

A Single Market for Electricity?

Aggregate Productivity Gains through Production Reallocation*

Paula Patzelt[†]

LSE

Job market paper
September 24, 2026

Latest version online

Abstract

When an input cannot be traded freely, production may sort across locations by input costs rather than productivity advantages. This paper estimates the aggregate productivity cost of electricity market fragmentation in Europe. Electricity price differences reflect both natural advantages and policy choices and therefore are not necessarily inefficient. I develop a model in which gains from integration depend on how much production moves, where it moves, and how national electricity prices respond. To estimate how much production reallocates, I use a quasi-natural experiment in Norway, where new interconnectors exposed the south to Europe's gas crisis, opening unexpectedly large and persistent regional price gaps. Production shifts out of high-price regions in proportion to sector electricity intensity, almost entirely through the number of active firms rather than the scale of continuing ones. In the calibrated 23-economy model of Europe, a full electricity union raises manufacturing productivity by 13 basis points and by 35 bp in the most electricity-intensive divisions. Over 80% of the gain comes from production reallocating towards more productive locations and clusters rather than firms adjusting in place. The value of integration almost triples by 2040 as industry electrifies and the expansion of renewables widens price dispersion. Integration also insures against energy crises and halves generation emissions at current carbon prices.

JEL codes: Q43, Q41, E23, O47, R32.

Keywords: Energy Markets, Productivity, Production Location, Spatial Misallocation.

*For their invaluable advice I am particularly indebted to my supervisor Ricardo Reis and my advisors Maarten De Ridder and Jonathon Hazell. For helpful comments and discussions I am grateful to Philippe Aghion, Robin Burgess, Swati Dhingra, Isabel Figueiras, Alessandro Gavazza, Luca Looser, Benjamin Moll, Mar Reguant, Emiliano Rinaldi, Silvana Tenreyro, and John Van Reenen. Special thanks to Ranik Raaen Wahlstrøm for sharing data on financial statements of companies in Norway. I also thank seminar and conference participants at LSE Macroeconomics, EEE Energy Markets Mini-Conference, EEE Hour, QMUL PhD Workshop, HEC Economics PhD Conference, XXIX Workshop on Dynamic Macroeconomics, CESifo Junior Workshop in International Economics, 18th Nordic Summer Symposium in Macroeconomics and Finance, and Oslo Macro Conference.

[†]Contact: p.h.patzelt@lse.ac.uk.

1 Introduction

A location's cost advantage in an input can be used in two ways: export the input, or attract production that uses it. The two are substitutes, so when the input cannot be traded, production moves instead (Samuelson, 1948, Mundell, 1957). Both are subject to frictions: infrastructure and regulation set the cost of trade, and sunk investment and agglomeration attach a firm to its location. These frictions allow input price gaps to persist. Production then sorts partly on where inputs are cheap rather than on where it is most productive, so the same firms and factors produce lower aggregate output than they could (Restuccia and Rogerson, 2008, Hsieh and Klenow, 2009). But price differences are not exogenous wedges, instead they can reflect natural advantages and change as production moves, so the productivity loss cannot be read off observed dispersion. With market integration, a location can export the input rather than attract production. Quantifying the resulting productivity gain requires a model of how strongly production responds to price differences and how much of the cost advantage survives reallocation.

Electricity is a clear example of such an input: despite being uniform, its price varies widely across space, even within the EU and US, where most goods move freely. Increasing reliance on renewables is tying electricity prices to geography because generation depends on local resources and electricity cannot be stored long-term at scale, so its price is set by the system generating it. Prices reflect both natural advantages in hydro, wind, and solar and policy choices over the generation mix, including nuclear reliance and renewable subsidies. Limited cross-border transmission allows price gaps to persist, and the generation mix determines how much they narrow when production moves towards cheaper power. Within Europe, wholesale prices differ persistently across national borders (Figure 1), while cross-border trade covers only about 5–10% of consumption. Grid operators plan major investments in transmission to integrate these markets, but appraisals focus on expected electricity-system benefits, leaving the effect on where production locates unquantified.

This paper asks what integrating an input market is worth in productivity when production moves, using electricity in Europe to answer the question. I measure how much productivity Europe loses under its fragmented electricity market relative to full integration. I then ask how much realistic transmission expansion could recover, and how the gains change as renewables expand and industry electrifies. This puts a number on the productivity gain from integration that policy reports miss, due to the change in where production takes place. In the pre-crisis European electricity system, full integration raises manufacturing productivity by 13 basis points even though the average electricity price does not fall. More than 80% of the gain comes from production reallocation. In the projected 2040 economy, electricity is cheaper overall, but a full union raises manufacturing productivity by 35 basis points, nearly three times today's gain.

The paper proceeds in three steps. First, I develop a model of firms choosing their scale and operating location across segmented electricity markets. They trade off electricity and labour costs,

weighted by sector-specific intensities, against inherent productivity advantages and agglomeration externalities. In equilibrium, agglomeration pulls firms into country clusters, but entry also bids up local wages and electricity prices. The strength of electricity congestion depends on a country's generation mix: prices respond strongly when fossil plants are marginal but remain flat when renewables or nuclear are marginal. The mix, shaped by natural resources and policy, determines not only the initial price level but also the capacity of a country to attract more production before its price advantage erodes.

The value of market integration in the model relies on three steps. The location multiplier governs how much production moves for a given cost difference, netting the agglomeration pull against firms' attachment to particular places. Country-specific electricity supply elasticities, together with the wage response, determine how much of the initial cost difference survives as production moves. The aggregate productivity effect then depends on where production ends up, how concentrated it becomes, how evenly costs are spread, and how much fuel the electricity system uses. I measure the supply side and calibrate the productivity channels, but need to estimate the location multiplier from firms' responses to persistent price differences.

The second step estimates how firm activity responds to persistent electricity price gaps using a quasi-natural experiment in Norway. New interconnectors between Southern Norway and continental Europe came online in 2021, just as the energy crisis drove continental prices to record highs. Despite no direct fossil reliance of the hydro-rich Norwegian electricity system, the increase in continental demand to substitute for gas raised prices in the Norwegian south, much more than expected when the cables were commissioned. Limited north-south transmission capacity kept prices in Northern Norway low, and the resulting price gap persists beyond the crisis. Within this common national setting, I compare sectors with different pre-crisis electricity intensity exposed to the same zone-level price shock, accounting for shocks shared across sectors in a zone and common sector trends.

