for both instruments, with the peak bracketed by rates of 14 and 18.5 percent. Global welfare under Chinese deployment support turns down after 18.5 percent, whereas global welfare under Chinese equipment support continues to rise through the highest evaluated rate of 26 percent. These are one-instrument profiles with all other green policies set to zero.

C.6 Source Restrictions and Tariff Alternatives

The main source-restriction exercise permanently removes U.S. access to solar equipment produced in China and Southeast Asia. It is a broad decoupling experiment rather than a literal representation of a particular trade rule: the country aggregation does not identify ownership, component content, or transshipment. The restriction creates no tariff revenue or quota rent, and learning follows the production that remains after sourcing adjusts.

Table C.8 compares the main restriction with five alternatives: a China-only restriction, a 2023–2032 restriction, a finite tariff under two revenue allocations, and the China–Southeast Asia restriction with the 2022 tariffs removed. In each regime, policy rates are solved again to attain the common fiscal target. The China-only restriction has little effect because direct observed-tariff U.S. solar sourcing from China is small; most of the main restriction's cost comes from excluding Southeast Asia. The temporary restriction preserves the signs and has smaller effects than the permanent case.

The form of the restriction affects incidence. A finite tariff produces smaller losses than a prohibitive barrier when the United States retains the revenue. Reassigning that revenue to exporters leaves the real allocation and global consumption welfare unchanged but shifts the country effects. Within the China–Southeast Asia restriction, removing pre-existing tariffs does not change the signs of the incremental policy effects. Adding energy use in production leaves the tariff-free restriction's global-welfare loss essentially unchanged: the permanent restriction lowers global consumption welfare by 0.0028 percent, within 0.00001 percentage points of the estimate when production energy use is omitted.

Across the policy variants, the global welfare gain from deployment support exceeds that from

Table C.7: Tariff-Free National and Global Welfare across Policy Rates

Policy	Rate (percent)	CEV (10^{-4} percentage points) Global	National
U.S. deployment subsidy	0.50	0.2985	0.1548
	1.00	0.5932	0.2783
	1.50	0.8839	0.3692
	2.00	1.1702	0.4263
	2.75	1.5912	0.4458
	3.50	2.0012	0.3818
	6.00	3.2757	−0.4967
	10.00	4.9400	−4.4417
	13.00	5.7826	−9.9651
U.S. equipment subsidy	2.0	0.1574	0.4534
	5.0	0.3740	0.9503
	7.0	0.5010	1.1235
	9.0	0.6099	1.1354
	11.5	0.7140	0.8669
	14.0	0.7729	0.2003
	19.0	0.7016	−2.7996
	25.0	0.1009	−10.9483
China deployment subsidy	5.0	15.6122	75.3277
	8.0	23.6323	111.2676
	12.0	32.0566	144.1142
	14.0	35.0169	152.3774
	16.0	36.9461	154.0107
	18.5	37.6435	145.1383
	21.0	36.0737	121.9610
	26.0	24.2628	21.5211
China equipment subsidy	6.0	21.5054	47.4648
	9.0	32.2811	65.0878
	12.0	42.9254	77.1716
	14.0	49.8762	81.4579
	16.0	56.6488	82.1505
	18.5	64.7660	77.0708
	21.0	72.3558	64.1333
	26.0	85.1917	8.0451

Notes: Each row varies one policy in isolation in the tariff-free benchmark. Consumption equivalents exclude climate damages. The reported points bracket all four national peaks. The Chinese equipment-subsidy grid does not bracket a global-welfare peak, so no global optimum is reported for that instrument.

equipment support, and the U.S. consumption loss is smaller under deployment, although the global gap narrows under broad restrictions. China gains from both policies in every reported regime.

