Green Industrial Policy in a Globalized Economy [*†]

Sungwan Hong
University of Pittsburgh

July 2026

Abstract

Renewable-energy equipment is traded internationally and its production is concentrated in a few countries, while the electricity it generates is largely local. I develop a dynamic multi-country trade model with learning-by-doing to study how the target and location of green subsidies shape their international effects. At the same fiscal cost, U.S. deployment subsidies generate a larger global gain and a smaller domestic loss than equipment-production subsidies. In China, deployment subsidies deliver the larger domestic gain, but equipment subsidies generate the larger global gain by expanding production and learning at a major exporter and lowering future equipment costs abroad. National incentives favor the opposite assignment: U.S. equipment support and Chinese deployment support. Ignoring equipment trade makes U.S. equipment subsidies appear more valuable globally than deployment subsidies and removes most of the advantage of subsidizing production at the export hub. Equipment trade also shapes the international incidence of emissions: deployment subsidies can raise foreign emissions by bidding up equipment prices, whereas equipment subsidies at an established exporter can lower them by reducing delivered costs abroad.

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 the equipment it uses. The U.S. Inflation Reduction Act (IRA) combines both approaches, and related programs operate in Europe, China, India, and Japan.[1] The effects of these policies depend not only on the instrument, but also 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 and learning toward the supporting country and can lower future installation costs in the foreign markets it serves.[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 therefore change not only current sourcing and income, but also the future geography of comparative advantage and the cost of renewable deployment. Fossil-fuel trade adds a second international channel through energy prices and carbon leakage.

The main result is that policy assignment depends on a country's position in the equipment market. Holding fiscal cost fixed, deployment support in the United States produces the larger global gain and the smaller U.S. loss. U.S. equipment support instead draws production away from established foreign suppliers, so its learning benefit is largely offset by production distortion. In China, deployment support produces the larger domestic gain, but equipment support produces the larger global gain: it expands production and learning at a major exporter and reduces future equipment costs for foreign buyers.

Equipment trade is central to this ranking. Removing it makes U.S. equipment support more valuable globally than deployment support and eliminates most of the extra global return to locating production support in China rather than the United States. The emissions response also differs across instruments and locations. Deployment support lowers emissions at home but can raise them abroad as equipment demand bids up equipment prices, crowds out foreign installations, and low-

[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.

ers the world fossil price. Equipment support can instead lower the supported origin's delivered price abroad, with production learning extending that cost reduction into the future. Alongside the familiar fossil-price channel, trade in clean equipment therefore creates a second channel of international emissions transmission: deployment support can crowd out renewable investment abroad by tightening the global equipment market, whereas equipment support at an established export hub can accelerate it by lowering delivered costs. U.S. equipment support does not produce the same foreign-emissions decline because it expands a smaller production base partly by displacing established suppliers.

National and global incentives do not select the same policies. When each country chooses a single instrument to maximize its own welfare, the United States prefers equipment support and China prefers deployment support. For global welfare, the preferred assignment is the opposite: U.S. deployment support combined with Chinese equipment support. This gap arises because governments value domestic production, learning, and trade-price gains but do not fully internalize the lower equipment costs and emissions reductions that accrue abroad.

A stylized IRA-scale package combining U.S. deployment and equipment subsidies reduces discounted world emissions by 1.90 $GtCO_2$, but lowers U.S. and global 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 analytical results organize these findings around three channels: trade-price incidence, production learning, and carbon. A carbon tax together with origin-specific learning subsidies implements the planner allocation when both instruments are available. Deployment support is instead a second-best way to reward the carbon and learning benefits of renewable investment when direct instruments are restricted. A separate routing result formalizes the importing-market–export-hub distinction: equipment support in an import-dependent market creates domestic experience by diverting orders from foreign suppliers, whereas support at an export hub expands output on a broader sales base. Bilateral sourcing and the learning value of the production displaced elsewhere determine the net effect.

The quantitative model 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, with a smaller but economically meaningful local component; variation in exporters' exposure to foreign feed-in-tariff expansions provides corroborating evidence. Wind learning is disciplined by price and manufacturer-location data.

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 in which many origins serve many destinations. Producer support may therefore expand world experience or reallocate production away from other suppliers. The world price–experience relation disciplines frontier improvement, while bilateral exporter performance disciplines the geographic component of learning and how much of the resulting cost decline the subsidizing country captures.

Clean growth and climate policy. A large literature studies clean growth, renewable deployment, and climate policy (Acemoglu et al., 2012, 2016; Fried, 2018; Bistline et al., 2023; Casey et al., 2023; Donald, 2024; Gentile, 2024). Hémous (2016) studies unilateral policy and the evolution of clean comparative advantage. Kotchen and Maggi (2025) characterize carbon taxes and green-energy production subsidies in a two-country economy, emphasizing fossil-price leakage, lobbying, and international agreements. This paper instead distinguishes subsidies to renewable deployment from subsidies to the traded equipment used for deployment; bilateral sourcing and production learning determine their incidence over time. In a closed-economy supply-chain model, Aghion et al. (2025) show that targeted technology subsidies can remain useful alongside carbon pricing when technology adoption is complementary across sectors. The mechanism here is instead production learning in a traded clean capital good. A close quantitative antecedent is Arkolakis and Walsh (2023), who combine global learning with a detailed spatial electricity system. Relative to that framework, the model resolves equipment origins and bilateral sourcing, distinguishes world from origin-specific experience, and compares destination-based deployment subsidies with origin-based equipment subsidies. Hahn et al. (2026) provide a complementary welfare analysis of U.S. climate policies that includes market-wide learning benefits.

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 (Young, 1991; Bai et al., 2024). Here, a temporary policy changes bilateral equipment sourcing, origin production accumulates into

experience, and experience 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

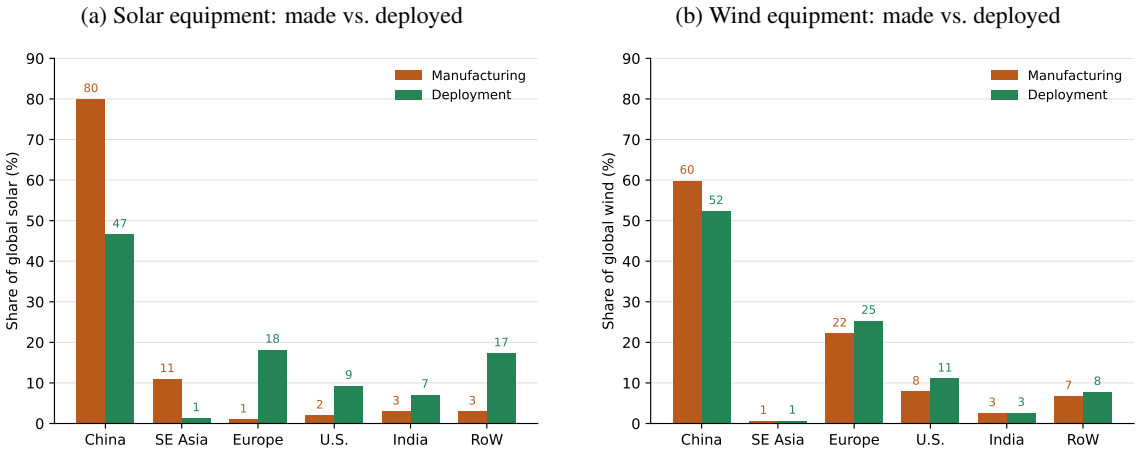
2.1 Production, Deployment, and Trade in Renewable-energy Equipment

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 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%. This descriptive comparison should not be read as a measure of bilateral wind-equipment trade. In particular, country shares of manufacturing and capacity additions do not reveal cross-border sourcing of blades, nacelles, towers, and other components. The model's sourcing patterns are therefore disciplined by bilateral equipment-expenditure matrices rather than inferred from the gaps in Figure 1.

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

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.

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

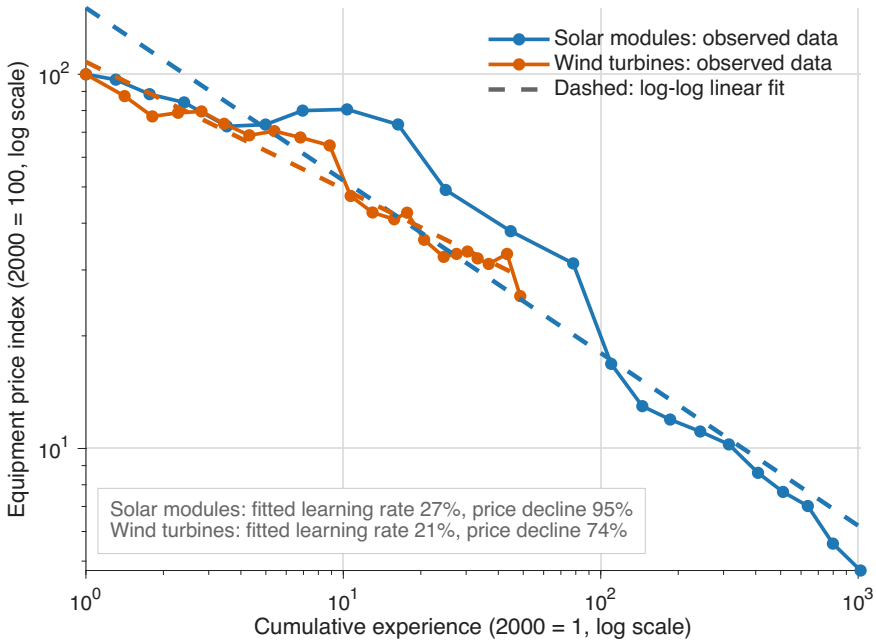
2.2 Equipment Cost Declines and Learning-by-Doing

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.

[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 it does not capture every component or production network and is not a comprehensive measure of wind-supply-chain openness.

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.

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 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 not only current

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

prices and production but also where production experience accumulates and who benefits from the resulting cost decline. The next section formalizes that link between bilateral sourcing, equipment production, 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. Bilateral sourcing determines where equipment is produced and where production experience accumulates; experience in turn lowers future equipment costs and changes comparative advantage. Electricity generation and fossil displacement occur locally, while renewable equipment, fossil fuel, and non-energy goods are traded. Conditional on inherited capital and experience, the static equilibrium determines prices, trade, and production; investment and equipment production then update the dynamic state. The same environment underlies the incidence results in Section 4 and the quantitative counterfactuals.

The within-period production and sourcing block is a multi-sector Eaton–Kortum economy with input–output linkages (Eaton and Kortum, 2002; Caliendo and Parro, 2015). Renewable equipment is a traded investment good, 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 adds a further state: its global component nests the learning block in Arkolakis and Walsh (2023), while its origin-specific component 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 an object in the equipment sector associated with renewable technology r. Solar and wind generation are local. The non-energy composite is used for consumption, production, investment, balance-of-system inputs, and operating 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 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, and absorption clear 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, $\sum_n w_{n,t} L_n = 1$.

Preferences. The representative household supplies labor inelastically. Consumption combines the non-energy good 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, and the annual discount factor is $\beta = 0.96$. The one-period discount factor for nominal payoffs is $M_{n,t+1} = \beta P^C_{n,t} C_{n,t} / (P^C_{n,t+1} C_{n,t+1})$.

3.2 Production and Bilateral Trade

The equipment sectors link deployment in one country to production in another. 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

$$c^{M^r}_{i,t} = B^r_i w^{\zeta^{L,r}_i}_{i,t} (P^E_{i,t})^{\zeta^{E,r}_i} (P^O_{i,t})^{1-\zeta^{L,r}_i-\zeta^{E,r}_i}, \qquad r \in \mathcal{R}.$$

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

$$F^s_{i,t} = \varphi^F_i \left(\frac{p^F_{i,t}}{w_{i,t}}\right)^{\varepsilon_F},$$

where φ^F_i 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^j_{ni} be the iceberg cost, $\tau^j_{ni,t}$ an import tariff, and $s^j_{i,t}$ an origin subsidy, which is zero outside renewable equipment. Define delivered unit cost by

$$d^j_{ni,t} = \kappa^j_{ni}(1+\tau^j_{ni,t})(1-s^j_{i,t})c^j_{i,t}.$$

The Eaton–Kortum price index and bilateral expenditure share are

$$\begin{aligned} P^j_{n,t} &= \gamma^j \left[\sum_i A^j_{i,t}(d^j_{ni,t})^{-\theta^j}\right]^{-1/\theta^j}, \\ \pi^j_{ni,t} &= \frac{A^j_{i,t}(d^j_{ni,t})^{-\theta^j}}{\sum_m A^j_{m,t}(d^j_{nm,t})^{-\theta^j}}, \qquad j \in \mathcal{J}^T, \end{aligned} \tag{2}$$

where γ^j is the usual Fréchet price constant. Appendix D rewrites these equations in exact changes around the observed 2022 expenditure shares used in computation. Let $X^j_{n,t}$ denote purchaser expenditure in destination n. Gross origin revenue available to pay factors, including the subsidy top-up, is

$$Y^j_{i,t} = \frac{1}{1-s^j_{i,t}} \sum_n \frac{\pi^j_{ni,t} X^j_{n,t}}{1+\tau^j_{ni,t}}, \tag{3}$$

Destination spending is allocated across origins by $\pi^j_{ni,t}$, while equation (3) clears origin sales. Domestic factor payments and tariff revenue follow from cost minimization and accrue to the household.

3.3 Renewable Investment and Local Energy

Renewable investment provides the demand bridge between local deployment and the world equipment market. A developer in country n combines traded equipment with local balance-of-system inputs to form a renewable investment bundle. Denote the bundle's unit cost by $C^{\text{proj},r}_{n,t}$. Its expenditure decomposition is

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

Here B_{nr} is a destination–technology installation-cost shifter, ω^r is equipment's share of installation cost, and $X^{B,r}_{n,t}$ is local balance-of-system spending. The equipment term is allocated across origins by $\pi^{M^r}_{ni,t}$. Thus $\omega^r \pi^{M^r}_{ni,t}$ measures how renewable-investment spending in destination n reaches supplier i.