Where electricity became expensive, production reallocated away from electricity-intensive activity. Municipality-level electricity use, revenue, and profits all rise less in high-price zones, while total firm counts barely differ: the adjustment is compositional, shifting activity toward less electricity-intensive sectors as exposed firms exit or contract. The model-aligned specification puts an elasticity on this response: firm numbers fall with electricity costs with an elasticity of about sixteen times the sector's electricity cost share, so a 10% price increase leaves roughly 2% fewer firms in the median sector and 6% fewer in basic metals. Almost all of the response runs through firm numbers rather than the rescaling of continuing firms. Existing evidence, largely from transitory or firm-specific energy-cost shocks, finds that firms absorb them in place. A persistent, location-wide price gap instead triggers reallocation through the extensive margin, which governs long-run industrial geography.

To assess whether the Norwegian evidence applies across European countries, I perform two

checks. Price gaps across national borders are similar to those across bidding-zone boundaries within countries in that they largely depend on transmission capacity, and the Norwegian divergence is large enough to cover typical cross-border gaps. In the European Orbis panel, profits respond more to electricity prices in sectors that use more of it, particularly when price differences persist. These patterns support the cost-exposure mechanism beyond Norway, while the Norwegian experiment remains the source of the location response used in the model.

The third step calibrates the model to 23 European economies and thirteen sectors, using the estimated location multiplier. Electricity prices and wages adjust as production shifts and electricity is traded across countries up to transmission capacities. I compare the current and planned 2040 grids with a full electricity union, where a single price clears the European market. A country's cheap power can then be exported through the grid instead of attracting production. The model measures what this change is worth on terms of productivity.

Electricity is 1.7% of manufacturing gross output, so a uniform 1% fall in electricity prices would raise manufacturing productivity by about 1.7 basis points through the cost channel alone.¹ On the 2015–2019 generation mix, a full electricity union raises manufacturing productivity by 13 basis points. This is equivalent to electricity becoming eight percent cheaper everywhere, although the average demand-weighted price does not fall in the union. Within the four most electricity-intensive divisions the integration gain is 35 bp.² Freezing firm locations and equalising prices delivers only 2 bp of the manufacturing gain, so 83% is due to production reallocation, the margin the Norwegian experiment identifies. Most of this arrives through production moving towards inherently more productive countries, and the rest through reinforcing productive clusters. Neither the planned expansion of transmission nor a US-equivalent level of interconnection reaches half of the union's gain. Moving renewable capacity toward the best natural resources can even lower TFP on the segmented grid, because the cheap power strands behind transmission bottlenecks.

The gain from integration scales with the electricity share of costs and demand-weighted price dispersion, which are both set to grow over the next fifteen years. Industrial electrification raises the electricity share of costs and widens the differences in exposure across sectors, while the renewable build-out expands low-cost electricity supply differently across countries, increasing price dispersion. Building out renewables without integration pulls production towards countries that build most, and manufacturing productivity falls by 7 bp. Integration reverses the loss by sharing the low-cost capacity. With projected electrification, the union delivers 35 bp of manufacturing productivity in the 2040 economy, and 86 bp within the four most electricity-intensive divisions. Redirecting the planned integration budget towards bottlenecks around Italy captures

¹The closest spatial-misallocation estimates sit on the same order of magnitude: harmonising US state taxes, a wedge several times larger as a share of costs, raises welfare by 0.6% (Fajgelbaum et al., 2019), and the housing-constraint gains of Hsieh and Moretti (2019) fall to basis points once corrected (Greaney, 2026).

²Economy-wide, across sectors most of which buy almost no electricity, the union is worth 6 basis points, of which 2 bp is the fuel saving of a merged merit order. The within-sector figures are pure location channels, since the group normalisation excludes the residual.

85% of the full union manufacturing gain.

Two further benefits accrue from integration. The electricity union insures against energy crises: at 2021–2023 fuel prices, the economy-wide gain of a full union increases by about 60% relative to the calm pre-crisis baseline. Integration also moves fossil dispatch to cheaper plants, which reduces emissions when the carbon allowance price is above 13 euros per tonne of CO₂. At today’s carbon price, a full union halves generation emissions. Hence integration and carbon pricing are complements, with interconnections determining where generation can move, and the carbon price determining whether it moves to cleaner generation technologies.

Literature A recent strand of the misallocation literature identifies specific, observable distortions and aggregates quasi-experimental estimates of their effects into aggregate TFP losses (Bau and Matray, 2023, Sraer and Thesmar, 2023). Within spatial settings, Fajgelbaum et al. (2019) measure the losses from US state tax dispersion, and Hsieh and Moretti (2019) trace aggregate losses to housing-supply constraints that keep workers out of productive cities. This is the closest analogue to the mechanism studied here, with electricity prices in place of housing rents as the binding input cost.³ An electricity price gap is not necessarily a distortion. Gaps can reflect natural advantages and change as production moves, so the gain from removing them cannot be inferred from observed dispersion.

A quantitative literature measures the gains from integrating goods and factor markets. Costinot and Donaldson (2016) measure the gains from integrating US agricultural markets, where heterogeneous factors reallocate across crops and counties according to comparative advantage; Redding (2016) shows that allowing factors to move materially changes the welfare gains from goods-market integration; and Caliendo et al. (2021) quantify the joint integration of goods and labour markets in the 2004 EU enlargement. These papers value integration by what moves once the market opens. Here, electricity can move across borders, reducing the need for production to move towards cheap power.

Most estimates of the value of electricity market integration focus on outcomes inside the electricity sector: where to build generation, how to dispatch it, how to connect it. Newbery, Strbac and Viehoff (2016) evaluate welfare gains from market coupling and congestion rents, Cicala (2022) measure efficiency losses from non-market dispatch, and Gonzales, Ito and Reguant (2023) analyse the returns to transmission investment. Arkolakis and Walsh (2023, 2024) study where clean generation capacity should be built and its local wage effects along dynamic transi-

³See also Gaubert (2018) on how firm sorting and agglomeration jointly shape spatial equilibria; Bartelsman, Haltiwanger and Scarpetta (2013), Restuccia and Rogerson (2017), Brandt, Tombe and Zhu (2013), Moll (2014), Gopinath et al. (2017), Bigio and La’O (2020), Yang, Yang and Sun (2018), Xu and Tan (2021), Choi (2020), Baqaee and Farhi (2020), Tombe and Zhu (2019) on the measurement, sources, and aggregate implications of misallocation across countries and settings, Midrigan and Xu (2014) on the relative importance of incumbent reallocation versus entry and technology adoption, and Asker, Collard-Wexler and De Loecker (2014, 2019) on the interpretation of observed input-price dispersion and supply-curve-based misallocation measurement in commodity markets.

tion paths. I measure how integration changes where electricity-using industries produce, and the productivity of the resulting allocation.