Table C.8: Source-Restriction Robustness Variants

Trade regime	U.S. policy	Rate (percent)	World log CEV	U.S. CEV (percent)	China CEV	Discounted CO_2 ($MtCO_2$)
Panel A: Total effect relative to the factual no-policy transition						
China and Southeast Asia excluded (reference)	None	—	−0.0022	−0.0065	−0.0004	140
China excluded only	None	—	≈ 0	−0.0001	≈ 0	0
Windowed 2023–2032, reopens after	None	—	−0.0009	−0.0025	−0.0005	74
Finite 50% tariff, revenue to U.S.	None	—	−0.0012	−0.0025	−0.0012	125
Finite 50% tariff, rent to exporters	None	—	−0.0012	−0.0054	−0.0007	125
China/SEA restriction, 2022 tariffs removed	None	—	−0.0028	−0.0071	−0.0025	264
Panel B: Policy effect relative to no policy under the same trade regime						
China excluded only	Deployment subsidy	14.58	0.0009	−0.0008	0.0052	−419
China excluded only	Equipment subsidy	29.99	0.0002	−0.0034	0.0060	−408
Windowed 2023–2032	Deployment subsidy	15.40	0.0008	−0.0012	0.0056	−428
Windowed 2023–2032	Equipment subsidy	28.54	0.0007	−0.0024	0.0067	−451
Finite 50% tariff	Deployment subsidy	14.99	0.0008	−0.0007	0.0052	−417
Finite 50% tariff	Equipment subsidy	28.86	0.0005	−0.0032	0.0070	−455
China/SEA restriction, 2022 tariffs removed	Deployment subsidy	14.96	0.0008	−0.0013	0.0054	−409
China/SEA restriction, 2022 tariffs removed	Equipment subsidy	28.96	0.0007	−0.0019	0.0056	−413

Notes: Panel A compares each restriction with the factual transition. Panel B compares policy with no policy within each regime; rates are solved to attain the same fiscal target in each trade regime. The windowed case reopens after 2032; negative emissions entries denote reductions.

Table C.9: Local Pigouvian Equipment Rates on the Observed-Tariff Path

Origin	Technology	2023 rate $\mu_Q = .99$	Output-weighted 2023–2032 rate $\mu_Q = .99$	$\mu_Q = 1$
United States	Solar	43.0	38.2	38.3
China	Solar	43.9	39.0	39.2
United States	Wind	15.2	13.7	13.4
China	Wind	18.2	16.1	15.9

Notes: Entries are ad valorem percentages; window averages use origin output weights. Rates value future learning on the observed-tariff no-policy path through 2100. They are not a solved planner allocation and exclude the carbon tax.

C.7 A Local Pigouvian Benchmark

Lemma 1 characterizes the equipment subsidy that prices the learning externality along a planner allocation. For a scale comparison, I instead evaluate that marginal benefit on the no-policy reference path. I use the expenditure-normalized social discount factor

$$m^W_{t,u} = \beta^{u-t} \frac{E^W_t}{E^W_u}, \qquad E^W_t = \sum_n P^C_{n,t} C_{n,t},$$

apply the world and origin-specific experience weights in equation (12), and divide the resulting per-unit learning benefit by the origin unit cost. Table C.9 reports both the first-period rate and the origin-output-weighted average over the policy window.

The benchmark puts the policy experiments in perspective. A common 30 percent equipment rate is below the local solar-learning wedge but above the wind wedge, so it cannot be read as a uniform approximation to the first best. The similarity across the two experience-retention assumptions also