For the experience law, let $a^r_{i,t} \equiv (A^{M^r}_{i,t})^{1/\theta^r}$ and $\tilde{c}^{M,r}_{i,t} \equiv c^{M^r}_{i,t}/a^r_{i,t}$. Origin real equipment output

is the producer-revenue index deflated by this productivity-adjusted unit cost:

$$y_{i,t}^{M,r} = \frac{Y_{i,t}^{M^r}}{\widetilde{c}_{i,t}^{M,r}}. \tag{5}$$

This origin-based index measures factory-gate output in common resource-cost units.

Local energy. Installed capacity produces nontraded electricity,

$$\begin{aligned} G_{n,t}^r &= z_n^r K_{n,t}^r, \qquad r \in \mathcal{R}, \\ r_{n,t}^r &= P_{n,t}^{G,r,p} z_n^r - m_n^r P_{n,t}^O. \end{aligned} \tag{6}$$

Here z_n^r is generation productivity, $P_{n,t}^{G,r,u}$ is the energy-nest user price, $P_{n,t}^{G,r,p}$ is the price received by the generator, and m_n^r is operating input per unit of capacity. The two prices coincide absent a renewable-generation production subsidy. Generation has constant returns to installed capacity.

Solar and wind first form a renewable composite,

$$R_{n,t} = \left[(\eta_n^S)^{1/\rho} (G_{n,t}^S)^{(\rho-1)/\rho} + (\eta_n^W)^{1/\rho} (G_{n,t}^W)^{(\rho-1)/\rho} \right]^{\rho/(\rho-1)}, \qquad \eta_n^S + \eta_n^W = 1, \tag{7}$$

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

Here $F_{n,t}^d$ is fossil absorption in destination n, and $P_{n,t}^E$ is the dual price of energy services. The η_n^r'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 $\mathcal{E}_{n,t} = C_{n,t}^E + E_{n,t}^O + \sum_r E_{n,t}^{M^r}$, where $E_{n,t}^O$ and $E_{n,t}^{M^r}$ are the energy inputs of the non-energy and equipment sectors and all users face the same composite price $P_{n,t}^E$. Household demand assigns expenditure share μ_n^E to energy, and producer energy demand follows from the unit-cost functions above. Appendix A.3 gives the dual prices and expenditure shares. The annual block does not resolve dispatch, storage, or transmission; Appendix B.6 discusses that 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

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

Here $\mathcal{T}_{n,t}$ is an international transfer and $\Lambda_{n,t}$ the lump-sum policy tax. 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 that external position; setting $\overline{nx}_n = 0$ gives balanced trade. Given the state and investment, market clearing determines wages, prices, expenditure, sourcing, and origin production. Appendix D gives the complete dividend, external-balance, absorption, and market-clearing equations.

3.5 Investment and Learning Dynamics

The within-period equilibrium maps deployment across destinations into equipment production across origins. Two laws of motion carry those choices forward: 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

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

Here δ_k is depreciation, χ^k_n the installation scale, and the inherited stock a fixed 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, and 2023 is the first choice year for each; observed 2022 investment enters only through the inherited stocks.

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

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

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

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

The denominator $Q^r_{i,2022}$ is the beginning-of-2022 stock accumulated through 2021, with $Q^r_{W,2022} = \sum_i Q^r_{i,2022}$. The first equilibrium state is beginning-of-2023 experience, $Q^r_{i,2023} = \mu_{Q,r} Q^r_{i,2022} + y^{M,r}_{i,2022}$, so 2022 production enters once before the first policy choice. Here $\nu^r_{\text{tot}} \geq 0$ is the Fréchet-location elasticity of total experience, $g_r = \nu^r_{\text{tot}} / \theta^r$ is the corresponding price–experience elasticity with $\theta^r \equiv \theta^{M^r}$, 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 $\lambda_L^r = 0$ nests the global-learning specification of Arkolakis and Walsh (2023); positive λ_L^r adds an origin-specific component that can change comparative advantage. Individual producers take both stocks as given, making equation (12) a country–sector Marshallian externality.

Let $P_{n,t}^{\text{inv},r}$ denote the pre-subsidy acquisition price of one unit of technology-r renewable investment. Because investment is purchased as the renewable investment bundle, $P_{n,t}^{\text{inv},r} = C_{n,t}^{\text{proj},r}$, while the price paid by the fund is $\widetilde{P}_{n,t}^{\text{inv},r} = (1 - s_{n,t}^{I,r}) P_{n,t}^{\text{inv},r}$. The installed-capital value is $q_{n,t}^r = \widetilde{P}_{n,t}^{\text{inv},r} / \Phi_{I,n,t}^r$, where Φ_I and Φ_K are the two derivatives of the installation technology. The fund's renewable-capital Euler equation is

$$q_{n,t}^r = M_{n,t+1} \left[r_{n,t+1}^r + q_{n,t+1}^r \left(1 - \delta_r + \Phi_{K,n,t+1}^r \right) \right], \qquad r \in \mathcal{R}. \tag{13}$$

Non-energy investment is chosen on the same margin, with installed value $q_{n,t}^D = P_{n,t}^O / \Phi_{I,n,t}^D$ and Euler equation

$$q_{n,t}^D = M_{n,t+1} \left[r_{n,t+1}^D + q_{n,t+1}^D \left(1 - \delta_D + \Phi_{K,n,t+1}^D \right) \right]. \tag{14}$$

The renewable rental r^r is given by equation (6), while r^D is the rental on non-energy capital. All three Euler equations use the same household discount factor and the same installation technology. Learning does not appear in equation (13) because firms and the fund take A^{M^r}, aggregate experience, and the resulting trade shares as given. 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 instruments enter on opposite sides of the investment–trade mapping. An origin equipment subsidy $s_{i,t}^{M,r}$ lowers the delivered cost of all equipment produced in i and therefore changes sourcing in every destination. It can raise production and local experience at i while reducing both at competing origins; equations (2), (5), and (11) retain that full displacement. The destination-side instrument is a deployment subsidy, $s_{n,t}^{I,r}$, implemented as an investment tax credit (ITC). It lowers renewable-investment cost 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, $(1 - s_{n,t}^{I,r}) C_{n,t}^{\text{proj},r} I_{n,t}^r$, while the government of the same country finances the subsidy. The

[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 pays inframarginal as well as newly induced generation. Appendix B.1 maps a PTC into an equivalent marginal investment incentive, while Appendix D records its distinct fiscal bill.

fiscal bills are

$$\Lambda_{n,t}^{I,r} = s_{n,t}^{I,r} C_{n,t}^{\text{proj},r} I_{n,t}^{r},$$
$$\Lambda_{i,t}^{M,r} = \frac{s_{i,t}^{M,r}}{1 - s_{i,t}^{M,r}} \sum_{n} \frac{\pi_{ni,t}^{M^r} X_{n,t}^{M^r}}{1 + \tau_{ni,t}^{M^r}}, \qquad \Lambda_{n,t} = \sum_{r} (\Lambda_{n,t}^{I,r} + \Lambda_{n,t}^{M,r}).$$

The first expression is country n's deployment-subsidy rate times its technology-r renewable-investment expenditure. In the second, the sum is the tariff-exclusive amount paid by buyers worldwide for technology-r equipment from origin i. Buyers finance the share $1 - s_{i,t}^{M,r}$ of gross producer revenue, so the origin government supplies the remaining share $s_{i,t}^{M,r}$; its top-up is therefore $s_{i,t}^{M,r}/(1 - s_{i,t}^{M,r})$ times buyer-funded receipts. Finally, $\Lambda_{n,t}$ sums country n's deployment and equipment subsidies across solar and wind. All policy spending is financed by the contemporaneous lump-sum tax in equation (9).

The two instruments also transmit emissions differently across countries. A deployment subsidy raises equipment demand, and the resulting capacity first displaces fossil use in the installing country. In general equilibrium, the demand increase can bid up equipment costs and crowd out renewable investment elsewhere, while the fall in the world fossil price creates conventional carbon leakage; subsequent global learning partly offsets these effects. An equipment subsidy instead lowers the supported origin's effective delivered cost in every destination. Production shifts toward that origin, but lower current equipment prices and the future cost reductions generated by learning can raise renewable investment and reduce emissions in importing countries. The model therefore allows international emissions spillovers through both fossil markets and clean-equipment markets.

Given the inherited 2023 state and policy paths, an equilibrium satisfies producer and developer cost minimization, household and fund optimality, government budgets, market clearing, and the capital and experience laws of motion. Equipment producers take aggregate experience as given. Temporary green subsidies eventually return to zero, while regime-specific trade policy, other trade costs, and transfers remain in place. Appendix D states the terminal closure, and Appendix C reports horizon checks.

3.7 Welfare and Carbon Accounting

Fossil absorption generates territorial CO_2 emissions,

$$Z_t = \sum_{n} \epsilon_n F_{n,t}^{d}. \tag{15}$$

Here ϵ_n is territorial CO_2 per unit of fossil absorption. Carbon damages do not enter private utility or equilibrium choices, but they are incorporated later through an ex post carbon-inclusive welfare calculation based on the policy-induced emissions path. In the reference and green-subsidy experiments,

no carbon price enters firms' first-order conditions or the equilibrium allocation; Appendix C.8 separately introduces a fossil-use tax. For policy p, country n's consumption-welfare change is

$$\Delta U_n^{\text{cons}} = \sum_{t=1}^{\infty} \beta^{t-1} \left(\log C_{n,t}^{p} - \log C_{n,t}^{0} \right). \tag{16}$$

Let ℓ_n^{cons} denote the corresponding country log consumption equivalent and let $\overline{\mathcal{E}}_{n,T}$ be country n's present-value baseline expenditure under the same post-2100 welfare-tail assumption used for every policy. Let SCC_t be 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

$$\begin{aligned} \ell_W^R &= \sum_n \omega_n \ell_n^{\text{cons}}, & \omega_n &= \frac{\overline{\mathcal{E}}_{n,T}}{\sum_m \overline{\mathcal{E}}_{m,T}}, \\ \ell_W^{R,SCC} &= \ell_W^R - \frac{1}{\overline{\mathcal{E}}_{W,T}} \sum_{t=1}^{T} \beta^{t-1} SCC_t \left(Z_t^p - Z_t^0 \right), & \overline{\mathcal{E}}_{W,T} &= \sum_n \overline{\mathcal{E}}_{n,T}. \end{aligned} \tag{17}$$

The fixed weights define a distribution-sensitive incidence index; the Pigouvian benchmark in Section 4 instead uses a cooperative planner objective. Changing the SC–CO_2 path or damage shares revalues a solved allocation without changing equilibrium. Appendix A.4 explains the post-2100 welfare adjustment, level-CEV conversion, and country-damage allocation.

Because the full transition contains many general-equilibrium feedbacks, the next section differentiates the equilibrium around the reference path. The resulting envelope and routing formulas isolate the bilateral sourcing, learning, and fossil-use moments that organize the quantitative exercises.

4 International Incidence and Policy Assignment

Section 3 implies that green industrial policy affects national welfare through trade-price incidence, production learning, and carbon. To isolate these mechanisms, this section uses a tariff-free reference economy. It first derives a local welfare envelope for the full transition, then states a direct first-best benchmark, and finally traces how equipment and deployment subsidies route production through the supplier network. The envelope decomposes national consumption welfare into international-price incidence and production learning, and appends carbon damages when evaluating global welfare. These channels provide the vocabulary used throughout the propositions and the quantitative decomposition of Section 6.2. The results apply to arbitrary positions in the supplier network; the United States and China enter later as empirically distinct cases.

4.1 International Price Incidence, Learning, and Carbon

Consider a one-dimensional policy perturbation indexed by p around a zero-green-subsidy reference transition. Let p scale a complete policy sequence. Bars denote reference values, and dots are total derivatives of the equilibrium transition with respect to p, including all induced price, sourcing, investment, capital, experience, and fossil-use responses. Let $E_{n,t} = P^C_{n,t} C_{n,t}$ be consumption expenditure and define bilateral sales in the tariff-free reference by

$$V^j_{ni,t} = \pi^j_{ni,t} X^j_{n,t}.$$

For origin i, let $\ell^j_{i,t} = \log[(1 - s^j_{i,t}) c^j_{i,t}]$ denote the origin-side price component, including an origin subsidy but excluding the Fréchet productivity shifter.

The two noncarbon money-metric components at date t admit a plain reading: $\mathcal{I}^p_{n,t}$ captures international price and transfer incidence—the net effect of export receipts, import costs, and external transfers—while $\mathcal{L}^p_{n,t}$ values the equipment-productivity path that current production sets in motion. Formally,

$$\begin{aligned} \mathcal{I}^p_{n,t} &= \sum_j \sum_d \bar{V}^j_{dn,t} \dot{\ell}^j_{n,t} - \sum_j \sum_i \bar{V}^j_{ni,t} \dot{\ell}^j_{i,t} + \dot{\mathcal{T}}_{n,t}, \\ \mathcal{L}^p_{n,t} &= \sum_{r \in \mathcal{R}} \sum_i \frac{\bar{V}^{M^r}_{ni,t}}{\theta^r} \frac{d \log A^{M^r}_{i,t}}{dp}. \end{aligned} \tag{18}$$

The first line is the net change in export receipts and import costs, evaluated at baseline bilateral sales, together with the transfer term.[8] The learning term values the equipment-productivity path generated by equations (11)–(12).[9]

Let $\bar{E}_{n,t}$ denote household expenditure in the reference equilibrium, and define

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

The weights ϖ_n are the infinite-horizon counterparts of the baseline-expenditure weights used to aggregate country consumption equivalents. Let ℓ^{cons}_n denote private consumption welfare; the superscript R below denotes reference-expenditure weighting. Consider a tariff- and subsidy-free competitive reference satisfying the private equilibrium conditions and $\eta_F = \varepsilon_F / (1 + \varepsilon_F)$. Policy changes are financed lump sum, and inherited states are held fixed.