A historical literature documents that energy access has redrawn industrial geography before (Michaels, 2011, Fiszbein et al., 2020, Severnini, 2023, Kline and Moretti, 2014*a*). In today’s electricity system, whether energy advantages reshape industrial geography depends on how far the grid lets cheap power travel.⁴

A reduced-form literature studies how firms respond to electricity and energy prices. Davis et al. (2013) document that electricity prices vary widely even within the US, largely because plants are located in different places. Kahn and Mansur (2013) and Saussay and Sato (2024) estimate firm responses to energy costs, while Colmer et al. (2025) find that plants regulated under the EU Emissions Trading Scheme adjust within-plant rather than relocate. Studies point mainly to adjustment in place through the intensive margin, with less evidence on changes in firm numbers at the extensive margin.⁵ The Norwegian experiment exposes firms to a large, persistent price gap across zones, whose uneven incidence reflects long-standing grid constraints rather than firms’ choices. This contrast suggests why the margin of adjustment differs: firms can absorb temporary or firm-specific cost changes in place, whereas a persistent price gap facing an entire zone changes the number of firms active in electricity-intensive sectors. I identify both margins from the same shock and use the model to trace the location response into aggregate productivity.

Section 2 documents three stylised facts on prices, natural potential, and electricity intensity that motivate the modelling choices. Section 3 develops the model and derives the location equation the empirics estimate. Section 4 presents the Norwegian experiment and the cross-country evidence. Section 5 calibrates the multi-country model and confronts it with untargeted moments, and Section 6 prices integration in today’s system and in the 2040 one. Section 7 takes up three further questions: the distance to US-level integration, insurance against energy crises, and emissions. Section 8 concludes.

2 Motivating stylised facts

Three empirical patterns motivate the modelling choices in Section 3 and discipline the analysis throughout the paper. The first two characterise electricity prices: their persistent dispersion

⁴See also Gutberlet (2014), Crafts and Wolf (2014), Brey (2025), Glaeser, Kerr and Kerr (2015), Ellison and Glaeser (1999), Ellison, Glaeser and Kerr (2010), Klein and Crafts (2012), Combes, Duranton and Gobillon (2008), Greenstone, Hornbeck and Moretti (2010), Kline and Moretti (2014*b*) on the long-run interaction between energy, agglomeration, and industrial location, and on the empirical identification of agglomeration externalities; and Caliendo, Dvorkin and Parro (2019), Fajgelbaum and Redding (2022), Head and Mayer (2014) as methodological precedents for multi-region spatial general equilibrium counterfactuals applied to policy shocks.

⁵See also Davis et al. (2008), Linn (2008), Gray, Linn and Morgenstern (2019), Marin and Vona (2021), Wolverton, Shadbegian and Gray (2022), Allcott, Collard-Wexler and O’Connell (2016), Manderson and Kneller (2020), Panhans, Lavric and Hanley (2017), Tan et al. (2025), Loiacono, Rizzo and Stagnaro (2025), Hampole et al. (2025), Shapiro (2025) on firm-level adjustment to electricity prices, reliability, and energy-policy shocks across the US and European manufacturing sectors.


across segmented national markets and the relative contributions of natural endowments and policy choices to that dispersion. The third characterises the firm side: heterogeneity in electricity intensity across sectors and economies. Together these facts justify the model’s main ingredients — segmented national electricity markets, a distinction between natural and policy-induced sources of comparative advantage, and sector-specific electricity intensity. The remaining ingredient, how strongly production reallocates when electricity prices differ, is identified causally in Section 4.

2.1 Fact 1: Persistent dispersion across segmented markets


Wholesale electricity prices differ by several multiples across European bidding zones, and these differences persist year after year rather than reverting after temporary shocks. Figure 1 documents the cross-sectional and time-series dimensions. Panel (a) maps the 2024 dispersion across bidding zones. The cheapest zones (Nordic hydro areas) averaged roughly one-third the price of the most expensive (Italy, South-East Europe, Poland). Panel (b) plots annual average prices by country. The ordering is persistent across two decades: Italy the most expensive major market throughout, the Nordics and Iberia the cheapest. The 2021–2023 energy crisis amplified differences due to different reliance on fossil fuels across countries. Together, these patterns motivate treating electricity markets as segmented, with persistent price gaps.

Figure 1: Persistent wholesale electricity price dispersion across Europe

(a) Average wholesale price by bidding zone, 2024



(b) Average wholesale prices by country, 2000–2025



Notes: Day-ahead wholesale prices from Bloomberg in EUR/MWh. Panel (a): annual average for 2024, aggregated at the bidding-zone level. Panel (b): annual averages by country, aggregated across bidding zones within country.


This dispersion is not undone by trade. Cross-border electricity trade averages roughly 10%

of total EU electricity consumption, far below the level that would arbitrage price differences away. Where interconnections are small relative to demand, price gaps are larger.⁶ Even where interconnectors are sizeable for the smaller market, they are often small relative to demand on the larger side, leaving room for sizeable price gaps across well-connected borders. These patterns motivate modelling electricity prices as clearing nationally with cross-border trade as a capacity-constrained margin.

2.2 Fact 2: Generation mix and renewable potential

Not all price dispersion has the same source. Countries differ in natural advantages (solar irradiance, wind resources, hydro potential), but their electricity systems also reflect which resources they have developed and whether they rely on nuclear or fossil generation. Figure 2 shows both the electricity countries actually generate and how renewable generation relates to natural potential. This motivates distinguishing natural advantage from generation choices in the model, without yet assigning observed price gaps to either.

Figure 2: Generation mix and renewable potential, 2015–2019



Notes: Panel (a) shows the generation mix by country and generation source in TWh, for the 23 countries in the quantification sample and Norway. Annual data from Ember, averaged over 2015–2019. Coal is hard coal plus lignite. Other comprises other fossil fuels (including Estonian oil shale), bioenergy and other renewables. Countries are ordered by total generation. Panel (b) plots total technical renewable potential relative to electricity demand against realised renewable generation relative to electricity demand by country. The x-axis is on a logarithmic scale; tick labels are factors of annual demand, so a value of 10 means technical renewable potential equal to ten times annual electricity demand. Technical potential is time-invariant and equals the sum of solar PV, wind, and hydro potential. Renewable generation and electricity demand are averages over 2015–2019. Appendix figure A.2 shows total renewable potential by country and source.