Table C.10: Carbon Taxes in Equilibrium at the Flat $210 Valuation

Experiment	Noncarbon world CEV (%)	U.S. consumption CEV (%)	Carbon benefit (% / $bn)	World total (%)	Emissions ($GtCO_2$)	Foreign offset (%)
World tax, $210/t	−2.60	—	+6.69 / 121,026	+4.09	−576.3	—
+ local-Pigouvian learning	−2.79	—	+6.99 / 126,478	+4.20	−602.3	—
+ IRA-scale package	−2.66	—	+6.75 / 122,170	+4.09	−581.8	—
U.S. tax, $50/t	−0.04	−0.147	+0.18 / 3,187	+0.14	−15.2	47.0
U.S. tax, $100/t	−0.12	−0.476	+0.30 / 5,441	+0.18	−25.9	47.4
U.S. tax, $150/t	−0.22	−0.872	+0.39 / 7,068	+0.17	−33.7	47.7
U.S. tax, $210/t	−0.35	−1.354	+0.47 / 8,479	+0.12	−40.4	47.8

Notes: Changes are relative to the no-policy transition over 2023–2100 with observed tariffs. Carbon benefits value discounted world emissions at $210 per ton; world total adds this benefit to noncarbon world CEV. U.S. CEV excludes climate damages. Foreign offset is the discounted foreign emissions increase divided by the U.S. reduction. Tax revenue is rebated lump-sum in the taxing region.

shows that the comparison is not driven by the small annual depreciation of knowledge.

C.8 Carbon Pricing in Equilibrium

This appendix reports the carbon-tax counterfactuals referenced in the main text. A specific fossil-use tax $t^F_{n,t} = \epsilon_n \, SCC$ per unit of real fossil absorption enters the equilibrium: fossil demand responds to the user price $P^F_n + t^F_n$ at every use margin, the consumer price index and the producers' energy-input prices are built from the same user price, the fossil producer continues to receive P^F_n, and the revenue is rebated lump-sum to region n's household. Because observed tariffs are retained and the subsidy schedule is not re-optimized against this objective, the rows are constrained benchmarks rather than a first best; the learning row applies an origin- and technology-specific equipment-subsidy schedule derived from Lemma 1 and iterated to a fixed point on the taxed path.

The unilateral rows show why the $210 rate is a global-damage benchmark rather than an optimal U.S. tax: among the four evaluated rates, carbon-inclusive world welfare peaks near $100 (0.18 percent) and falls to 0.12 percent at $210 because the noncarbon cost grows faster than the abatement benefit. At $210, U.S. discounted emissions fall by 77.4 $GtCO_2$ while rest-of-world emissions rise by 37.0 $GtCO_2$—a discounted foreign-emissions offset of 47.8 percent over 2023–2100—as the world fossil producer price falls by about 6.2 percent. The offset remains between 47 and 48 percent throughout the evaluated range; it reflects large, non-marginal shocks and is sensitive to the calibrated fossil-supply elasticity. The $210 tax raises about $9.0 trillion in present-value revenue for the U.S. household. Like deployment support, the unilateral tax also raises clean-equipment demand, but Table C.10 does not separately decompose the resulting production response. The tax abates far more than the subsidy experiments, although the instruments are not directly comparable: the tax applies through 2100 whereas the subsidy windows close in 2032.

Table C.11: Observed-Tariff Robustness to the Explicit Transition Horizon

Policy	$T = 78$	$T = 148$	Difference
	(10^{-4} percentage points)		
U.S. deployment subsidy	7.43	7.28	−0.15
U.S. equipment subsidy	−0.19	−0.37	−0.18
China deployment subsidy	12.49	12.32	−0.17
China equipment subsidy	21.83	21.58	−0.26

Notes: Entries are in 10^{-4} percentage points and use the observed-tariff specification. They report world consumption CEVs over 2023–2100, excluding carbon and the stationary tail. Both horizons use the same policy rates. The exercise tests the common transition window, not convergence of every terminal state.

C.9 Learning Stock, Horizon, and Accounting Revaluations

In the observed-tariff specification, using $\mu_Q = 1$ instead of 0.99 changes the material policy effects by less than 5 percent and leaves their rankings unchanged; policy effects that are already close to zero move more in relative terms while remaining close to zero. Table C.11 extends the explicit transition from $T = 78$ to $T = 148$ and evaluates both paths over 2023–2100 without the stationary welfare tail. The three material effects change by 1.2–2.1 percent and retain their ranking; U.S. equipment support remains close to zero.