Proposition 1 (Welfare decomposition)**.** *For a marginal change in either subsidy, country n's con-*

[8] In the quantitative normalization the fixed external transfer has zero derivative, but it is retained here for the general accounting identity.

[9] Appendix A.1 derives the additional Harberger term with pre-existing tariffs. The IRA and source-restriction applications restore observed equipment tariffs.

sumption equivalent satisfies

$$\dot{\ell}_n^{\text{cons}} = \frac{1}{\mathcal{B}_\beta} \sum_{t=1}^{\infty} \frac{\beta^{t-1}}{\bar{E}_{n,t}} \left[\mathcal{I}_{n,t}^p + \mathcal{L}_{n,t}^p \right], \tag{19}$$

The corresponding carbon-inclusive global welfare derivative is

$$\dot{\ell}_W^{R,SCC} = \frac{1}{\mathcal{B}_\beta} \sum_n \varpi_n \sum_{t=1}^{\infty} \frac{\beta^{t-1}}{\bar{E}_{n,t}} \left[\mathcal{I}_{n,t}^p + \mathcal{L}_{n,t}^p \right] - \frac{1}{\sum_n \bar{\mathcal{E}}_n} \sum_{t=1}^{\infty} \beta^{t-1} SCC_t \dot{Z}_t. \tag{20}$$

where the last term appends the ex-post carbon account. If there exist positive country scales a_n and a common positive time profile b_t such that $\bar{E}_{n,t} = a_n b_t$, international incidence cancels in the global aggregation, leaving learning and carbon as the only first-order terms.

The proportional-expenditure condition includes a static economy and a common balanced-growth path. Outside it, international incidence need not cancel because reference-expenditure aggregation assigns different welfare values to a dollar received in different countries and dates. At a zero-subsidy reference, the subsidy payment and its lump-sum financing cancel to first order; the conventional policy-distortion term appears away from zero. Other competitive margins—including installation frictions, intermediate inputs, and fossil-price responses—enter through private technologies and add no separate first-order term. Fossil trade still affects international incidence and $\dot{Z}_t$. Table C.1 provides the finite-policy counterpart. Its international-price and learning columns integrate $\mathcal{I}_{n,t}^p$ and $\mathcal{L}_{n,t}^p$ along the policy path; policy cost appears away from zero, and carbon is reported separately.

Appendix A.1 gives the derivation, the extension with pre-existing tariffs, the transfer-neutral special case, and the boundary error associated with a finite numerical horizon.

Efficiency benchmark. The same envelope gives a compact direct implementation of the constrained first best. Consider a cooperative planner with Pareto weights and international lump-sum transfers, so that the social value of a unit of the numeraire is common across countries at each date. Let $\mathcal{B}_{i,t}^{L,r}$ denote the discounted marginal value of the future world and origin-specific learning generated by one additional unit of current technology-r equipment output at origin i. Appendix equation (A.7) gives its full expression.

Lemma 1 (Pigouvian implementation)**.** *In the tariff-free environment of Proposition 1, suppose country–date lump-sum transfers support the planner's international consumption allocation and learning and carbon are the only externalities. A direct implementation combines an origin-specific equipment subsidy that prices learning with a fossil-use tax that prices carbon; no deployment subsidy is required. If all dollar objects, including the date-t marginal damage SCC_t, are expressed in the model numeraire, the implementing wedges satisfy*

$$s_{i,t}^{M,r,*} \tilde{c}_{i,t}^{M,r} = \mathcal{B}_{i,t}^{L,r}, \qquad \tau_{n,t}^{C,*} = \epsilon_n SCC_t, \qquad s_{n,t}^{I,r,*} = 0. \tag{21}$$

Here $\tau^{C}_{n,t}$ is a specific tax on fossil absorption at every use margin, household and intermediate. No additional subsidy to renewable generation is required.

This implementation is the production-learning counterpart of the familiar carbon-tax-plus-clean-innovation-subsidy result (Acemoglu et al., 2012; Hémous, 2016): the fossil tax prices emissions, while the origin subsidy prices the future productivity created by current equipment output. Because equipment output at every active origin contributes to world experience, this cooperative implementation generally subsidizes each producing origin; it is not a collection of unilateral national best responses. The first equality prices the future world and origin-specific learning generated by current production; the second prices territorial emissions at marginal damage. To implement the learning wedge as the model's ad valorem subsidy requires $0 \leq \mathcal{B}^{L,r}_{i,t} / \widetilde{c}^{M,r}_{i,t} < 1$ at every active origin and date.[10] Because the carbon tax applies to fossil use at every margin, the equipment unit cost in the subsidy base is carbon-tax-inclusive; the specific wedge $s^{M,r,*}_{i,t} \widetilde{c}^{M,r}_{i,t} = \mathcal{B}^{L,r}_{i,t}$ is unchanged, while the implied ad valorem rate is correspondingly lower in fossil-intensive origins.

Once those two margins are priced directly, the private renewable-capital Euler equation coincides with the planner's condition. If either direct instrument is unavailable, a deployment subsidy can partly price a renewable installation's carbon benefit and its sourcing-weighted average learning benefit. Because it operates at the destination, however, a single deployment subsidy generally cannot reproduce the vector of origin-specific learning wedges.

4.2 Direct Policy Routing through the Supplier Network

The gravity block traces how each instrument changes renewable-investment costs and equipment production. Holding current origin costs, productivities, trade costs, and destination equipment spending fixed, an origin subsidy changes renewable-investment cost in destination n according to

$$\frac{\partial \log C^{\text{proj},r}_{n,t}}{\partial s^{M,r}_{i,t}} = -\frac{\omega^{r} \pi^{M^r}_{ni,t}}{1 - s^{M,r}_{i,t}}. \tag{22}$$

The exposure is the product of equipment's share of installation cost and the destination's sourcing share from the supported origin.

To describe production, let $\widetilde{c}^{M,r}_{i,t} = c^{M^r}_{i,t} / (A^{M^r}_{i,t})^{1/\theta^r}$. Let $\ell^{R,L}_{W}$ denote the local learning component of global welfare around the reference transition, holding international-price incidence separate. Let $\Omega^{r}_{i,t}$ denote the contribution of one additional unit of current equipment output in origin i to this

[10] Without this bound, $\mathcal{B}^{L,r}_{i,t}$ remains the required specific subsidy per unit of output, but it cannot be represented by the model's positive-price factor $1 - s^{M,r}_{i,t}$.

learning component through future world and origin-specific experience. Formally,

$$\Omega^r_{i,t} \equiv \sum_{u>t} \sum_h \frac{\partial \ell^{R,L}_W}{\partial \log A^{M^r}_{h,u}} \frac{\partial \log A^{M^r}_{h,u}}{\partial y^{M,r}_{i,t}} \bigg|_0 ,$$

where the derivatives are evaluated along the reference transition. This value incorporates experience retention, future sourcing, and both learning components in equation (12); Appendix equation (A.6) gives the expanded expression. It need not equal the planner value $\mathcal{B}^{L,r}_{i,t}$ above.[11] Define origin h's learning value per dollar of receipts by

$$\kappa^r_{h,t} \equiv \frac{\Omega^r_{h,t}}{\tilde{c}^{M,r}_{h,t}}.$$

For a supported origin i, let $\bar{S}^r_{i,t} = \sum_n \bar{V}^{M^r}_{ni,t}$ denote its total equipment receipts, $\chi^r_{ni,t} = \bar{V}^{M^r}_{ni,t} / \bar{S}^r_{i,t}$ its destination sales weights, and

$$\kappa^r_{-i|n,t} \equiv \frac{\sum_{h \neq i} \bar{\pi}^{M^r}_{nh,t} \kappa^r_{h,t}}{1 - \bar{\pi}^{M^r}_{ni,t}}$$

the average learning value of competing suppliers in destination n.

An equipment subsidy redirects orders toward the supported origin and away from competing suppliers. A deployment subsidy instead expands equipment spending in the installing country and distributes the new orders according to that country's sourcing shares. The next proposition compares their direct learning effects.

Proposition 2 (Direct learning effects of the two subsidies)**.** *At the tariff-free, zero-subsidy reference with interior sourcing shares, hold current X, c, A fixed. The direct learning contribution of a marginal equipment subsidy at origin i is*

$$\begin{aligned} \mathcal{D}^{r,\Omega}_{i,t} &\equiv \sum_h \Omega^r_{h,t} \frac{\partial y^{M,r}_{h,t}}{\partial s^{M,r}_{i,t}} \bigg|_{X,c,A} \\ &= \bar{S}^r_{i,t} \Bigg[\underbrace{\kappa^r_{i,t}}_{\substack{\text{own-output} \\ \text{learning}}} + \underbrace{\theta^r \sum_n \chi^r_{ni,t} \left(1 - \bar{\pi}^{M^r}_{ni,t}\right) \left(\kappa^r_{i,t} - \kappa^r_{-i|n,t}\right)}_{\substack{\text{production} \\ \text{reallocation}}} \Bigg] . \end{aligned} \tag{23}$$

[11] Global welfare values the resulting country consumption changes using reference-expenditure weights. The planner instead values future productivity gains using the common shadow value of the numeraire supported by international transfers. Even at the same allocation, the two values coincide only when these welfare weights per dollar align at every date.

It is negative if and only if

$$\theta^r \sum_n \chi^r_{ni,t}\left(1-\bar{\pi}^{M^r}_{ni,t}\right)\left(\kappa^r_{-i|n,t}-\kappa^r_{i,t}\right) > \kappa^r_{i,t}.$$

It is nonnegative whenever $\kappa^r_{i,t} \geq \kappa^r_{-i|n,t}$ *in every destination served by origin* i*, and is strictly positive when, in addition,* $\kappa^r_{i,t} > 0$*. By contrast, one additional dollar of equipment spending induced by a deployment subsidy in destination* n *is routed according to*

$$\mathcal{Y}^r_{n,t} = \sum_i \frac{\bar{\pi}^{M^r}_{ni,t}}{\bar{C}^{M,r}_{i,t}}, \qquad \mathcal{Y}^{r,\Omega}_{n,t} = \sum_i \Omega^r_{i,t} \frac{\bar{\pi}^{M^r}_{ni,t}}{\bar{C}^{M,r}_{i,t}} = \sum_i \bar{\pi}^{M^r}_{ni,t}\kappa^r_{i,t}. \tag{24}$$

Here $\mathcal{Y}^r_{n,t}$ *is world equipment output per dollar of equipment purchaser spending, and* $\mathcal{Y}^{r,\Omega}_{n,t}$ *is the associated contribution of future learning to global welfare.*

Equation (23) separates scale from composition. Let $B^r_{i,t}$ denote its bracketed term. A large sales base $\bar{S}^r_{i,t}$ magnifies the effect of a given subsidy rate but does not determine its sign. At zero subsidy, marginal fiscal cost is also $\bar{S}^r_{i,t}$, so the direct learning return per marginal subsidy dollar is $B^r_{i,t}$; at a common marginal fiscal cost, origin i generates more direct learning value than origin j if and only if $B^r_{i,t} > B^r_{j,t}$. More precisely, the direct effect is positive if and only if the own-output term $\kappa^r_{i,t}$ exceeds the sales-weighted learning value lost by reallocating production away from higher-κ suppliers; equality makes the direct effect zero. An equipment subsidy in an importing market therefore lowers the direct learning contribution when the reverse inequality holds. A production hub need not satisfy the positive-effect condition: inherited experience can raise its sales base while lowering the marginal value of another unit of origin-specific experience. National welfare also includes the fiscal bill and international price incidence in equation (19).[12]

The deployment expression provides the parallel comparison. Conditional on the same dollar of additional equipment spending, destination n generates more direct learning value than destination m if and only if $\sum_i(\bar{\pi}^{M^r}_{ni,t} - \bar{\pi}^{M^r}_{mi,t})\kappa^r_{i,t} > 0$. Thus a destination has a larger $\mathcal{Y}^{r,\Omega}_{n,t}$ when its sourcing weights fall on origins that create more learning value per dollar; a destination with a larger displacement of fossil use has a larger carbon benefit. Equipment orders and production respond immediately, while new capacity and its first fossil-use response occur in the destination when that capacity comes online. The bilateral sourcing vector determines how induced origin-specific experience is distributed across suppliers, while its domestic element gives the share retained in the destination. In the data, China resembles a production hub with large origin receipts, exports, and inherited experience, whereas the United States resembles a large importing destination. These are measured positions in the supplier network rather than assumptions imposed on country types.

[12] Unlike Proposition 1, Proposition 2 gives conditional current-period routing derivatives, not total equilibrium responses to a policy rate. The transition solution determines the policy-induced change in equipment spending and the subsequent general-equilibrium feedbacks.

Why local electricity matters. These deployment results use the fact that renewable generation is local. New equipment orders can be filled by suppliers in any origin, but the installed capacity and its first displacement of fossil generation remain in the destination. If electricity were widely traded, the supplier-routing formulas would remain, but a deployment subsidy would operate through a second network: power trade would distribute generation, fossil displacement, electricity-price incidence, and generator revenue across connected markets. The installing country would then no longer identify the location of the first carbon response.

The results leave three quantitative questions: where each policy routes current equipment production, how strongly that production lowers future costs, and where lower renewable-investment costs displace fossil use. The calibration therefore measures bilateral sourcing and origin costs, the level and geographic scope of learning, equipment shares in installation cost, and renewable–fossil substitution. Section 6 then computes the finite, nonlinear transition, including endogenous wages, fiscal bases, investment, and later production. Tariff interactions enter only in the applications that restore observed policy wedges.

5 Calibration

The analytical results identify the moments that matter for policy assignment: bilateral sourcing determines where production moves; the level and geographic scope of learning determine the future cost response; and installation-cost composition and energy substitution determine where cheaper equipment displaces fossil use. I calibrate these objects 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 production and absorption margins; the wind matrix combines generating sets and selected parts with installation-based absorption. The resulting equipment expenditures are $139.6 billion for solar and $58.0 billion for wind, with China supplying 80.1 and 59.8 percent, respectively. 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. The benchmark accounts satisfy resource balance without an auxiliary residual-demand category.