Panel (a) shows average annual generation by source and country in TWh. It shows both the

⁶Appendix figure A.1 compares annual interconnection capacity on selected borders with electricity consumption on each side of the border, and then reports the corresponding annual average wholesale price gap.


mix each country uses and the scale of its contribution to European supply. In the model, these differences affect not only national price levels but also how prices respond when demand rises.

Panel (b) plots realised renewable generation against technical potential, both scaled by electricity demand. The two are only weakly related: countries with similar measured potential generate very different amounts of renewable electricity. Potential alone therefore cannot stand in for the electricity system firms face. How much of observed price dispersion reflects natural resources, generation choices, or the limited integration documented in Fact 1 remains a quantitative question, answered in Section 6.2..

2.3 Fact 3: Heterogeneity in electricity cost shares across sectors and countries

Sectors differ widely in the share of electricity in costs, and country economies inherit this variation through their industrial composition. Figure 3 shows electricity cost shares at the NACE 2-digit level, averaged across European countries over 2010–2011. The measure is electricity purchases divided by gross output, also used as exposure variable in the empirical specifications and the quantitative model. The cross-sector gradient is large, spanning roughly 0.3% to 5.6% of gross output, and is the basis for the identification strategy in Section 4. Sectoral differences and varying sectoral composition imply a range of roughly 0.6% to 1.3% of gross output at the country level.⁷

Figure 3: Electricity shares across sectors, 2010-2011



Notes: 2010–2011 EU average electricity purchases divided by gross output by industry, constructed from Exiobase data. Data exclude electricity supply (NACE D35).

⁷Appendix figure A.3 aggregates sector cost shares using each country's gross-output composition, showing predicted country-level electricity exposure.

Together, the three facts set up the following analysis. Electricity prices differ persistently across segmented markets, renewable generation does not simply track natural resources, and sectors differ substantially in their exposure to price gaps. Large electricity bills can affect whether production continues in a location. For example, Hydro stopped primary aluminium production at its Slovalco smelter in Slovakia after it could not secure a competitive electricity contract (Norsk Hydro, 2022).⁸ Selected crisis cases do not tell us how firms generally respond to persistent price differences. Section 4 estimates the response across firms in sectors with different electricity shares.

3 Theory: production location choices with national electricity prices

Pricing a single market for electricity requires a framework in which prices are equilibrium outcomes rather than fixed wedges. The model combines firms choosing production locations across countries, local labour markets that clear within each country, and national electricity markets whose prices are determined endogenously by merit-order dispatch over predetermined generation fleets. In equilibrium, agglomeration concentrates firms, while wage and electricity congestion disperse them, with the generation mix governing the relative strength of the electricity congestion force.

The model delivers three key objects that feed into the productivity effects of market integration: a location multiplier determining how sensitive firm location is to cost differences, an electricity supply elasticity that sets how much electricity prices respond to increased demand in each country based on the generation mix, and a general equilibrium dampening combining them with the rest of the economy.

3.1 Setup

Segmented factor markets. Institutional features of European markets motivate the baseline assumptions. The economy has J countries indexed j and S sectors indexed by s . Electricity markets are segmented, with a national wholesale price p_j^E , consistent with cross-border trade accounting for only about 10% of consumption.⁹ Countries differ in their mix of electricity generation capital across renewables, nuclear, and fossil (K_j^R, K_j^N, K_j^F). Generation capital is predetermined, reflecting commissioning and connection times of more than five years, and asset lifetimes up to 80 years, longer than the horizon over which production adjusts location (Gumber, Zana and Steffen, 2024, International Energy Agency and Nuclear Energy Agency, 2020).

⁸Similarly, BASF curtailed and then structurally closed ammonia capacity at Ludwigshafen (DE), redirecting supply through its global network (BASF, 2021, 2023). Yara curtailed European ammonia production in 2021–2022, sourcing from plants outside Europe (Yara International, 2021, 2022).

⁹See Fact 1 (Section 2.1) for persistent price differences and for limited cross-border electricity trade. This assumption is relaxed with the introduction of cross-border electricity transmission capacity to allow for constrained trade in Section 6.

Each country has a fixed labour endowment $\bar{L}_j$. Labour is immobile, reflecting the fact that only about 3–4% of EU citizens reside in another member state (Eurostat, 2026a, European Commission, 2026).¹⁰ National wages w_j adjust to clear local markets. Both factor prices are therefore endogenous national outcomes rather than exogenous cost wedges, which is the central departure from location models such as Fajgelbaum et al. (2019).

Integrated product market. Within each sector, a continuum of firms produces differentiated varieties under monopolistic competition. All firms sell into a common integrated sectoral market with CES demand elasticity $\sigma > 1$ and a common sectoral demand shifter B_s . Free movement of goods within the single market makes a producer’s customer base largely independent of where in the EU it produces, so location is a pure cost decision, isolated from the market-access forces that drive standard economic geography models.¹¹ The total mass of sector- s firms M_s is taken as given, with the focus on the share locating in country j , $m_{js} \equiv M_{js}/M_s$.

3.2 The firm problem

Technology, profits, and factor demands. A firm f in sector s in country j produces with labour, electricity, and a residual input under Cobb–Douglas technology, $y_{fjs} = \phi_{fjs} L_{fjs}^{\alpha_s} E_{fjs}^{\beta_s} X_{fjs}^{1-\alpha_s-\beta_s}$, and productivity ϕ_{fjs} , specified below. Constant returns fix cost shares at the sector level, so electricity intensity is a characteristic of the sector rather than a firm-level choice (figure 3).¹² The residual input collects capital, materials, and other non-labour inputs, bought at a price p^X common to all countries, so labour and electricity are the only inputs whose price varies with location.¹³

Cost minimisation yields marginal cost proportional to $w_j^{\alpha_s} (p_j^E)^{\beta_s} (p^X)^{1-\alpha_s-\beta_s} / \phi_{fjs}$.¹⁴ Under monopolistic competition with CES demand, firms charge a constant markup $\sigma/(\sigma - 1)$ over marginal cost, and operating profits and factor demands take the form

$$\pi_{fjs} = B_s \phi_{fjs}^{\sigma-1} \underbrace{w_j^{-\alpha_s(\sigma-1)} (p_j^E)^{-\beta_s(\sigma-1)}}_{\text{location attractiveness}}, \quad \ell_{fjs} = \frac{\alpha_s(\sigma-1)}{w_j} \pi_{fjs}, \quad e_{fjs} = \frac{\beta_s(\sigma-1)}{p_j^E} \pi_{fjs}. \quad (1)$$

¹⁰This number is especially low compared with the US, where roughly 30% of working-age Americans live outside their state of birth (Arpaia et al., 2016).