Evaluating each solution through its own terminal date and stationary tail also preserves the ranking. The already small U.S. equipment effect declines substantially with the longer horizon, so the paper does not assign it a precise sign.

In the subsidy exercises, the carbon price does not enter the equilibrium allocation: the flat \$210 per ton SC–CO_2 benchmark revalues the physical emissions deviations after the equilibrium has been solved, an accounting exercise rather than an equilibrium robustness check. By contrast, Appendix C.8 introduces a fossil-use tax and re-solves the equilibrium.

Appendix D Exact Hat Algebra in the Trade Block and Computational Implementation

D.1 Within-period trade block

The Fréchet productivity distributions, the equipment unit-cost function, the delivered cost $\kappa^j_{ni}(1 + \tau^j_{ni,t})(1 - s^j_{i,t})c^j_{i,t}$, and the level Eaton–Kortum price index and bilateral shares are those of Section 3.2, equation (2). This subsection records the remaining unit costs and rewrites the gravity block

in exact changes. The non-energy unit cost and fossil supply curve are

$$c^O_{i,t} = B^O_i \left[w^{1-\gamma_K}_{i,t} (r^D_{i,t})^{\gamma_K} \right]^{\gamma^V_i} (P^E_{i,t})^{\zeta^{E,O}_i} (P^O_{i,t})^{1-\gamma^V_i-\zeta^{E,O}_i}, \tag{D.1}$$

$$c^F_{i,t} = p^F_{i,t}, \qquad \widehat{F}^s_{i,t} = \left(\frac{\widehat{p}^F_{i,t}}{\widehat{w}_{i,t}} \right)^{\varepsilon_F}. \tag{D.2}$$

Here γ^V_i is the non-energy value-added share, γ_K is capital's share within value added, $\zeta^{E,O}_i$ is the non-energy sector's energy-input share, and $1-\gamma^V_i-\zeta^{E,O}_i$ is its own-composite intermediate share. The variable r^D is the rental on non-energy capital, and the fossil origin price clears the reduced-form supply curve. The share η_F of fossil revenue is paid to labor and $1-\eta_F$ is resource rent.

The reference and policy transitions are solved relative to the 2022 historical boundary. Following Dekle et al. (2007), I use exact hat algebra only within the within-period gravity block. Let bars denote the 2022 reference equilibrium and define

$$\widehat{A}^j_{i,t} = \frac{A^j_{i,t}}{\bar{A}^j_i}, \qquad \widehat{d}^j_{ni,t} = \frac{1+\tau^j_{ni,t}}{1+\bar{\tau}^j_{ni}} \frac{1-s^j_{i,t}}{1-\bar{s}^j_i} \frac{c^j_{i,t}}{\bar{c}^j_i}.$$

The level gravity equations (2) imply

$$\widehat{P}^j_{n,t} = \left[\sum_i \bar{\pi}^j_{ni} \widehat{A}^j_{i,t} (\widehat{d}^j_{ni,t})^{-\theta^j} \right]^{-1/\theta^j},$$

$$\pi^j_{ni,t} = \frac{\bar{\pi}^j_{ni} \widehat{A}^j_{i,t} (\widehat{d}^j_{ni,t})^{-\theta^j}}{\sum_m \bar{\pi}^j_{nm} \widehat{A}^j_{m,t} (\widehat{d}^j_{nm,t})^{-\theta^j}}. \tag{D.3}$$