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 parameter and local share are calibrated to the empirical range. Let $a^r_i = (A^{M^r}_i)^{1/\theta^r}$ denote the unit-productivity shifter used in equipment cost. The mapping into the model is

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

Thus a trade-elasticity sensitivity holds (g_r, b_r) fixed and remaps both structural learning parameters.

For solar, I estimate

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

where $P^{S,\text{real}}_{W,t}$ is the real world module price, $Q^{S,\text{beg}}_{W,t}$ 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. 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 $(\nu^S_{\text{tot}}, \lambda^S_L) = (2.501, 0.123)$.

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. Appendix B.5 reports the estimates, data construction, sensitivity to stock timing, initial experience, sample windows, and a policy-based instrument; Appendix Table B.4 summarizes the mapped learning parameters.

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. I therefore use $(\nu^W_{\text{tot}}, \lambda^W_L) = (1.00, 0.50)$, which implies a 14.6 percent price decline per doubling and lies within the empirical range (Wiser et al., 2024). The 0.50 weight on origin-specific experience is a rounded calibration. The analogous manufacturer-location regression gives an own-country coefficient near 0.6 but is based on only six producer clusters. Consistent with a substantial local component, 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 the much larger outside experience pool still produces meaningful aggregate spillovers. These firm-level estimates do not map directly into a country-level weight, so I use 0.50 as a midpoint and consider 0.25 and 0.75 in the robustness analysis.

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 capacity in 2021 and 2022 determines beginning- and end-of-2022 solar and wind stocks. Ember 2022 generation divided by end-of-2022 capacity determines the effective country–technology

generation coefficient z_n^r in equation (6). The benchmark assumes constant returns to installed capacity, $G_n^r = z_n^r K_n^r$.

The inner solar–wind aggregator is Cobb–Douglas, $\rho = 1$.[13] 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, $X^F_{n,2022} + E^R_{n,2022}$, is divided between households and production. Household energy expenditure is the residual

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

The household energy weight and the renewable share of the common energy composite are then

$$\mu^E_n = \frac{X^{E,H}_{n,2022}}{P^C_{n,2022} C_{n,2022}}, \qquad \omega^R_n = \frac{E^R_{n,2022}}{X^F_{n,2022} + E^R_{n,2022}}.$$

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^R_n 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^r_n is chosen so the capital law maps observed 2021 capacity and matrix-implied 2022 renewable-investment spending into observed end-of-2022 capacity. Ap-

[13] Appendix B.4 reports the suggestive estimate $\hat{\rho} = 0.645$ (s.e. 0.554). It is consistent with the Cobb–Douglas benchmark but is not used as a structural estimate.

pendix B.7 describes the regularization used for three small wind cells. 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).

I impose the fossil-supply elasticity $\varepsilon_F = 1$ and use the competitive mapping $\eta_F = \varepsilon_F/(1 + \varepsilon_F) = 0.5$ for the labor-payment share, as required by Proposition 1. Sensitivity exercises vary 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 reports the horizon extension.

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; geospatial slopes reported separately
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

5.5 Calibration Moments and External Diagnostics

Panel A of Table 2 compares the balanced equipment matrices with their source-data shares. These are calibration objects rather than validation moments: 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 reports external diagnostics from the model's reference transition, which retains observed 2022 tariffs but imposes no deployment or equipment subsidy after 2022. The 2023 origin-output comparison overlaps the exporter data used to estimate b_S, and 2024 lies within the price sample used to estimate g_S, so these are not holdout tests. The transition 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.

Table 2: Calibration Moments and External Diagnostics

Moment	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. Observed-tariff, non-targeted transition diagnostics			
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. Panel B uses the reference transition with observed 2022 tariffs and no post-2022 deployment or equipment subsidy. It is intended as a diagnostic rather than a holdout test. Model solar output combines cells and modules, whereas the observed China series covers modules.

6 Policy Assignment in Global Equipment Markets

The main exercise compares deployment and equipment subsidies in an importing market (the United States) and a production hub (China) at a common fiscal cost. It removes equipment tariffs while retaining calibrated technologies and nonpolicy trade costs and allowing sourcing to adjust.[14] 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 identify the 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. 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 $45.8 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.

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

Table 3: Policy Instruments and Locations at a Common Fiscal Cost

Policy	Rate (percent)	World log CEV (percentage points)	U.S. CEV	China CEV	Noncarbon value (2022 $bn)	Carbon benefit
Panel A: tariff-free benchmark with equipment trade						
U.S. deployment	12.83	0.0008	−0.0011	0.0048	13.8	76.6
U.S. equipment	30.00	≈ 0	−0.0025	0.0032	−0.3	62.9
China deployment	4.21	0.0012	0.0006	0.0067	21.8	112.8
China equipment	5.82	0.0020	0.0009	0.0044	36.2	118.7
Panel B: tariff-free benchmark without equipment trade						
U.S. deployment	14.36	0.0008	−0.0005	0.0048	13.6	77.5
U.S. equipment	26.85	0.0011	−0.0005	0.0065	20.6	86.2
China deployment	4.21	0.0011	0.0005	0.0069	19.9	111.3
China equipment	8.79	0.0016	0.0007	0.0090	29.3	123.5

Notes: Policies operate during 2023–2032. Each rate is solved within its trade regime to match a $45.8 billion fiscal present value, and each effect is measured relative to that regime's no-policy transition. "Deployment" denotes the deployment ITC and "equipment" the production subsidy. Panel B removes solar- and wind-equipment trade but retains fossil and non-energy trade and global learning. The noncarbon value converts the world log CEV using the present value of reference world consumption. The carbon benefit values discounted 2023–2100 emissions reductions at $210 per ton of CO_2 and is not included in the CEV columns.

Panel A reveals a location-dependent ranking while holding spending fixed. In the United States, deployment support raises global welfare by 0.0008 percent, while the effect of equipment support is economically close to zero. Both policies reduce U.S. consumption, but the loss is smaller under deployment support. A U.S. equipment subsidy expands a relatively small domestic production base partly by displacing output from established foreign suppliers.[15] China nevertheless gains from this policy. Its equipment output falls by $9.6 billion in present value during the subsidy window, but rises by $13.7 billion over the full horizon. Additional U.S. production raises the global experience stock and lowers subsequent equipment costs, so Chinese output rebounds and China benefits as both a large producer and a large renewable investor while the United States bears the fiscal cost. In China, equipment support produces the larger global gain, 0.0020 percent against 0.0012 percent for deployment, and slightly more abatement. Deployment support nevertheless produces the larger Chinese gain. The equipment subsidy lowers future costs across China's large export market, while the deployment subsidy raises domestic renewable investment and demand for China's own equipment sector. Across these four experiments, discounted emissions fall by 0.03–0.07 percent of discounted reference emissions through 2100. At $210 per ton, those reductions are worth $63–119 billion,

[15]The tariff-free benchmark removes explicit equipment tariffs but retains calibrated bilateral trade costs, including transport, supply-chain, and potentially non-tariff frictions. In its 2023 reference, China supplies 5.7 percent of U.S. solar-equipment purchases and 25.4 percent of wind-equipment purchases—15.2 percent of combined spending, or 22.7 percent of imported equipment spending. Tariff removal alone therefore does not make China the dominant U.S. supplier. Lower remaining frictions would make U.S. equipment support displace more Chinese output, worsening China's export loss and the negative learning-composition effect; the magnitudes and signs require re-solving the equilibrium.

compared with the common $45.8 billion fiscal cost.

Panel B removes equipment trade and re-matches each fiscal cost. The U.S. global ranking reverses: equipment support now yields 0.0011 percent, compared with 0.0008 percent for deployment. The U.S. welfare losses under the two instruments also become nearly identical. Without cross-origin sourcing, U.S. production support no longer draws output away from foreign suppliers, although global learning still transmits future cost reductions across countries. China's gain from U.S. equipment support consequently rises from 0.0032 to 0.0065 percent: China has no U.S. export sales to lose but continues to receive the global learning spillover. In equation (23), shutting down equipment trade removes the displacement term and leaves the positive own-origin learning term; the ranking reversal is consistent with this no-trade limit of Proposition 2. The China–U.S. difference in the global return to equipment support falls from 0.0020 to 0.0005 percentage points, while the global effect of U.S. deployment support remains almost unchanged. Equipment trade therefore matters both for instrument choice in an importing market and for the value of locating production support at an export hub.

6.2 International Emissions and Production Spillovers

The aggregate emissions effects conceal different foreign responses. Because electricity is local while renewable equipment and fossil fuels are traded, a subsidy reaches other countries through equipment demand and prices, production learning, and the world fossil price. Table 4 separates responses in the subsidizing country and abroad in the four equal-budget experiments.

Deployment subsidies concentrate emissions reductions in the subsidizing country but route equipment demand through its supplier network. U.S. emissions fall by 451 $MtCO_2$, while emissions in the rest of the world rise by 86 $MtCO_2$. The corresponding split for Chinese deployment support is -667 and $+130$ $MtCO_2$. Additional deployment draws equipment toward the subsidizing market, while lower world fossil demand reduces the fossil price and encourages fossil use elsewhere. Later learning partly offsets these forces. Under U.S. deployment support, foreign equipment output rises by $86 billion in present value, more than four times the U.S. increase; under Chinese deployment support, most of the production increase remains in China because it already supplies much of its own market.

Equipment subsidies concentrate production in the subsidized origin. U.S. and Chinese equipment output rises by $163 billion and $199 billion, while output in the rest of the world falls by $37 billion and $11 billion. The foreign emissions response is much smaller than under deployment support: it is a 5 Mt increase for the U.S. subsidy—economically close to zero—and an 82 Mt reduction for the Chinese subsidy. An equipment subsidy directly lowers the supported origin's effective delivered price in importing markets, and production learning extends that reduction into the future. These foreign benefits are strongest when support expands an origin that already serves a large foreign customer base. The model therefore does not imply that any equipment subsidy lowers

foreign emissions.[16] International markets transmit deployment support through both fossil-price leakage and clean-equipment crowd-out; under equipment support, lower current delivered prices and subsequent learning can offset or reverse those foreign effects.

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

	ΔCO_2 (Mt, discounted)		Δ equipment output (PV, 2022 $bn)	
Policy	Policy country	Rest of world	Policy country	Rest of world
U.S. deployment subsidy	−450.9	+86.1	+20.2	+86.3
U.S. equipment subsidy	−304.6	+5.0	+162.7	−37.3
China deployment subsidy	−666.8	+129.9	+102.7	+11.0
China equipment subsidy	−483.9	−81.5	+199.0	−10.7

Notes: These are the tariff-free equal-budget experiments in Table 3, Panel A, with direct energy use in non-energy and equipment production. CO_2 entries are discounted full-horizon changes; equipment entries are discounted changes in solar-plus-wind real output valued at 2022 unit costs.

Table 5 connects the quantitative results to the welfare envelope in Proposition 1. For each policy, it evaluates the marginal welfare components at intermediate subsidy rates and accumulates them from zero to the reported rate. The learning column combines the global and origin-specific components; Appendix Table C.1 reports them separately and describes the numerical integration.

Table 5: Welfare Channels at a Common Fiscal Cost

Policy	Intl. prices	Learning	Policy cost	Other	Total
Panel A: Global consumption welfare					
U.S. deployment subsidy	0.00	11.68	−4.72	0.66	7.63
U.S. equipment subsidy	−0.01	11.18	−11.93	0.58	−0.18
China equipment subsidy	0.01	20.79	−1.78	1.02	20.04
Panel B: Policy-country consumption welfare					
U.S. deployment subsidy	−2.51	7.48	−16.38	0.96	−10.47
U.S. equipment subsidy	9.96	11.90	−47.20	0.79	−24.55
China equipment subsidy	45.52	68.29	−72.93	3.34	44.21

Notes: Entries are in 10^{-4} percentage points of consumption-equivalent welfare. These are three tariff-free equal-budget experiments from Table 3. "Learning" combines the global and origin-specific components. "Other" collects finite-policy terms omitted by the local envelope; small discrepancies in displayed sums are due to rounding. Climate damages are excluded.

The global panel shows three different balances between learning and policy cost. Learning more

[16] This differs from the "reverse leakage" result in Kotchen and Maggi (2025). In their model, a subsidy to green-energy production lowers the common fossil-fuel price and therefore reduces foreign fossil output when green and fossil energy are substitutes. Here the subsidy applies to traded renewable equipment, so the direct and learning-induced declines in installation costs compete with contemporaneous sourcing, equipment-market displacement, and fossil-price responses, leaving the foreign-emissions effect unsigned.

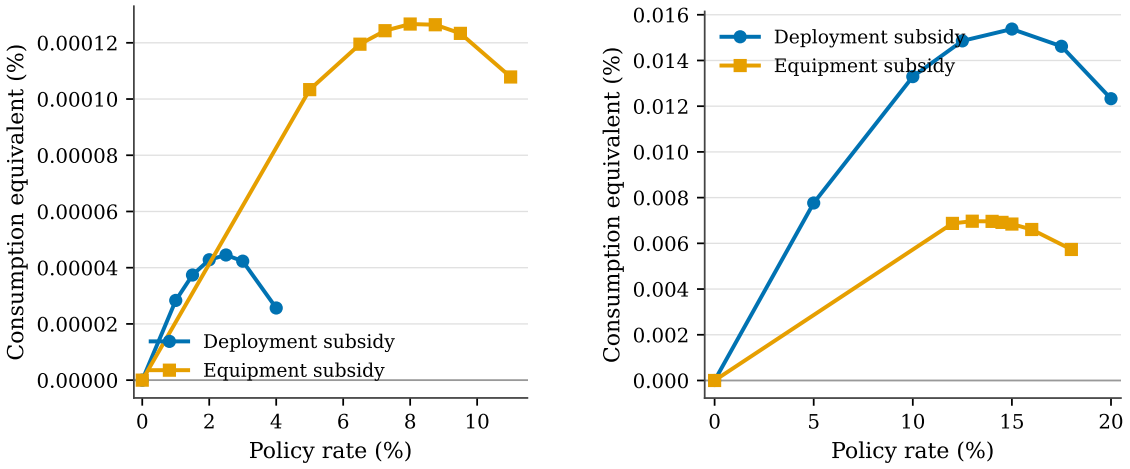
than covers the cost of U.S. deployment support, nearly equals the cost of U.S. equipment support, and is largest relative to cost for Chinese equipment support. International-price incidence nearly cancels globally but remains important for the subsidizing country. China captures enough favorable price incidence and learning to offset the cost of equipment support, whereas the United States does not under either instrument. The location ranking therefore reflects both the reach of the production base and the reallocation required to expand it.