¹¹ B_s embeds the sectoral price index P_s , expenditure Y_s , and the common residual-input price p^X . It cancels from location comparisons and is tracked explicitly only in the quantitative counterfactuals. Appendix B records its expression.

¹²The production elasticities (α_s, β_s) are not the expenditure shares recorded in the data: with a constant markup, input expenditure is a fraction $(\sigma - 1)/\sigma$ of revenue, so a share measured over gross output is $(\sigma - 1)/\sigma$ times the corresponding elasticity. Section 4 states the mapping where it is used, in the model-consistent specification.

¹³The residual input is how the model treats inputs whose prices do not vary across locations. It affects the level of costs and profits, but its common price drops out of every location comparison and is absorbed into B_s . Its share is a residual, so the two location-relevant shares satisfy $\alpha_s + \beta_s < 1$. Appendix B records the closure holding the residual-input block fixed, so the model is not a general equilibrium of capital and materials markets.

¹⁴The omitted factor of proportionality and the common residual-input price are sector constants shared by all countries and are absorbed into B_s . Appendix B derives marginal cost, profits, and the composition of B_s in full.

Profits decompose into a sectoral scale term B_s , a country–sector productivity term ϕ_{js} , and a location-attractiveness term decreasing in wages and electricity prices, with sector-specific elasticities $\alpha_s(\sigma - 1)$ and $\beta_s(\sigma - 1)$. Electricity-intensive (high- β_s) sectors are therefore more sensitive to electricity prices: the same country-level price change generates differential location responses across sectors, which is the source of identification in the empirical analysis. Labour and electricity demand are both proportional to operating profits, so π_{fjs} is a sufficient statistic for firm scale and feeds directly into the market-clearing conditions below.¹⁵

Productivity and the location choice. I assume firm productivity has a systematic component that varies across locations, and the location choice is additionally shifted by an idiosyncratic firm–location match. Operating productivity in (j, s) is $\ln \phi_{js} = \bar{\Phi}_{js} + \gamma \ln M_{js}$, and from the profit function a firm’s normalised log profit in j is $\frac{1}{\sigma-1} \ln \pi_{fjs} = \bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E + \text{const}_s$.¹⁶ A firm ranks locations by maximising the location payoff in terms of profits and location match,

$$\max_j \underbrace{\bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E}_{\equiv v_{js}, \text{ location-varying log profit}} + a_{fjs}, \quad a_{fjs} \stackrel{iid}{\sim} \text{Type I EV}(0, \psi), \quad (2)$$

where $\bar{\Phi}_{js}$ is a country–sector base productivity capturing fundamentals unrelated to electricity (institutions, capital intensity, technology, workforce composition), $\gamma \ln M_{js}$ is an agglomeration externality with elasticity $\gamma \geq 0$ capturing thick-labour-market, input-sharing, and knowledge-spillover channels (Ellison, Glaeser and Kerr, 2010, Combes, Duranton and Gobillon, 2008, Greenstone, Hornbeck and Moretti, 2010). a_{fjs} is an idiosyncratic firm–location match (local ties, established relations, or other factors that tie a particular firm to a particular place) following a Gumbel distribution, which shifts where a firm locates but not its operating productivity once there.

Every firm in (j, s) has common operating productivity that enters profits and factor demands above.¹⁷ The dispersion parameter $\psi > 0$ governs how strongly firms respond to systematic cost differences: high ψ means idiosyncratic match quality dominates, $\psi \rightarrow 0$ means firms sort

¹⁵When p_j^E rises, firm-level electricity demand falls through two channels: direct input substitution and a scale contraction working through π_{fjs} . The combined elasticity $-[1 + \beta_s(\sigma - 1)]$ reappears in the electricity market-clearing condition and contributes to the GE dampening below.

¹⁶Normalising by $(\sigma - 1)$ puts the match draw in the same units as the location payoff and keeps σ out of the location ranking, while leaving it explicit in profits, factor demands, and market clearing. The sector constant collects B_s , the common residual-input price, and the Cobb–Douglas cost constant, all of which are identical across countries and drop from the ranking.

¹⁷The match term enters the discrete location problem exactly as a Gumbel productivity draw and is treated as a non-pecuniary firm–location payoff, the firm-side counterpart of the idiosyncratic location preferences in Suárez Serrato and Zidar (2016). Because it does not load on realised cell output, country–sector productivity equals log operating productivity $\bar{\Phi}_{js} + \gamma \ln M_{js}$, and the firm-selection margin that arises when establishments operate with their own draw — the “firm-location” term in Suárez Serrato and Zidar (2016, eq. 8–9) and the representative-productivity term $\bar{z} = z^0 M^{-1/\varepsilon_F}$ in Fajgelbaum et al. (2019) — is absent by construction; with positive agglomeration this selection term would otherwise net against γ in average productivity.

perfectly on fundamentals.

The Gumbel assumption is what makes the location problem tractable: it turns a discrete choice over countries into a closed-form, log-linear relationship between firm counts and local costs. Given the draw and the profit function, location shares satisfy a logit expression:

Lemma 1 (Logit location shares).

$$m_{js} \equiv \frac{M_{js}}{M_s} = \frac{\exp(v_{js}/\psi)}{\sum_k \exp(v_{ks}/\psi)}, \quad v_{js} \equiv \bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E \quad (3)$$

The share of firms in sector s locating in country j is increasing in base productivity $\bar{\Phi}_{js}$ and in agglomeration $\gamma \ln M_{js}$, and decreasing in wages and electricity prices, with greater sensitivity in electricity-intensive sectors.

3.3 Partial equilibrium

The structure so far is already enough to answer the paper's question in its simplest form: if a country's electricity price falls by one percent, how much production moves there, before anything else adjusts? Logit location shares and the profit function deliver the answer in one line, and it depends on a two-parameter composite that the empirical work is built to recover.