Equation (D.3) uses the observed 2022 bilateral expenditure shares directly, so the baseline Fréchet locations and iceberg costs need not be recovered separately. Baseline tariffs are embedded in both the observed shares and the denominator of the tariff change; counterfactual tariff rates enter through their gross-rate ratio. For non-equipment sectors, productivity is fixed unless an exogenous path is specified. For equipment, equation (12) determines the change in the Fréchet location parameter, while $a^r_{i,t} = (A^{M^r}_{i,t})^{1/\theta^r}$ is the corresponding unit-productivity shifter used in equation (5). In the implementation, the stored equipment cost is $\tilde{c}^{M,r} = c^{M^r}/a^r$, tariff-exclusive receipts are $R^{\text{net}} = \sum_n V_{ni}$, gross factor receipts are $R^{\text{gross}} = R^{\text{net}}/(1-s^M)$, and real output is $y^{M,r} = R^{\text{gross}}/\tilde{c}^{M,r}$.

Capital and experience are not solved in hats. They evolve in 2022-normalized levels according to equations (10) and (11). Observed 2022 capital and equipment production form the historical boundary, and 2023 is the first choice year governed by the Euler equations; no Euler condition is imposed on observed 2022 investment.

D.2 Resource and Policy Accounting

Border-value bilateral sales and tariff revenue are

$$V^j_{ni,t} = \frac{\pi^j_{ni,t} X^j_{n,t}}{1+\tau^j_{ni,t}}, \qquad TR_{n,t} = \sum_{j\in\mathcal{J}^T}\sum_i \frac{\tau^j_{ni,t}}{1+\tau^j_{ni,t}}\pi^j_{ni,t}X^j_{n,t}.$$

Together with origin revenue in equation (3), cost minimization allocates factor income as

$$r^D_{n,t}K^D_{n,t} = \gamma_K\gamma^V_n Y^O_{n,t}, \qquad \Pi^F_{n,t} = (1-\eta_F)Y^F_{n,t},$$
$$w_{n,t}L_n = \sum_r \zeta^{L,r}_n Y^{M^r}_{n,t} + (1-\gamma_K)\gamma^V_n Y^O_{n,t} + \eta_F Y^F_{n,t}.$$

Under the implemented constant-returns technology, renewable operating surplus equals the rental on installed capacity:

$$\Pi^r_{n,t} = P^{G,r,p}_{n,t}G^r_{n,t} - m^r_n P^O_{n,t}K^r_{n,t} = r^r_{n,t}K^r_{n,t}.$$

The national fund's dividend is therefore

$$\begin{aligned}\text{Div}_{n,t} = {} & r^D_{n,t}K^D_{n,t} + \sum_r \Pi^r_{n,t} + \Pi^F_{n,t} + TR_{n,t} \\ & - P^O_{n,t}I^D_{n,t} - \sum_r (1-s^{I,r}_{n,t})C^{\text{proj},r}_{n,t}I^r_{n,t}.\end{aligned}$$

The remaining traded and local absorption terms are

$$\begin{aligned}X^F_{n,t} &= P^F_{n,t}F^d_{n,t}, \\ X^O_{n,t} &= P^O_{n,t}(C^O_{n,t}+I^D_{n,t}) + \sum_r\left[(1-\omega^r)C^{\text{proj},r}_{n,t}I^r_{n,t} + m^r_n P^O_{n,t}K^r_{n,t}\right] \\ &\quad + \sum_r (1-\zeta^{L,r}_n - \zeta^{E,r}_n)Y^{M^r}_{n,t} + (1-\gamma^V_n - \zeta^{E,O}_n)Y^O_{n,t}, \\ X^E_{n,t} &= P^E_{n,t}\mathcal{E}_{n,t} = P^E_{n,t}C^E_{n,t} + \zeta^{E,O}_n Y^O_{n,t} + \sum_r \zeta^{E,r}_n Y^{M^r}_{n,t}.\end{aligned} \tag{D.4}$$

Together with equation (4), these terms exhaust final, intermediate, operating, and investment demand. Energy-market clearing requires

$$P^F_{n,t}F^d_{n,t} = \gamma^F_{n,t}X^E_{n,t}, \qquad \sum_r P^{G,r,u}_{n,t}G^r_{n,t} = (1-\gamma^F_{n,t})X^E_{n,t}, \tag{D.5}$$

with renewable spending split by the inner CES. Non-energy capital and the fossil resource clear domestically.