6.3 Policy Scale and National–World Incidence

The equal-budget experiments compare policies at a common scale. Figure 3 instead reports tariff-free unilateral rate profiles for all four country–instrument pairs, with every other green policy set to zero. In the United States, welfare peaks near a 2.5 percent deployment subsidy and an 8 percent equipment subsidy. The peak gain under equipment support is about 2.8 times the peak gain under deployment support, although both are small. At those nationally preferred rates, however, deployment support produces the larger global gain: about 0.0002 percent, compared with 0.0001 percent under equipment support. The U.S. instrument ranking therefore differs from the global ranking even before the policies are scaled to a common budget. At the $45.8 billion anchor, the required rates rise to 12.83 and 30.00 percent, beyond the point at which either U.S. policy benefits the subsidizing country, and deployment support then produces the smaller loss. The distinction reflects curvature rather than a contradiction—the national return to U.S. equipment support rises at first but declines rapidly as the production distortion grows.[17]

In China, welfare peaks near a 15 percent deployment subsidy and a 13 percent equipment subsidy. The peak national gain is more than twice as large under deployment support: most of the resulting equipment demand falls on Chinese suppliers, so China combines cheaper renewable investment with domestic production learning. Equipment support instead lowers the prices paid by foreign buyers, and the associated international benefit is only partly captured by China. Global welfare under deployment support peaks near 17.5 percent. For Chinese equipment support, it peaks near 35 percent, bracketed by 30 and 40 percent. That rate has a fiscal present value of $634 billion—about fourteen times the common-budget experiment—and lowers Chinese consumption welfare by 0.0390 percent. The nationally preferred programs also differ sharply in fiscal scale: the U.S. peaks cost $4–6 billion, compared with $123–229 billion for the Chinese peaks. These are one-instrument national optima, not equal-effort comparisons; the much larger world-welfare benchmark illustrates the cross-border learning benefit that China does not internalize.

[17] Restoring observed equipment tariffs reverses the small-rate U.S. instrument ranking. The tariff wall raises the national return to deployment support and depresses the return to equipment support; I treat this as a second-best extension rather than the central assignment result.



(a) United States

(b) China

Figure 3: Tariff-Free Unilateral Policy Returns by Country and Instrument

Notes: Markers are solved unilateral grid points in the tariff-free benchmark; lines connect them for readability. Each panel reports the policy country's consumption equivalent and sets all other green policies to zero. Climate damages are excluded. The panels bracket each country's preferred rate; Appendix Table C.5 extends the Chinese equipment-subsidy grid to locate the global-welfare peak.

6.4 Joint Policies in an Importing Market and a Production Hub

The assignment problem also arises when both countries act. Panel A of Table 6 combines the equal-budget policies from Table 3. Each component uses the rate that costs $45.8 billion when implemented alone; the jointly realized bill differs slightly from twice that amount because each policy changes the other's tax base. U.S. deployment combined with Chinese equipment support produces the largest global gain, 0.0027 percent, the largest emissions reduction, 0.93 $GtCO_2$, and the smallest U.S. loss. Each country is better off than when implementing its own component alone, but the global interaction is mildly subadditive in all four cells. The result therefore reflects reciprocal international incidence rather than technological complementarity. The United States remains below its no-policy welfare in every equal-budget cell, so these packages are not self-enforcing without climate benefits or transfers.

Panel B evaluates the same four instrument pairs at each country's unilateral one-instrument preferred rate: 2.5 and 8 percent for U.S. deployment and equipment support, and 15 and 13 percent for Chinese deployment and equipment support. The fiscal scales are very different, so this panel addresses instrument assignment rather than equal effort. The world-best assignment again combines U.S. deployment with Chinese equipment support. Private instrument choices point in the opposite direction: at the displayed rates, the United States prefers equipment support for either Chinese instrument, and China prefers deployment support for either U.S. instrument. This selects U.S. equipment and Chinese deployment, the lowest-global-welfare corner among the four. The compar-

ison is a fixed-rate instrument-choice exercise rather than a full Nash equilibrium in instruments and rates.

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

U.S. instrument	China instrument	World log CEV	U.S. CEV	China CEV	Discounted CO_2	Fiscal PV	U.S./China fiscal PV
		(percentage points)			($MtCO_2$)	(2022 $bn)	
Panel A: rates matched to a $45.8 billion unilateral budget for each component							
Deployment (12.83%)	Deployment (4.21%)	0.0020	−0.0005	0.0114	−901	91.8	45.9 / 45.9
Deployment (12.83%)	Equipment (5.82%)	0.0027	−0.0002	0.0091	−932	92.9	46.4 / 46.5
Equipment (30.00%)	Deployment (4.21%)	0.0012	−0.0019	0.0098	−835	91.7	45.9 / 45.9
Equipment (30.00%)	Equipment (5.82%)	0.0020	−0.0015	0.0076	−850	89.7	44.8 / 44.9
Panel B: unilateral one-instrument preferred rates							
Deployment (2.5%)	Deployment (15%)	0.0035	0.0024	0.0161	−2,315	235.6	6.4 / 229.1
Deployment (2.5%)	Equipment (13%)	0.0046	0.0024	0.0077	−1,455	129.8	6.6 / 123.2
Equipment (8%)	Deployment (15%)	0.0034	0.0025	0.0155	−2,293	233.4	4.3 / 229.1
Equipment (8%)	Equipment (13%)	0.0045	0.0025	0.0072	−1,425	126.1	3.8 / 122.3

Notes: Both policies operate during 2023–2032 in the tariff-free benchmark. Panel A fixes each rate at its unilateral equal-budget value. Panel B fixes each rate at the country's unilateral one-instrument welfare peak. Joint fiscal bills need not equal the sum of unilateral bills because policies change each other's tax bases. Consumption equivalents exclude climate damages; negative CO_2 entries denote emissions reductions.

6.5 A Stylized IRA-Scale Package

For this positive policy application, I restore observed 2022 equipment tariffs. I represent the two principal margins of the IRA by a 30 percent deployment ITC and a 30 percent subsidy to domestic solar- and wind-equipment production.[18] Both operate during 2023–2032.[19] Table 7 reports the two components and the jointly solved package.

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

Experiment	Fiscal PV	World log CEV	U.S. CEV	China CEV	Discounted CO_2
	(2022 $bn)		(percentage points)		($MtCO_2$)
Deployment component, 30%	188.6	0.0003	−0.0100	0.0140	−1,109
Equipment component, 30%	52.4	0.0002	−0.0033	0.0055	−387
Combined 30/30 package	359.4	−0.0025	−0.0294	0.0278	−1,903

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.

[18] The deployment rate is the full statutory ITC rate under Sections 48 and 48E. Section 45X instead specifies component-level production credits rather than a uniform ad valorem rate. Along an integrated solar supply chain, these 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 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. Appendix B.1 describes the mapping and scope of the experiment.

[19] This is a stylized rate experiment rather than a legislative budget score. Statutory production credits are component specific, and the model covers only solar and wind.

The package has small aggregate consumption effects but substantial international incidence. It lowers the U.S. consumption equivalent by 0.0294 percent, raises China's by 0.0278 percent, and raises the Europe–United Kingdom aggregate by 0.0124 percent. Global welfare before carbon damages falls by 0.0025 percent. Discounted world emissions fall by 1.90 $GtCO_2$, about 0.2 percent of discounted reference emissions through 2100. Accounting for energy used in equipment and other production reduces the package's emissions effect by about 5 percent. Valuing the net reduction at the flat $210 benchmark gives about $400 billion in avoided damages and a 0.0196 percent increase in carbon-inclusive world welfare.

The deployment component raises U.S. equipment demand, whereas the equipment component shifts production toward U.S. suppliers. When combined, equipment support lowers renewable-investment costs and expands the deployment-credit base, while deployment raises equipment output and expands the production-subsidy base. The joint fiscal present value is therefore larger than the sum of the two component-only bills. The 30 percent equipment rate is a policy-scale benchmark, not an estimate of the first best.[20]

6.6 Trade Restrictions and Policy Assignment

I now restore observed 2022 equipment tariffs and consider a permanent restriction under which U.S. buyers cannot source solar equipment from China or Southeast Asia.[21] The restriction creates no tariff revenue or quota rents; wind equipment and other trade links remain unchanged.

A source restriction operates differently from an equipment subsidy. The subsidy lowers one origin's cost in every destination, whereas the restriction removes selected origins only from the U.S. market and may redirect purchases toward domestic or third-country suppliers. I first compare the restriction-only transition with the factual transition. I then re-match the two U.S. policy rates to the $52.4 billion fiscal target within each trade regime, measuring policy effects relative to no policy under the same regime. Table 8 reports both comparisons.

[20] Evaluating only the marginal learning benefit in Lemma 1 on the no-policy path gives equipment rates of about 38–39 percent for solar and 14–16 percent for wind. These rates neither characterize an optimal investment credit nor include the carbon tax. Appendix Table C.7 reports the calculation.

[21] Recent U.S. measures move in this direction without imposing a blanket exclusion. The Uyghur Forced Labor Prevention Act presumptively bars imports produced wholly or partly in Xinjiang or by listed entities, while the Department of Commerce has found that certain China-linked cells and modules completed in Cambodia, Malaysia, Thailand, and Vietnam circumvent the China solar AD/CVD orders (U.S. Customs and Border Protection, 2025; U.S. Department of Commerce, 2023). The experiment is more stringent than the 2023–2032 restriction reported in the appendix.

Table 8: Trade Restrictions and Policy Assignment

Trade regime	U.S. policy	Rate (percent)	World log CEV	U.S. CEV (percentage points)	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 in every row. Panel A is relative to the factual transition. Panel B is relative to no policy under the same trade regime, with rates re-matched to $52.4 billion. Thus the restricted-sourcing emissions entries in Panel B are incremental to a restricted no-policy path that already emits 140 $MtCO_2$ more than the factual path. The permanent restriction generates no revenue or rents.

The restriction alone raises the U.S. solar installation-cost index by about 8 percent and lowers U.S. consumption by 0.0065 percent. Southeast Asia supplies 66.8 percent of total U.S. solar-equipment purchases, or 83.7 percent of imports, whereas direct sourcing from China is small. Global consumption welfare falls by 0.0022 percent as installation costs rise and production and learning are reallocated, and discounted emissions increase by 140 $MtCO_2$.

The restriction weakens the deployment subsidy. Higher delivered equipment costs contract U.S. renewable investment, so matching the same fiscal outlay requires the subsidy rate to rise from 14.58 to 15.21 percent. Even at that higher rate, the U.S. loss widens from 0.0007 to 0.0009 percent and the world gain falls slightly. The subsidy must first offset the sourcing-induced increase in installation cost, while much of the additional equipment demand continues to accrue to permitted foreign suppliers.

The interaction with equipment support runs in the opposite direction. Import substitution expands the domestic production base, so the rate required to spend $52.4 billion falls from 30.00 to 28.33 percent. Production on this larger protected base raises U.S. origin-specific experience, reducing the U.S. loss from 0.0033 to 0.0020 percent and raising the world gain from 0.0002 to 0.0008 percent. The U.S. effect nevertheless remains negative. The restriction protects sales only in the U.S. market; it does not improve U.S. suppliers' access to foreign destinations, where established Chinese and Southeast Asian suppliers remain available. Moreover, the calibrated origin-specific weight in solar learning is 0.123, so most experience-induced productivity growth is shared across origins rather than captured by the expanding U.S. producer. China consequently receives a slightly larger gain under both restricted-sourcing policy experiments. This is consistent with its small direct exposure to the restriction—China supplies only 0.6 percent of U.S. solar-equipment purchases in the calibration—and its receipt of global learning spillovers; most of the direct trade loss falls instead

on Southeast Asia.

The larger emissions reductions in the restricted-sourcing rows of Panel B require a separate qualification. They are measured relative to the restricted no-policy path, not relative to the factual path. Because the restriction alone raises discounted emissions by 140 $MtCO_2$, combining it with the deployment or equipment subsidy reduces emissions by about 270 or 296 $MtCO_2$ relative to factual no policy. These reductions are smaller than the corresponding 398 and 387 $MtCO_2$ reductions under factual sourcing. The Panel B comparison therefore does not imply that shifting production to the United States lowers emissions because U.S. manufacturing is cleaner. It instead measures a somewhat larger marginal policy response from a higher-cost, lower-deployment baseline, reinforced in the equipment case by the larger domestic learning base. Appendix C.6 reports a China-only case, a temporary restriction, finite tariffs, and alternative allocations of tariff revenue.

6.7 Carbon Pricing and Green Industrial Policy

The preceding experiments target learning and the location of clean-equipment production. They do not price the broader carbon externality. Appendix C.8 therefore solves carbon-tax counterfactuals in which a flat $210 per ton tax on fossil use enters the equilibrium and the taxing region rebates the revenue lump sum. A worldwide tax reduces discounted emissions by 576 $GtCO_2$. Its noncarbon consumption cost is 2.6 percent of world consumption, but the avoided emissions valued at $210 produce a net global gain of about 4.1 percent. The much larger effect is not a like-for-like return comparison: the tax acts on fossil use throughout the economy, whereas the industrial subsidies operate through learning in a small traded-input sector.

The two instruments are complementary. Adding the local-Pigouvian learning schedule to the worldwide carbon tax raises the carbon-inclusive gain from 4.09 to 4.20 percent. The carbon tax prices emissions, while the equipment subsidies reward production experience that lowers future clean-investment costs. Because observed tariffs remain and the learning schedule is evaluated on the taxed transition rather than on a solved planner allocation, this exercise illustrates the division of labor in Lemma 1 but is not a first best.