Totally differentiating the log of the location share in Lemma 1, with wages and electricity prices treated as exogenous, gives the direct response of the firm mass to cost changes. Substituting the agglomeration feedback, whereby entry raises operating productivity by $\gamma d \ln M_{js}$, yields the partial equilibrium (PE) location equation.¹⁸

Proposition 1 (Location equation in partial equilibrium).

$$d \ln M_{js} = \frac{1}{\psi - \gamma} \left[-\alpha_s d \ln w_j - \beta_s d \ln p_j^E \right] + \text{const}_s \quad (4)$$

The location multiplier $1/(\psi - \gamma)$ captures the convergent sum of the agglomeration feedback loop: each firm's location response is amplified by the further agglomeration it induces. The constant const_s is common to all countries within the sector and is pinned by the fixed sector mass, $\sum_j m_{js} d \ln M_{js} = 0$.

The location multiplier $1/(\psi - \gamma)$ is the first of the model's key objects. It nets two forces that decide how much production reallocates in response to cost differences. Agglomeration γ pulls firms into the same location with the cheapest cost combination given their input mix. A positive γ leads to amplification in response to a cost change. A fall in p_j^E attracts firms ($d \ln M_{js}$ rises),

¹⁸Throughout the analytical results, $d \ln x$ denotes the log-change in variable x between two nearby stationary states. Sector-specific terms common across countries drop out in country differences, while the base productivity $\bar{\Phi}_{js}$ is a fixed fundamental that shapes the level of the firm distribution but drops out of every comparative static.

which raises productivity through γ , which attracts more firms, and so on.¹⁹ Match dispersion ψ instead sorts firms according to their individual draws. Because the total sector mass is fixed, only dispersion in prices reshapes production, and a price change common to all countries leaves every location comparison unchanged.

Section 4 estimates the location multiplier from the cross-sector gradient in electricity intensity. Sectors differ in β_s so a price change common to a location moves them differentially, and the slope of that differential identifies $1/(\psi - \gamma)$.²⁰ Quantification in Section 5 splits it into dispersion ψ given external estimates of agglomeration γ .

3.4 Electricity supply

European wholesale electricity markets are segmented into mostly national bidding zones, with persistent cross-country price disparities documented in Section 2.²¹ The supply-side of the model formalises the mechanism behind these price differences. Each country j 's electricity system combines three generation technologies with different marginal costs. The installed generation capacity by technology relative to demand then determines which technology is marginal and hence sets the market price through merit-order pricing.

Low-cost generation. Renewables (composed of solar, wind, and hydro) generate electricity at zero marginal cost, $c^R = 0$.²² Their generation output is $\theta_j^R K_j^R$, which is determined by installed generation capacity K_j^R , and additionally varies with the fixed natural productivity parameter θ_j^R , capturing geographic potential by country. Nuclear generation is fixed and fully determined by installed capacity K_j^N . It has a low constant marginal cost $c^N > 0$.

Nuclear and renewable electricity generation together form low-cost generation, both supplying electricity exogenously at a low and fixed marginal cost.

$$\bar{E}_j^0 \equiv \theta_j^R K_j^R + K_j^N. \quad (5)$$

Total effective low-cost capacity $\bar{E}_j^0$ can hence be decomposed into a natural advantage component (θ_j^R), reflecting geography, and a policy component, reflecting policy decisions on the generation mix (K_j^N, K_j^R).

¹⁹The condition $\psi > \gamma$ rules out a local tipping force and ensures that the loop converges. See Appendix B for the single-sector derivation. In general equilibrium the same condition suffices: wage and electricity congestion are stabilising, and the dampening factor lies between zero and one automatically (Appendix B.8).

²⁰Because the shares recorded in the data are the production elasticities scaled by the markup, the estimated interaction coefficient is $\sigma/(\sigma - 1)$ times the multiplier, and Section 4 states the mapping with the specification.

²¹Appendix A provides institutional detail.

²²The engineering marginal cost is zero, but the wholesale price is pinned by a positive opportunity price $p_j > 0$ (the value of selling the surplus elsewhere).

Fossil generation. Fossil capacity K_j^F generates at increasing marginal cost, with utilisation of the capacity adapting to electricity demand. The fossil fleet in each country consists of a continuum of plants with heterogeneous efficiencies at which they convert fuel into electricity, an electricity-generation counterpart of heterogeneous firm productivities in Melitz (2003). They are dispatched from most to least efficient in order to meet electricity demand, which leads to an upward-sloping fossil segment of the supply curve.

Efficiency is tracked by its inverse, the heat rate: the fuel a plant needs per unit of electricity produced, and hence the quantity that translates directly into marginal cost. I assume heat rates follow a common Pareto distribution with shape $\zeta > 1$ and lower bound $h_{\min}$, the heat rate of the frontier plant, so that most of the fleet sits close to the frontier and a thin tail is markedly less efficient. All countries draw from the same distribution, and cross-country variation arises through the number of active plants and the utilisation cutoff.²³

The plant at position k in the efficiency-ranked fleet has heat rate $h(k) = h_{\min}(1 - k)^{-1/\zeta}$, the quantile function of the Pareto distribution. The position k of the marginal active plant in the fleet gives the utilisation rate of the fleet, $u_j \equiv (E_j^D - \bar{E}_j^0)/K_j^F \in (0, 1)$, the fraction of fossil capacity currently dispatched. The marginal cost of fuel generation therefore is given by

$$c_j^F = c_0^F + p^F h_{\min}(1 - u_j)^{-1/\zeta}, \quad (6)$$

where p^F is the fuel price and $c_0^F > 0$ is the non-fuel component of fossil marginal cost (variable operations and maintenance).²⁴ This marginal cost is increasing and convex in utilisation u_j : each additional unit of fossil generation requires activating a less fuel-efficient plant, so the cost increase accelerates as the fleet approaches full capacity.

Merit-order pricing. Pricing follows the merit order principle, where the wholesale price equals the marginal cost of the last generation activated to satisfy demand. The resulting supply schedule is piecewise, as illustrated in figure 4. Renewables ($c^R = 0$ marginal cost) are dispatched first, nuclear ($c^N > 0$) next, fossil ($c^F \gg 0$) last.