The investment credit changes who pays for renewable investment, not its gross resource cost: the fund payment plus $\Lambda^{I,r}$ equals $C^{\text{proj},r}I^r$. For equipment support, purchaser spending net of desti-

nation tariffs plus $\Lambda^{M,r}$ equals the origin factor payment in equation (3). Dividing that payment by the productivity-adjusted unit cost gives equation (5); the subsidy top-up is not counted twice. In the PTC branch,

$$P^{G,r,p}_{n,t} = (1+\tau^{D,r}_{n,t})P^{G,r,u}_{n,t}, \qquad \Lambda^{PTC,r}_{n,t} = \tau^{D,r}_{n,t} P^{G,r,u}_{n,t} G^{r}_{n,t}. \tag{D.6}$$

Here $\tau^{D,r}_{n,t}$ is the ad valorem generation-credit rate. Energy demand uses the user price, renewable surplus and rentals use the producer price, and the investment Euler uses the unsubsidized acquisition price because $s^{I,r} = 0$. Contemporaneous lump-sum taxes finance policy spending, and these accounts exhaust resource use.

Fixed transfers $\mathcal{T}_{n,t} = \overline{nx}_n \sum_m w_{m,t} L_m$ reproduce observed external imbalances without creating an asset position. Border-value external balance requires

$$\sum_{j\in\mathcal{J}^T}\sum_i \frac{\pi^j_{ni,t} X^j_{n,t}}{1+\tau^j_{ni,t}} - \sum_{j\in\mathcal{J}^T}\sum_m \frac{\pi^j_{mn,t} X^j_{m,t}}{1+\tau^j_{mn,t}} = \mathcal{T}_{n,t}.$$

Transfers consume no resources and carry no Euler equation. Fossil extraction is a within-period fixed-resource margin, and the model has no bond or net-foreign-asset state.

D.3 Post-2100 Values and Numerical Solution

The household discount factor and the installed-renewable-capital value used in the private Euler equation are

$$M_{n,t+1} = \beta \frac{P^C_{n,t} C_{n,t}}{P^C_{n,t+1} C_{n,t+1}}, \qquad q^r_{n,t} = \frac{\widetilde{P}^{\text{inv},r}_{n,t}}{\Phi^r_{I,n,t}},$$

$$\Phi^r_{I,n,t} = \alpha \chi^r_n \left(\frac{I^r_{n,t}}{K^r_{n,t}}\right)^{\alpha-1}, \qquad \Phi^r_{K,n,t} = (1-\alpha)\chi^r_n \left(\frac{I^r_{n,t}}{K^r_{n,t}}\right)^{\alpha}, \qquad r \in \mathcal{R}.$$

The economic model has an infinite horizon. The computation solves an explicit transition through a distant terminal year, after which policy rates are zero. The last renewable Euler row uses a stationary private-return approximation,

$$q^r_{n,T}\left[1 - \beta \Gamma^{q,r}_{n,T}\left(1 - \delta_r + \Phi^r_{K,n,T}\right)\right] = \beta r^r_{n,T}.$$

Here $\Gamma^{q,r}_{n,T} \equiv q^r_{n,T+1}/q^r_{n,T}$ is the assumed post-2100 growth factor for installed-capital value. The benchmark sets $\Gamma^{q,r}_{n,T} = 1$. When experience is fully cumulative and learning remains active, the alternative specification lets $\Gamma^{q,r}_{n,T}$ reflect the sourcing-weighted decline in future equipment prices implied by experience growth at the end of the explicit horizon. The non-energy Euler row uses the

same stationary closure,

$$q^D_{n,T}\left[1-\beta\,\Gamma^{q,D}_{n,T}\left(1-\delta_D+\Phi^D_{K,n,T}\right)\right]=\beta r^D_{n,T},$$

with $\Gamma^{q,D}_{n,T}\equiv q^D_{n,T+1}/q^D_{n,T}=1$ in the benchmark.