Unilateral carbon pricing also generates international incidence. When the United States alone applies a $210 tax through 2100, discounted world emissions fall by 40.4 $GtCO_2$. U.S. consumption welfare falls by 1.35 percent before assigning the global climate benefit across countries. The noncarbon global effect is -0.35 percent, while avoided emissions valued at $210 contribute 0.47 percent, yielding a net global gain of 0.12 percent. Across the four evaluated U.S. tax rates, the global gain peaks near $100 at 0.18 percent. This is a global-damage valuation rather than an optimal U.S. tax. More moderate unilateral U.S. carbon prices produce smaller global emissions reductions, while leakage remains between 47 and 48 percent. The appendix reports the rate ladder, leakage through the world fossil market, and the accompanying reallocation of equipment production.

7 Conclusion

This paper studies how the instrument and location of green subsidies shape their international incidence when renewable equipment is traded but electricity generation remains local. Bilateral sourcing determines where a policy changes current production; global and origin-specific learning determine how those changes affect future costs; and local generation determines where renewable investment first displaces fossil use. In the benchmark, U.S. deployment support produces a positive global effect and a smaller U.S. loss than U.S. equipment support, whose global effect is approximately zero. In China, equipment support produces the larger global gain and more abatement, while deployment support produces the larger Chinese gain. Removing equipment trade reverses the U.S. global ranking and sharply narrows the difference between locating equipment support in China and in the United States. Equipment trade therefore affects policy choice and location, not only the distribution of a fixed aggregate gain. It also changes carbon leakage: deployment support raises foreign emissions in both locations, whereas lower delivered prices and subsequent learning under Chinese equipment support reduce emissions in foreign buyer markets.

Joint policies reinforce the assignment result but also reveal its strategic limit. At equal unilateral budgets, global welfare is highest when the United States supports deployment and China supports equipment production, although the United States remains worse off than under no policy. At nationally preferred one-instrument rates, private instrument choices instead select U.S. equipment and Chinese deployment support, the lowest-global-welfare combination among the four evaluated packages. Because the nationally preferred Chinese programs are far larger than the U.S. programs, this comparison concerns instrument assignment rather than common fiscal cost.

The IRA application shifts consumption welfare away from the United States and reduces discounted world emissions by 1.90 $GtCO_2$; its carbon-inclusive global effect is positive when those emissions are valued at $210 per ton. However, these industrial subsidies are not substitutes for carbon pricing. They target learning in a small traded-input sector, while a carbon price acts on the broader fossil-use margin. Their case instead rests on how production support and deployment support use international equipment markets to distribute learning, costs, and emissions reductions across countries.

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," *Working Paper*, 2023.

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.

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 electricity 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.

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 Electricity 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) 6
A.3 Energy Aggregation and Expenditure Shares 6
A.4 Consumption Welfare and the Ex-post Carbon Account 7

Appendix B Calibration: Data and Estimation **9**
B.1 Renewable Energy Provisions in the IRA 9
B.2 Geographic and Time Coverage 10
B.3 Equipment Trade and Benchmark Sourcing 11
B.4 Evidence on the Inner-Nest Elasticity 12
B.5 Learning Calibration and Supporting Estimates 13
B.6 Intermittency, Value Factors, and Annual Aggregation 17
B.7 Measuring the Initial Productive Stocks 18
B.8 Emissions Accounting 19
B.9 Social Cost of Carbon 20

Appendix C Mechanisms and Robustness **20**
C.1 Welfare Decomposition 20
C.2 Learning and Equipment Trade 21
C.3 Structural Sensitivity 23
C.4 Policy Duration and the PTC Mapping 23
C.5 Unilateral Policy-Rate Profiles 24
C.6 Source Restrictions and Tariff Alternatives 25
C.7 A Local Pigouvian Benchmark 27

C.8 Carbon Pricing in Equilibrium 27
C.9 Learning Stock, Horizon, and Accounting Revaluations 28

Appendix D Exact-Change Gravity and Computational Implementation **29**
D.1 Within-period trade block 29
D.2 Resource and Policy Accounting 31
D.3 Post-2100 Values and Numerical Solution 32
D.4 Solution Algorithm 33

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.

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

$$\beta^{t-1}\frac{d\log C_{n,t}}{dp} = \mu_{n,t}\left[\dot{E}_{n,t} - \bar{E}_{n,t}\frac{d\log P^{C}_{n,t}}{dp}\right]. \tag{A.2}$$

Before differentiating, consolidate the fiscal accounts. For a deployment subsidy, $(1-s^{I,r}_{n,t})C^{\text{proj},r}_{n,t}I^{r}_{n,t} + \Lambda^{I,r}_{n,t} = C^{\text{proj},r}_{n,t}I^{r}_{n,t}$; for an origin subsidy, $Y^{M^r}_{i,t} - \Lambda^{M,r}_{i,t} = \sum_n V^{M^r}_{ni,t}$. 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

$$\dot{E}_{n,t} - \bar{E}_{n,t}\frac{d\log P^{C}_{n,t}}{dp} = \mathcal{I}^{p}_{n,t} + \mathcal{L}^{p}_{n,t} + \mathcal{H}^{p}_{n,t} + \sum_{k\in\mathcal{K}}\left(\bar{r}^{k}_{n,t}\dot{K}^{k}_{n,t} - \bar{P}^{\text{inv},k}_{n,t}\dot{I}^{k}_{n,t}\right), \tag{A.3}$$

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 $k \in \mathcal{K}$,

$$\begin{aligned} \dot{K}^{k}_{n,t+1} &= a^{k}_{n,t}\dot{K}^{k}_{n,t} + \bar{\Phi}^{k}_{I,n,t}\dot{I}^{k}_{n,t}, & \bar{P}^{\text{inv},k}_{n,t} &= \bar{q}^{k}_{n,t}\bar{\Phi}^{k}_{I,n,t}, \\ \bar{r}^{k}_{n,t}\dot{K}^{k}_{n,t} - \bar{P}^{\text{inv},k}_{n,t}\dot{I}^{k}_{n,t} &= -\bar{q}^{k}_{n,t}\dot{K}^{k}_{n,t+1} + \left(\bar{r}^{k}_{n,t} + \bar{q}^{k}_{n,t}a^{k}_{n,t}\right)\dot{K}^{k}_{n,t}. \end{aligned} \tag{A.4}$$

Non-energy capital enters this construction as a third stock with the same adjustment technology, not as a special case. 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.2)–(A.4) yields

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

At the tariff-free reference, $\mathcal{H}^p_{n,t} = 0$, which proves equation (19). Appending the separately defined carbon account, multiplying by ϖ_n, and summing over countries gives equation (20).

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 the sufficient special case stated in Proposition 1 without imposing it on the quantitative baseline.

The learning term already includes the entire induced productivity path and must not be added again as an output co-state. To obtain an equivalent representation of the global learning component, differentiate equations (11) and (12). One unit of technology-r output at origin i and date t has direct learning value

$$\Omega^r_{i,t} = \frac{1}{\mathcal{B}_\beta}\sum_{u>t}\beta^{u-1}\frac{\nu^r_{\text{tot}}}{\theta^r}\underbrace{(\mu_{Q,r})^{u-t-1}}_{\text{retention to } u}\Bigg[\underbrace{(1-\lambda^r_L)\frac{\sum_{n,h}(\varpi_n/\bar{E}_{n,u})\bar{V}^{M^r}_{nh,u}}{\bar{Q}^r_{W,u}}}_{\text{world-experience value}} \tag{A.5}$$

$$+\underbrace{\lambda^r_L\frac{\sum_n(\varpi_n/\bar{E}_{n,u})\bar{V}^{M^r}_{ni,u}}{\bar{Q}^r_{i,u}}}_{\text{origin-experience value}}\Bigg]. \tag{A.6}$$

Thus the direct learning contribution to global welfare can equivalently be written $\sum_{r,i,t}\Omega^r_{i,t}\dot{y}^{M,r}_{i,t}$. The first term values the effect of origin output on the common experience stock; the second values its effect on origin-specific experience. International-price incidence remains in $\mathcal{I}^p$, and any interaction with a pre-existing tariff remains in $\mathcal{H}^p$.

The exact proof uses the infinite horizon. 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 an endogenous non-energy margin the first of these terms carries three capital blocks rather than two. Alongside the solar and wind terminal rows, $\mathcal{R}^{K}_{n,T}$ now contains the approximation error in the stationary non-energy terminal row and 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 which the finite implementation retains. Non-energy capital is by a wide margin the largest stock in every region, so this component plausibly dominates $\mathcal{R}^{K}_{n,T}$ even though the policies studied here act directly on the renewable margins and only indirectly on non-energy accumulation. An exact derivative of the post-2100 value function would make all three terms zero. The calculations do not separately identify these directional terms; horizon-extension and symmetric-policy exercises assess their joint magnitude.

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

$$\mathcal{B}^{L,r}_{i,t} = \sum_{u>t} m^{W}_{t,u}(\mu_{Q,r})^{u-t-1} g_r \left[(1-\lambda^{r}_{L})\frac{\sum_h Y^{M^r}_{h,u}}{Q^{r}_{W,u}} + \lambda^{r}_{L}\frac{Y^{M^r}_{i,u}}{Q^{r}_{i,u}}\right]. \tag{A.7}$$

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 $u > t$,

$$\frac{\partial \log A^{M^r}_{h,u}}{\partial y^{M,r}_{i,t}} = (\mu_{Q,r})^{u-t-1}\nu^{r}_{\text{tot}}\left[\frac{1-\lambda^{r}_{L}}{Q^{r}_{W,u}} + \frac{\lambda^{r}_{L}}{Q^{r}_{i,u}}\mathbf{1}\{h=i\}\right].$$

At a tariff-free planner allocation, the Eaton–Kortum envelope values a log productivity increase at origin h by $Y^{M^r}_{h,u}/\theta^{r}$ units of the date-u numeraire. Multiplying by the derivative above, summing over future dates and origins, and applying $m^{W}_{t,u} = \lambda^{W}_{u}/\lambda^{W}_{t}$ gives $\mathcal{B}^{L,r}_{i,t}$. Equations (5) and the equipment-subsidy budget imply $\Lambda^{M,r}_{i,t} = s^{M,r}_{i,t}Y^{M^r}_{i,t} = s^{M,r}_{i,t}\tilde{c}^{M,r}_{i,t}y^{M,r}_{i,t}$, so equating the subsidy per unit of real output to its marginal learning value gives the first equality in equation (21). When $0 \le \mathcal{B}^{L,r}_{i,t}/\tilde{c}^{M,r}_{i,t} < 1$, the specific learning wedge is implemented by the admissible ad valorem rate $s^{M,r,*}_{i,t} = \mathcal{B}^{L,r}_{i,t}/\tilde{c}^{M,r}_{i,t}$. Otherwise the same specific wedge remains the Pigouvian object, but the ad valorem parameterization with $1 - s^{M,r}_{i,t} > 0$ 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 program. An alternative instrument mix is equivalent only if it reproduces the origin-specific 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.

Policy routing. For equipment technology r, differentiation of the gravity share at fixed X, c, A gives

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

At a zero-subsidy reference, real origin output is

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

Differentiating this expression and substituting the share derivative gives

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

The supported origin gains both the sourcing response and the subsidy top-up; every competing origin loses sales wherever the two origins overlap. Multiplying by $\Omega^r_{h,t}$, summing over origins, and using $\bar{V}^{M^r}_{ni,t} = \bar{\pi}^{M^r}_{ni,t} \bar{X}^{M^r}_{n,t}$ yields

$$\begin{aligned} \mathcal{D}^{r,\Omega}_{i,t} &= \sum_n \bar{V}^{M^r}_{ni,t} \left[\kappa^r_{i,t} + \theta^r \sum_{h \neq i} \bar{\pi}^{M^r}_{nh,t} \left(\kappa^r_{i,t} - \kappa^r_{h,t} \right) \right] \\ &= \bar{S}^r_{i,t} \left[\kappa^r_{i,t} + \theta^r \sum_n \chi^r_{ni,t} \left(1 - \bar{\pi}^{M^r}_{ni,t} \right) \left(\kappa^r_{i,t} - \kappa^r_{-i|n,t} \right) \right]. \end{aligned}$$

This proves equation (23); rearranging its bracket gives the stated sign condition. Equation (22) follows directly from the price-index derivative and the equipment cost share in equation (4).

For deployment, one additional dollar of purchaser expenditure in destination n creates $\bar{\pi}^{M^r}_{ni,t}$ dollars of receipts at origin i in the tariff-free benchmark. Dividing by $\bar{c}^{M,r}_{i,t}$, then summing across

origins, gives $\mathcal{Y}^r_{n,t}$; weighting each origin by $\Omega^r_{i,t}$ gives $\mathcal{Y}^{r,\Omega}_{n,t} = \sum_i \bar{\pi}^{M^r}_{ni,t}\kappa^r_{i,t}$. This proves equation (24) and Proposition 2. With a pre-existing tariff, the corresponding origin receipt is $\bar{\pi}^{M^r}_{ni,t}/(1+\bar{\tau}^{M^r}_{ni,t})$, as in the observed-tariff applications.

A.2 Composite Intermediate Goods (Dixit–Stiglitz)

Composite producers in sector j aggregate a unit continuum of varieties $\omega^j \in [0,1]$ via

$$Q^j_{n,t} = \left[\int_0^1 f^j_{n,t}(\omega^j)^{\frac{\sigma^j-1}{\sigma^j}}\, d\omega^j\right]^{\frac{\sigma^j}{\sigma^j-1}},$$

where $\sigma^j > 1$ is the elasticity across varieties. Cost minimization yields variety demand

$$f^j_{n,t}(\omega^j) = \left(\frac{p^j_{n,t}(\omega^j)}{P^j_{n,t}}\right)^{-\sigma^j} Q^j_{n,t},$$

and the sectoral price index

$$P^j_{n,t} = \left[\int_0^1 p^j_{n,t}(\omega^j)^{1-\sigma^j}\, d\omega^j\right]^{\frac{1}{1-\sigma^j}}.$$

These expressions underpin the notation in the main text.