Countries with high $\bar{E}_j^0$ relative to demand have lower electricity prices, reflecting resource-based advantages through renewables potential (θ_j^R) and generation policy (K_j^N, K_j^R). A country with abundant hydro resources such as Norway can operate at very low prices, France also has low average prices due to a large share of nuclear generation capacity, while Germany, where coal and gas typically set the marginal price, has higher prices dependent on fuel costs.²⁵

²³The fossil fleet technology rests on five assumptions discussed in Appendix B: common Pareto heat rates, equal-sized plants, linear plant technology, binary dispatch, and a common fuel price p^F .

²⁴Appendix B derives the supply curve and discusses details. The non-fuel component c_0^F responds neither to fuel prices nor to utilisation, so it results in fuel-price pass-through being less than one-for-one, and makes the price response to demand scale with the fuel share of the price rather than with the price itself.

²⁵This is the mechanism behind the bidding-zone divergence used for identification in Section 4.

Figure 4: Stylised electricity supply with merit order pricing



Supply elasticity. The second key object is the supply elasticity $\varepsilon_j \equiv \partial \ln p_j^E / \partial \ln E_j^D$: the percentage electricity price increase from a 1% increase in total demand. Lemma 2 pins it down in each regime.²⁶

Lemma 2 (Supply elasticity by marginal source). *Given generation capital (K_j^R, K_j^N, K_j^F) and electricity demand E_j^D , the local supply elasticity $\varepsilon_j \equiv \partial \ln p_j^E / \partial \ln E_j^D$ is:*

$$\varepsilon_j = \begin{cases} 0 & \text{if } E_j^D \leq \theta_j^R K_j^R + K_j^N \equiv \bar{E}_j^0 \quad (\text{low-cost marginal}) \\ \frac{\chi_j}{\zeta} \frac{E_j^D}{K_j^F (1-u_j)} > 0 & \text{if } E_j^D > \bar{E}_j^0 \quad (\text{fossils marginal}) \end{cases} \quad (7)$$

where $\chi_j \equiv \frac{p_j^F h_{\min}(1-u_j)^{-1/\zeta}}{p_j^E}$ is the fuel cost share of the electricity price, the fraction of the marginal plant's cost that is fuel rather than the fixed non-fuel component c_0^F .

When renewables or nuclear are marginal, the supply elasticity ε_j is zero, since the price is a constant. When fossil generation is marginal, the elasticity measures how tight the fossil margin is.²⁷ This is determined by three factors. First, the fuel cost share χ_j scales the elasticity, because only the fuel component of the marginal plant's cost responds to utilisation while the non-fuel component c_0^F is constant. Second, the Pareto shape ζ governs how fast efficiency deteriorates down the fleet, so a thin tail of inefficient plants (high ζ) makes the curve flat.

Third, spare fossil capacity $K_j^F (1 - u_j)$ relative to total demand E_j^D is the room the system has to absorb a demand increase. That ratio can be small and hence the elasticity ε_j large in different

²⁶Appendix B records the formal derivation, the key properties of the micro-founded supply curve, and the Pareto–Melitz analogy.

²⁷The elasticity equals the fuel cost share times total electricity demand divided by spare fossil capacity, scaled by fleet efficiency dispersion $1/\zeta$.

cases. Either because a fossil-dependent country runs a fossil fleet at high utilisation, or because a country with abundant low-cost generation is operating in its fossil segment. In the first case, the fleet is operating with marginally inefficient plants that have steeply increasing costs. In the second, the fossil segment is small, and an increase in demand leads to sharp price increases once that segment is in use. Both appear in the calibrated sample, such that a large low-cost share of generation does not automatically translate into a lower supply elasticity ε_j .²⁸

Generation mix as comparative advantage. The generation mix sets the level of a country's electricity price and also the slope of the supply curve. The level decides how attractive a location is to electricity-intensive production, while the slope decides how fast that attraction erodes as production arrives. Only a country with both a low price and a deep fossil buffer can host a large cluster of electricity-intensive industry without bidding its own advantage away, which constitutes a comparative advantage in hosting agglomeration.

A low price and low elasticity rarely go together. Adding low-cost capacity $\bar{E}_j^0$ within a country lowers its fossil utilisation u_j and so lowers both its price and its supply elasticity. Across countries, systems with abundant low-cost generation often hold only a thin fossil fleet, so their curve is steep while their price is low. France and Sweden have among the cheapest power in the sample and among the steepest supply curves, while fossil-heavy Germany and Poland have a smaller electricity price response to new demand (Appendix F). A country's price therefore says little about how much industry it can absorb, which is one reason that moving power across borders and moving production to power are not equivalent policies.

All ingredients into the key supply elasticity ε_j are observable, so the cross-country variation is measured rather than estimated.²⁹ Section 5 derives the primitives from data on generation, resource maps, demand, and the fossil cost stack.

3.5 General equilibrium

Integration is not a partial-equilibrium event: when production moves, wages rise where it arrives, and the electricity price it moved for rises with it. Closing the model determines how much of a price advantage survives being used, after accounting for equilibrium responses.

²⁸The high supply elasticity for low-cost systems operating in their fossil segment can be unintuitive, and is a consequence of the assumption that the fossil fleet in each country is drawn from the same Pareto distribution, and not truncated depending on its size. In the quantitative model, I test both versions of the plant sampling, and the currently used fixed distribution yields a better fit of the price geography across countries, which is why it is chosen as the baseline here.

²⁹Low-cost capacity $\bar{E}_j^0$ combines installed capacity by technology, observed in generation data, with natural potential θ_j^R approximated from resource maps. Fossil utilisation u_j is an equilibrium outcome given capacities and demand. The Pareto shape ζ is calibrated to the observed cost stack in Section 5 and tested in the sensitivity analysis.

Market clearing. Labour market clearing requires

$$\sum_s M_{js} \ell_{js}(w_j, p_j^E) = \bar{L}_j \quad (8)$$

in each country, and electricity market clearing requires

$$E_j^D = \sum_s M_{js} \beta_s (\sigma - 1) \bar{\pi}_{js} / p_j^E \quad (9)$$

with the electricity price pinned down by the supply curve (6).

Definition 1 (Stationary equilibrium). *A stationary equilibrium is a set of firm masses, wages, and electricity prices $\{M_{js}, w_j, p_j^E\}_{j,s}$ satisfying: (i) logit location shares (Lemma 1) with $v_{js} = \bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E$; (ii) labour market clearing (8) for each j ; (iii) electricity market clearing (9) for each j with the supply curve (6); and (iv) $\sum_j M_{js} = M_s$ for each s .*

The three-equation system. Around a fossil-marginal equilibrium (Lemma 2) the system log-linearises. Each country has one firm mass per sector, one wage and one electricity price, determined by three blocks: firms choose locations cell by cell, while the wage and the electricity price each clear a national market.