Household welfare uses the stationary last-consumption-gap tail in equation (A.14); the carbon account ends in 2100 and carries no stationary tail. The envelope proof of Appendix A.1 uses the infinite horizon, whereas the finite implementation has three distinct boundary approximations: the terminal private-capital return, the post-2100 evolution of learning, and the stationary last-consumption-gap tail. They need not be derivatives of a single post-2100 value function. Denote their combined directional error by

$$\mathcal{R}^{\text{bdry}}_{n,T}=\mathcal{R}^K_{n,T}+\mathcal{R}^Q_{n,T}+\mathcal{R}^C_{n,T}.$$

With endogenous non-energy capital, $\mathcal{R}^K_{n,T}$ includes its terminal approximation in addition to the solar and wind terms. It also contains the truncated telescoping term $-\mu_{n,T}\bar{q}^D_{n,T}\dot{K}^D_{n,T+1}$, which the infinite-horizon transversality condition would remove but the finite implementation retains. An exact derivative of the post-2100 value function would make all three terms zero. Because stock size alone does not determine the weighted boundary error, the calculations do not assign this residual to a particular capital block; the exercises in Appendix C.9 assess the combined quantitative importance of these boundary approximations.

I verify the goods, factor, fiscal, external-balance, capital, learning, and investment-Euler conditions for every reported solution. A tariff-free, balanced-trade benchmark without learning or carbon also reproduces the envelope cancellations derived above. The observed-tariff baseline contains policy wedges and is not treated as first best.

D.4 Solution Algorithm

Given a policy path, I solve for country–year sequences of the logarithms of wages, all three investment flows, and fossil prices from 2023 through 2100. One wage is omitted in each year under the world-labor-income numeraire, so with fifteen regions and a 78-year horizon the stacked unknown vector has $(5\times 15-1)\times 78=5{,}772$ entries. For any candidate vector, the model is evaluated forward from the inherited 2022 capital and experience stocks. In each year, a within-period fixed point solves the static block at the candidate wages and fossil prices: the unit-cost system of this appendix, including equation (D.1) and the feedback of the energy composite price into non-energy and equipment costs; the exact-change price indices and bilateral shares in equation (D.3); equipment demand and origin revenue in equations (4) and (3); the energy mix and dual prices in equation (A.12) with the clearing condition (D.5); the absorption identities in equation (D.4); and the household budget in equation (9). The candidate investment sequences update installed capacity through equation (10),

and equilibrium real equipment output from equation (5) updates the experience stocks and next year's equipment productivity through equations (11) and (12). Because the trade block is written in exact changes around observed 2022 shares, this evaluation never requires the levels of bilateral trade costs or fundamental productivity.

The outer residual system is square. For every year it stacks three blocks: labor-market clearing in all but the numeraire country, using the factor-income allocation of this appendix; the investment Euler equations in every country, comprising the solar and wind conditions (13) and the non-energy condition (14), whose final rows use the terminal conditions described above; and fossil-market clearing at each origin, equating the fossil demand implied by equation (D.5) to the supply curve in equation (D.2). The Euler residuals make current investment consistent with the entire future sequence of rents, installation costs, and capital values. The omitted labor-market condition is retained as an independent Walras check.

Following the sequence-space approach of Auclert et al. (2021), I solve this stacked system with a damped Newton method, reusing the Jacobian of the residual vector across policy counterfactuals and introducing large policy changes gradually. Reported solutions satisfy a maximum normalized equilibrium residual of 10^{-10} and are separately checked against the household budget (9), the goods-market and absorption identities of this appendix, the fiscal accounts, the capital and experience laws (10) and (11), and the omitted Walras condition.