A.3 Energy Aggregation and Expenditure Shares

Households purchase the nested energy composite. The dual price indices corresponding to the quantity nests in Section 3.3 are

$$\begin{aligned} P^R_{n,t} &= \left[\sum_{r\in\mathcal{R}} \eta^r_n (P^{G,r,u}_{n,t})^{1-\rho}\right]^{1/(1-\rho)}, \\ P^E_{n,t} &= \left[\eta_n (P^R_{n,t})^{1-\psi} + (1-\eta_n)(P^F_{n,t})^{1-\psi}\right]^{1/(1-\psi)}. \end{aligned} \tag{A.8}$$

Here $P^{G,r,u}_{n,t}$ is the user price of electricity generated with technology r, paid by the energy aggregator. It is distinguished from the producer price $P^{G,r,p}_{n,t}$ received by the generator when a generation subsidy is present; the two prices coincide otherwise. The Cobb–Douglas limits apply at unit elasticities. Total spending on energy services is equation (D.5), and its fossil share is

$$\gamma^F_{n,t} = \left[1 + \frac{\eta_n}{1-\eta_n}\left(\frac{P^F_{n,t}}{P^R_{n,t}}\right)^{\psi-1}\right]^{-1}. \tag{A.9}$$

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

$$\sigma_{n,t}^{r} = \frac{\eta_n^r (P_{n,t}^{G,r,u})^{1-\rho}}{\sum_{h \in \mathcal{R}} \eta_n^h (P_{n,t}^{G,h,u})^{1-\rho}}.$$

Thus fossil, solar, and wind expenditures sum exactly to $X_{n,t}^{E}$. At $\rho = 1$, $\sigma_{n,t}^{r} = \eta_n^r$; only the outer renewable–fossil share responds to relative prices. Equations (A.8) and (D.6) give the corresponding zero-profit and market-clearing conditions. The common energy price enters household demand and the energy-input demand of non-energy and equipment producers; because the composite shares are properties of the CES aggregator, the share algebra applies to total energy expenditure $X_{n,t}^{E}$ regardless of its split across users.

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 $d_{n,t} = \log C_{n,t}^{p} - \log C_{n,t}^{0}$, $H_T = \sum_{t=1}^{T} \beta^{t-1} + \beta^T/(1-\beta)$, and define

$$\ell_n^{\text{cons}} = \frac{1}{H_T} \left[\underbrace{\sum_{t=1}^{T} \beta^{t-1} d_{n,t}}_{\text{modeled transition through } T} + \underbrace{\frac{\beta^T}{1-\beta} d_{n,T}}_{\text{post-}T\text{ welfare tail}} \right]. \tag{A.10}$$

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.

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:

$$\ell_W^Z = -\frac{1}{\overline{\mathcal{E}}_{W,T}} \sum_{t=1}^{T} \beta^{t-1} \underbrace{SCC_t (Z_t^p - Z_t^0)}_{\text{change in global dollar damages}}. \tag{A.11}$$

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 $q_{n,t}^{\text{dam},(s)}$ of global dollar damages, with $\sum_n q_{n,t}^{\text{dam},(s)} = 1$, the corresponding national money-metric component is

$$\ell_n^{Z,(s)} = -\frac{1}{\overline{\mathcal{E}}_{n,T}} \sum_{t=1}^{T} \beta^{t-1} q_{n,t}^{\text{dam},(s)} SCC_t \Delta Z_t. \tag{A.12}$$

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 implementation computes the global account directly; equivalently, the reference-expenditure-weighted average of the national accounts reproduces it exactly because the damage shares sum to one in every period. The noncarbon global consumption-welfare measure is

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

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 KKT weights. Its local counterpart is equation (20).

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 $R^r_{n,t,u}$ be the marginal date-u generation revenue produced by capacity installed at t. Let $q^r_{n,t} = P^{\text{proj},r}_{n,t} / \Phi^r_{I,n,t}$ denote the pre-credit installed-capital cost. The ITC rate that produces the same marginal investment incentive as the PTC is

$$s^{\text{eq},r}_{n,t} = \frac{\sum_{u>t} M_{n,t,u} \tau^{D,r}_{n,u} R^r_{n,t,u}}{q^r_{n,t}}. \tag{B.1}$$

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 exercise measures the incidence of policy wedges rather than producing a legislative budget score of the IRA.

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.

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 and is not a legislative budget score. 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 $45.8 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 and trade regime, I

solve for the subsidy rate that matches this target in equilibrium. The resulting bills differ from the target by at most $0.04 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. The separate observed-tariff source-restriction exercise retains its original $52.4 billion target so that it remains comparable to the equipment component of the IRA application.

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, Australia–New Zealand, Southeast Asia, South America, Africa, the Middle East, and a synthetic rest of the world.
Definition of the aggregated regions. The aggregated regions are defined as follows:

- Europe: the 27 European Union member states plus Norway, Switzerland, and the United Kingdom.
- Australia–New Zealand: Australia and New Zealand.
- Southeast Asia: Bangladesh, Brunei, Cambodia, Indonesia, Laos, Malaysia, Myanmar, the Philippines, Singapore, Thailand, and Vietnam.
- South America: Argentina, Brazil, Chile, Colombia, Costa Rica, and Peru.
- Africa: Angola, Cameroon, Côte d'Ivoire, the Democratic Republic of the Congo, Egypt, Morocco, Nigeria, São Tomé and Príncipe, Senegal, South Africa, and Tunisia.
- Middle East: Israel, Jordan, Saudi Arabia, Türkiye, and the United Arab Emirates.
- Rest of the world: 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. Table B.1 documents how the source ICIO industries enter these composites. Energy-producing mining identifies the fossil-extraction good; all remaining non-equipment industries, including background electricity, enter the non-energy composite. Solar and wind generation are local goods whose 2022 scale is calibrated from observed generation and capacity rather than ICIO industries. The equipment sectors use C26 and C28 as input–output proxies,

Table B.1: Sector Aggregation

Model sector	Code in OECD-ICIO	Industry in OECD-ICIO
Non-energy composite	A01_02	Agriculture, hunting, forestry
Non-energy composite	A03	Fishing and aquaculture
Non-energy composite	B07_08	Mining and quarrying, non-energy producing products
Non-energy composite	B09	Mining support service activities
Fossil fuel energy	B05_06	Mining and quarrying, energy producing products
Non-energy composite	D	Electricity, gas, steam and air conditioning supply
Solar equipment (proxy)	C26	Computer, electronic and optical equipment
Wind equipment (proxy)	C28	Machinery and equipment, nec
Non-energy composite	C10T12	Food products, beverages and tobacco
Non-energy composite	C13T15	Textiles, textile products, leather and footwear
Non-energy composite	C16	Wood and products of wood and cork
Non-energy composite	C17_18	Paper products and printing
Non-energy composite	C19	Coke and refined petroleum products
Non-energy composite	C20	Chemical and chemical products
Non-energy composite	C21	Pharmaceuticals, medicinal chemical and botanical products
Non-energy composite	C22	Rubber and plastics products
Non-energy composite	C23	Other non-metallic mineral products
Non-energy composite	C24	Basic metals
Non-energy composite	C25	Fabricated metal products
Non-energy composite	C27	Electrical equipment
Non-energy composite	C29	Motor vehicles, trailers and semi-trailers
Non-energy composite	C30	Other transport equipment
Non-energy composite	C31T33	Manufacturing nec; repair and installation of machinery and equipment
Non-energy composite	E–T (excl. D)	All remaining industries

while their 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 share of `850300` (parts of electric motors/generators) and exclude HS `841290`. 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 row sums implied by the PPML-installation matrix; 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

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

[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).

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 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.2 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), statistically consistent with one. Within the model, ρ governs how the renewable composite rebalances in response to relative cost changes. The counterfactual analysis maintains the Cobb–Douglas case ($\rho = 1$); I do not use the auxiliary regression as a structural estimate.

Table B.2: 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.

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

$$g_k \equiv -\frac{\partial \log P_t^k}{\partial \log Q_{W,t}^k}, \qquad b_k \equiv \frac{\partial \log \hat{S}_{n,t}^k}{\partial \log Q_{n,t}^k}.$$

The exact mapping into equation (12) is

$$\nu_{\text{tot}}^k = \theta^k g_k, \qquad \lambda_L^k = \frac{b_k}{\theta^k g_k} = \frac{b_k}{\nu_{\text{tot}}^k}.$$

Thus the world curve disciplines total price progress, while the exporter panel disciplines the local 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 transition-state builder 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

$$\log P_{W,t}^{S,\text{real}} = a - g_S \log Q_{W,t}^{S,\text{beg}} + \beta_P \log P_t^{\text{poly,real}} + u_t,$$

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 $\hat{\nu}_{\text{tot}}^S = 2.501$ at $\theta^S = 5.38$ and a 27.5 percent price decline per doubling. Table B.3 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 $\hat{S}_{n,t}^S$ from a PPML gravity equation with importer-by-year and country-pair fixed effects. The importer-year effects absorb the

Table B.3: 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. $\nu^S_{mathrmtot} = 5.38 g_S$ and $LR^S = 1 - 2^{-g_S}$.

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. 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.

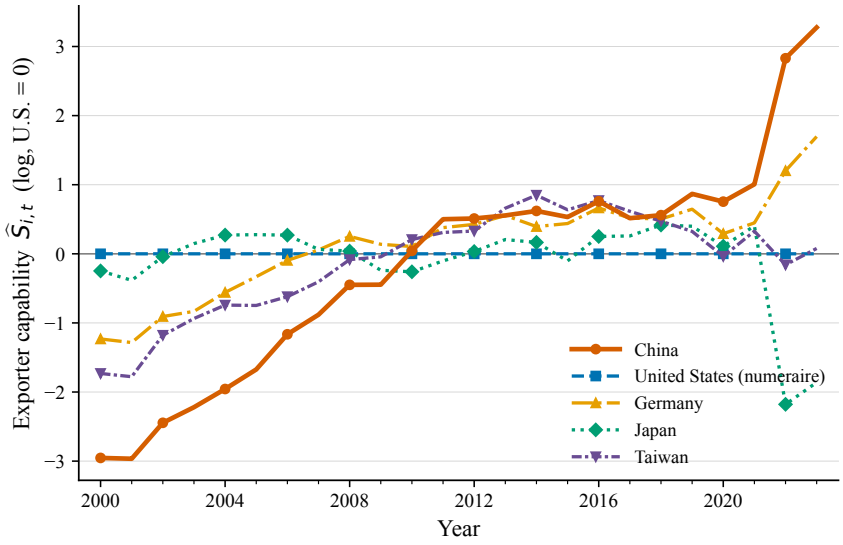


Figure B.1: Solar Exporter Capability

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.

China's 2018 policy reversal (descriptive). China's May 2018 "531" reform sharply reduced domestic 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; country production is constructed by mapping turbine-maker market shares to manufacturer-headquarters countries. 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.2 reports a joint lower envelope $(0.50, 0.25)$ and upper envelope $(1.20, 0.75)$, paired with the solar ranges described above.

Table B.4: 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. Wind values are calibration choices rather than comparable estimates.

These values are external calibration scenarios rather than estimates comparable to solar.

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 $Q^r_{W,2022} = \sum_n Q^r_{n,2022}$. When a trade-elasticity sensitivity is solved, the required remapping holds the empirical sufficient statistics fixed: 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.2 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 subannual dispatch, storage operation, curtailment, and hourly transmission constraints. The renewable–fossil substitution elasticity summarizes these integration costs in reduced form. 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.

The annual reduced form omits storage and within-day or seasonal dispatch. A fixed annual elasticity cannot steepen endogenously as renewable penetration grows. The model may therefore understate late-transition integration costs if value factors decline more sharply than the annual CES implies, or overstate them if storage becomes sufficiently cheap. The sensitivity analysis varies ψ between 2 and 4.

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 SCC-zero results in Section 6 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 SCC-zero tables exclude carbon damages, while equations (A.11) and (A.12) define the ex-post account 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 expands the main-text decomposition in Table 5 by separating global and origin-specific learning. For each policy, I evaluate the components of the marginal welfare envelope at intermediate subsidy rates and integrate them as the subsidy rises from zero to the rate used in the experiment. This converts the local envelope formula into a decomposition of the full policy effect: the entries add up to the welfare change between the no-policy transition and the equilibrium with the reported

subsidy. The two U.S. rows isolate the instrument comparison in an importing market; the China equipment row isolates production support at the export hub. The international-price term values changes in export receipts and import costs, the two learning terms value the productivity effects of world and origin-specific experience, and the policy-cost term records the deployment or production wedge net of the associated transfer. Carbon is excluded and reported separately in the main text.

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

Policy	Intl. prices	Global learning	Local learning	Policy cost	Remainder	Total
Panel A: Global consumption welfare						
U.S. deployment subsidy	0.00	9.31	2.37	−4.72	0.66	7.63
U.S. equipment subsidy	−0.01	9.42	1.76	−11.93	0.58	−0.18
China equipment subsidy	0.01	16.65	4.14	−1.78	1.02	20.04
Panel B: Policy-country consumption welfare						
U.S. deployment subsidy	−2.51	3.51	3.97	−16.38	0.96	−10.47
U.S. equipment subsidy	9.96	3.53	8.37	−47.20	0.79	−24.55
China equipment subsidy	45.52	48.25	20.04	−72.93	3.34	44.21

Notes: Entries are in 10^{-4} percentage points of consumption-equivalent welfare. The policies are the energy-extended, tariff-free equal-budget experiments in Table 3: rates are 12.83 percent for U.S. deployment, 30 percent for U.S. equipment, and 5.82 percent for China equipment. The decomposition uses 12-point Gauss–Legendre integration along the policy path. "Policy cost" is the ITC wedge for deployment and the buyer gain from the production wedge less the marginal fiscal outlay for equipment support. The remainder makes the finite-policy accounting exact and is at most 4.6 percent of the sum of absolute displayed channels; numerical quadrature errors are negligible. Climate damages are excluded.

Learning is positive in every row, but the scale of the production base determines how much new experience a fixed budget creates and where its value accrues. For U.S. equipment support, the global learning benefit and policy cost nearly cancel. The United States receives favorable international-price and learning effects, but they are too small to cover the fiscal cost of expanding domestic production. China equipment support produces more world learning for the same budget, and China's large sales base lets China capture substantial trade-price and learning gains. This accounting supports the main assignment result without treating the displayed channels as an exact first-order decomposition at a large policy: the remainder records the finite-policy terms omitted by the local envelope.

C.2 Learning and Equipment Trade

Table C.2 recomputes the four observed-tariff equal-budget policies under joint lower and upper learning envelopes. Within each scenario, the policy rate is re-matched to 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; at the upper envelope, China's domestic ranking switches from deployment to equipment support.

Table C.2: Equipment-Learning Robustness with Observed Tariffs

	Lower learning	Observed-tariff benchmark	Upper learning
Panel A: Fiscal-matched policy rate (percent)			
U.S. deployment credit	14.66	14.58	14.55
U.S. equipment subsidy	30.46	30.00	29.39
China deployment credit	4.71	4.68	4.69
China equipment subsidy	6.81	6.75	6.72
Panel B: Global consumption CEV (percent)			
U.S. deployment credit	0.0004	0.0008	0.0019
U.S. equipment subsidy	−0.0004	0.0002	0.0015
China deployment credit	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 credit	−0.0012	−0.0008	−0.0001
U.S. equipment subsidy	−0.0043	−0.0035	−0.0021
China deployment credit	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 (ν_S, b_S) of $(2.04, .221)$, $(2.50, .308)$, and $(3.52, .408)$ and wind pairs (ν_W, λ_W) of $(.50, .25)$, $(1, .50)$, and $(1.20, .75)$. Rates are fiscal-matched; climate damages are excluded.

Separate observed-tariff structural ablations isolate the mechanism rather than parameter uncertainty. Shutting down learning makes the global effect negative for the three policies examined. Retaining only origin-specific learning leaves only Chinese equipment support slightly positive. The tariff-free no-trade comparison in Table 3 provides the cleaner test of equipment trade because it re-matches fiscal costs within each trade regime.

The IRA-scale components also differ in how directly they move production. Equipment support shifts a larger share of solar and wind production toward the United States than deployment support. The combined package produces the largest shift and lowers the investment-weighted installation-cost index by 4.7 percent in 2032.

The technology split reinforces this interpretation. Solar accounts for most of the aggregate gain from U.S. deployment and Chinese equipment support, consistent with its more global learning and higher trade responsiveness. Wind's larger origin-specific learning weight directs more of its benefit toward the producing country. Because the technologies differ along several calibrated dimensions, this comparison is not a one-parameter experiment.

C.3 Structural Sensitivity

Table C.3 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.

Table C.3: Structural Sensitivity with Observed Tariffs

Parameter setting	U.S. deployment credit	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.

The small U.S. aggregate effects are parameter-sensitive. Across $\psi \in \{2,3,4\}$, the U.S. effect of its deployment credit is -0.0002, -0.0001, and 0.0002 percent. 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.

C.4 Policy Duration and the PTC Mapping

Table C.4 shows that longer support raises the global-welfare effect of a U.S. investment credit but at diminishing returns per fiscal dollar. As the window expands from five years to a permanent policy, 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.

Table C.4: Policy Duration and PTC Mapping with Observed Tariffs

Experiment	Policy horizon	World log CEV (percent)	Fiscal PV ($bn)	Discounted emissions ($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.

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.5 reports selected points from the tariff-free unilateral grids behind Section 6.3. U.S. welfare peaks near 2.5 percent for deployment support and 8 percent for equipment support. Chinese welfare peaks near 15 percent for deployment support and 13 percent for equipment support. Extending the Chinese equipment-subsidy grid locates the global-welfare peak near 35 percent, between lower values at 30 and 40 percent. These are one-instrument profiles with all other green policies set to zero.

Table C.5: Tariff-Free National and Global Welfare across Policy Rates

Policy	Rate (percent)	CEV (10^{-4} percentage points)	
		Global	National
U.S. deployment subsidy	1.0	0.7335	0.2836
	2.0	1.4524	0.4283
	2.5	1.8057	0.4452
	3.0	2.1547	0.4232
	4.0	2.8381	0.2567
U.S. equipment subsidy	5.0	0.4551	1.0334
	7.25	0.6373	1.2430
	8.0	0.6938	1.2667
	8.75	0.7479	1.2643
	11.0	0.8931	1.0785
China deployment subsidy	5.0	14.1388	77.6522
	10.0	25.6314	132.9350
	15.0	32.8364	153.7500
	17.5	34.1415	146.2040
	20.0	33.4430	123.2759
China equipment subsidy	12.0	41.2919	68.7212
	13.0	44.6633	69.6810
	14.0	47.9988	69.6124
	16.0	54.5319	66.0560
	18.0	60.8245	57.3181
	20.0	66.7979	42.5490
	25.0	79.6839	−27.6008
	30.0	87.9898	−162.0930
	35.0	88.7986	−389.6325
	40.0	77.7817	−751.0798

Notes: Each row varies one policy in isolation in the tariff-free benchmark. Consumption equivalents exclude climate damages. Figure 3 plots the subset needed to bracket each country's preferred rate; the additional Chinese equipment-subsidy rows locate the global-welfare peak.

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.6 compares this upper-bound experiment with a China-only restriction, a temporary restriction, finite tariffs, alternative ownership of tariff revenue, and a restricted-sourcing case in which the pre-existing 2022 equipment tariffs are also removed. Policy rates are re-matched to the common fiscal target in each regime.

Table C.6 reports five variants. The China-only restriction has little effect because direct calibrated U.S. sourcing from China is small; most of the main restriction's cost comes from excluding Southeast Asia. A restriction confined to 2023–2032 preserves the signs and is smaller 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, also removing the pre-existing equipment tariffs does not change the signs. Adding energy use in production leaves this result essentially unchanged: in the tariff-free economy, 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 these variants, the deployment point estimate remains larger and the U.S. loss remains smaller, although the gap in global-welfare effects narrows under broad restrictions. China gains under both instruments.

Table C.6: Source-Restriction Robustness Variants

Trade regime	U.S. policy	Rate (percent)	World log CEV	U.S. CEV	China CEV	Discounted CO_2 ($MtCO_2$)
			(percentage points)			
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 credit	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 credit	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 credit	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 credit	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 and re-matches rates to $52.4 billion. The windowed case reopens after 2032; negative emissions entries denote reductions.

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 common and local experience weights in equation (12), and divide the resulting per-unit learning benefit by the origin unit cost. Table C.7 reports both the first-period rate and the origin-output-weighted average over the policy window.

Table C.7: 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.

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 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 the local-Pigouvian schedule of Appendix Table C.7 evaluated at the fixed point on the taxed path.

Table C.8: 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$)	Leakage (%)
Worldwide carbon tax, $210/t	−2.60	—	+6.69 / 121,026	+4.09	−576.3	—
plus local-Pigouvian learning schedule	−2.79	—	+6.99 / 126,478	+4.20	−602.3	—
plus IRA-scale package	−2.66	—	+6.75 / 122,170	+4.09	−581.8	—
Unilateral U.S. carbon tax, $50/t	−0.04	−0.147	+0.18 / 3,187	+0.14	−15.2	47.0
Unilateral U.S. carbon tax, $100/t	−0.12	−0.476	+0.30 / 5,441	+0.18	−25.9	47.4
Unilateral U.S. carbon tax, $150/t	−0.22	−0.872	+0.39 / 7,068	+0.17	−33.7	47.7
Unilateral U.S. carbon tax, $210/t	−0.35	−1.354	+0.47 / 8,479	+0.12	−40.4	47.8

Notes: All entries are relative to the no-policy transition and cover 2023–2100. Dollar values in the experiment labels are carbon-tax rates; the carbon benefit in every row values discounted world emissions at $210 per ton. The noncarbon column is the Negishi-weighted world consumption-equivalent change; values under the baseline-expenditure weights are nearly identical. U.S. consumption CEV excludes climate damages. The world total adds the carbon benefit to noncarbon world welfare. Leakage is the rest-of-world emissions increase divided by the U.S. emissions reduction and is reported only for unilateral taxes. Tax revenue is rebated lump-sum to the taxing region's household. Observed tariffs are retained, so the rows are constrained benchmarks rather than a first best.

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 leakage rate of 47.8 percent over 2023–2100—as the world fossil producer price falls by about 6.2 percent. Leakage 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. The induced equipment demand is supplied mainly by foreign producers, so the unilateral carbon tax transmits abroad through the same margins as deployment support: abatement concentrates at home while equipment orders and the associated production experience accrue largely to exporters. 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.

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; cells whose benchmark effects are already close to zero move more in relative terms while remaining close to zero. Table C.9 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

Table C.9: Observed-Tariff Robustness to the Explicit Transition Horizon

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

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-Change Gravity and Computational Implementation

D.1 Within-period trade block

Each traded sector contains a continuum of constant-returns varieties with independently Fréchet productivity,

$$\Pr\left[z^j_{i,t}(\omega) \le z\right] = \exp\left[-A^j_{i,t} z^{-\theta^j}\right], \qquad j \in \mathcal{J}^T.$$

Equipment producers use labor, the non-energy composite, and local energy services,

$$c^{M^r}_{i,t} = B^r_i w^{\zeta^{L,r}_i}_{i,t} (P^E_{i,t})^{\zeta^{E,r}_i} (P^O_{i,t})^{1-\zeta^{L,r}_i-\zeta^{E,r}_i}, \qquad r \in \mathcal{R}.$$

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. For delivered cost $\kappa^j_{ni}(1+\tau^j_{ni,t})(1-s^j_{i,t})c^j_{i,t}$, the Eaton–Kortum price index and expenditure share are

$$P^j_{n,t} = \gamma^j \left[\sum_i A^j_{i,t}\left(\kappa^j_{ni}(1+\tau^j_{ni,t})(1-s^j_{i,t})c^j_{i,t}\right)^{-\theta^j}\right]^{-1/\theta^j},$$
$$\pi^j_{ni,t} = \frac{A^j_{i,t}\left(\kappa^j_{ni}(1+\tau^j_{ni,t})(1-s^j_{i,t})c^j_{i,t}\right)^{-\theta^j}}{\sum_m A^j_{m,t}\left(\kappa^j_{nm}(1+\tau^j_{nm,t})(1-s^j_{m,t})c^j_{m,t}\right)^{-\theta^j}}. \tag{D.3}$$

The reference and policy transitions are solved relative to the 2022 historical boundary. Exact changes are used only inside 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 imply

$$\widehat{P}^j_{n,t} = \left[\sum_i \bar{\pi}^j_{ni}\widehat{A}^j_{i,t}(\widehat{d}^j_{ni,t})^{-\theta^j}\right]^{-1/\theta^j},$$
$$\pi^j_{ni,t} = \frac{\bar{\pi}^j_{ni}\widehat{A}^j_{i,t}(\widehat{d}^j_{ni,t})^{-\theta^j}}{\sum_m \bar{\pi}^j_{nm}\widehat{A}^j_{m,t}(\widehat{d}^j_{nm,t})^{-\theta^j}}. \tag{D.4}$$

Equation (D.4) 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

$$\begin{aligned}X^F_{n,t} &= P^F_{n,t}F^d_{n,t}, \\ X^O_{n,t} &= P^O_{n,t}(C^O_{n,t}+I^D_{n,t}) + \sum_r\left[(1-\omega^r)C^{\text{proj},r}_{n,t}I^r_{n,t} + m^r_n P^O_{n,t}K^r_{n,t}\right] \\ &\quad + \sum_r (1-\zeta^{L,r}_n-\zeta^{E,r}_n)Y^{M^r}_{n,t} + (1-\gamma^V_n-\zeta^{E,O}_n)Y^O_{n,t}, \\ X^E_{n,t} &= P^E_{n,t}\mathcal{E}_{n,t} = P^E_{n,t}C^E_{n,t} + \zeta^{E,O}_n Y^O_{n,t} + \sum_r \zeta^{E,r}_n Y^{M^r}_{n,t}.\end{aligned} \tag{D.5}$$

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

$$P^F_{n,t}F^d_{n,t} = \gamma^F_{n,t}X^E_{n,t}, \qquad \sum_r P^{G,r,u}_{n,t}G^r_{n,t} = (1-\gamma^F_{n,t})X^E_{n,t}, \tag{D.6}$$

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,

$$P^{G,r,p}_{n,t} = (1+\tau^{D,r}_{n,t})P^{G,r,u}_{n,t}, \qquad \Lambda^{PTC,r}_{n,t} = \tau^{D,r}_{n,t} P^{G,r,u}_{n,t} G^{r}_{n,t}. \tag{D.7}$$

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.10). The private terminal Euler equation and welfare tail approximate values beyond 2100; the horizon exercises in Appendix C assess their quantitative importance. The carbon account ends in 2100 and carries no stationary tail.

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.4); equipment demand and origin revenue in equations (4) and (3); the energy mix and dual prices in equation (A.8) with the clearing condition (D.6); the absorption identities in equation (D.5); 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.6) 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. Thus a trial sequence determines the static allocations and the forward state path, and the residual vector records whether those allocations also satisfy factor-market clearing and the forward-looking investment conditions.

Following the sequence-space approach of Auclert et al. (2021), I solve this stacked system with a damped Newton method. A stationary approximation followed by damped time iteration supplies the initial reference path. I compute a finite-difference Jacobian of the residual vector with respect to the unknown sequences at that path, factorize it once, reuse the factorization across policy counterfactuals, and rebuild it only if convergence stalls; large policy changes are introduced gradually through homotopy in the policy rate. 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. Jacobian reuse makes the policy-rate grids and the path-integral decompositions tractable: a counterfactual solved from a converged neighboring policy typically requires only a handful of Newton steps.