Lemma 3 (The equilibrium system).

$$\text{(Location, each } j, s) \quad d \ln M_{js} = \frac{1}{\psi - \gamma} \left[-\alpha_s d \ln w_j - \beta_s d \ln p_j^E \right] + \text{const}_s \quad (\text{Prop. 1}) \quad (10)$$

$$\text{(Price, each } j) \quad d \ln p_j^E = \frac{\varepsilon_j \left[(1 + (\sigma - 1)\gamma) d \ln \bar{M}_j^E - (\sigma - 1) \bar{\alpha}_j^E d \ln w_j \right] + d \ln p_j^E|_{ED}}{1 + \varepsilon_j \left[1 + (\sigma - 1) \bar{\beta}_j^E \right]} \quad (11)$$

$$\text{(Wage, each } j) \quad d \ln w_j = \frac{(1 + (\sigma - 1)\gamma) d \ln \bar{M}_j^L - (\sigma - 1) \bar{\beta}_j^L d \ln p_j^E}{1 + (\sigma - 1) \bar{\alpha}_j^L} \quad (12)$$

where $d \ln p_j^E|_{ED}$ is the direct effect of a supply-side change on the wholesale price at a given level of demand.^{30,31} Variables $\bar{x}_j^Y \equiv \sum_s \omega_{js}^Y x_s$ represent averages of x weighted by the factor Y .³²

³⁰Generation capital enters only through this impulse and through ε_j : the counterfactuals of Section 6 move low-cost capacity $\bar{E}_j^0$ and the prices that set the fossil segment, while fossil capacity K_j^F is held fixed throughout.

³¹The location equation carries the sector constant of Proposition 1, which is what enforces the fixed sector mass. It does not drop from the weighted averages unless the country is small within each of its sectors, and the reductions below set it aside on that basis; the quantitative model carries it in full, and Section 5.5 measures its contribution as second order.

³²Write $\omega_{js}^L = L_{js}/\bar{L}_j$ and $\omega_{js}^E = E_{js}/E_j^D$ for the shares of country j 's labour and electricity used by sector s , and $\bar{x}_j^L \equiv \sum_s \omega_{js}^L x_s$ and $\bar{x}_j^E \equiv \sum_s \omega_{js}^E x_s$ for the corresponding averages of the cost shares $x \in \{\alpha, \beta\}$, with $d \ln \bar{M}_j^L$ and $d \ln \bar{M}_j^E$ the weighted averages of the sectoral mass changes.

Each equation carries one of the three forces that balance in equilibrium. The location equation (10) is Proposition 1 and acts as the concentration force. Agglomeration pulls firms into clusters through $\gamma \ln M_{js}$, which unchecked would concentrate each sector in a single country. It holds cell by cell, so the two factor prices are common to the country while the exposure to them is not, as a sector responds through its own α_s and β_s .

Each clearing condition weights the country's sectors by the factor it clears. The price equation (11) follows from electricity market clearing (9) against the supply curve and carries electricity congestion. Firm entry pushes demand up the merit-order supply curve, raising the price with a strength governed by the country-specific supply elasticity ε_j , and what presses on that curve is the electricity-weighted average intensity $\bar{\beta}_j^E$, so the country's price is shaped by the sectors that use the most electricity rather than by the average sector.³³ The wage equation (12) follows from labour market clearing (8) and carries the standard wage congestion force, weighted across sectors based on where labour is employed. Entry raises wage w_j and compresses profits. Both congestion equations are self-dampening through their denominators, since higher factor prices cut profits and so cut the demand that raised them in the first place.³⁴

With a single sector, the factor weightings coincide, $\bar{x}_j^L = \bar{x}_j^E = x$, and the system solves in closed form for a single dampening composite Ω_j (Appendix B.5, Proposition 4). With many sectors it does not, because the wage that clears the labour market is set by sectors according to employment, while the price that clears the electricity market is set by sectors according to electricity use.

The GE location multiplier. Every policy evaluated in the counterfactuals (Section 6) reaches firms as a shift in the national electricity supply curve: the replacement of a country's own price by a union-wide one, a build-out of low-cost capacity, or a move in world fuel prices. Each reaches the firm distribution through the direct price impulse it delivers at a given level of demand ($d \ln p_j^E|_{ED}$), damped by the concentration and congestion forces. The equilibrium response only differs by impulse size and incidence, such that it can be solved generally for any local supply-side shift in Proposition 2.³⁵ Appendix B records details and derivations.

³³For the analytical results the supply elasticity ε_j is evaluated at the initial equilibrium, while the quantitative model updates it along the counterfactual path, reinforcing the analytical findings because adding low-cost capacity simultaneously lowers both the price p_j^E and elasticity ε_j .

³⁴For instance, firm entry raises labour demand directly through the mass of firms in the country and indirectly through agglomeration-induced incumbent scale, with the denominator $1 + (\sigma - 1)\bar{\alpha}_j^L$ reflecting the self-dampening of higher wages through reduced profits and labour demand.

³⁵The system reduces in two steps. Labour market clearing gives the country's wage response $\lambda_j \equiv d \ln w_j / d \ln p_j^E < 0$, the same for every sector. Electricity market clearing, evaluated at that wage response, gives the country's electricity demand elasticity $\Theta_j \equiv -d \ln E_j^D / d \ln p_j^E$, and with the supply curve it delivers the equilibrium price change $d \ln p_j^E = d \ln p_j^E|_{ED} / (1 + \varepsilon_j \Theta_j)$. Substituting both into the location equation gives the proposition. Appendix B.8 derives λ_j and Θ_j and records the terms set aside.

Proposition 2 (Location equation in general equilibrium). Let $\lambda_j \equiv d \ln w_j / d \ln p_j^E$ and $\Theta_j \equiv -d \ln E_j^D / d \ln p_j^E$ denote the wage and electricity-demand responses of country j , which solve the two clearing conditions of Lemma 3 jointly. The firm mass response in sector s to a supply impulse is

$$d \ln M_{js} = - \underbrace{\frac{1}{\psi - \gamma}}_{\text{location multiplier}} \times \underbrace{[\beta_s + \alpha_s \lambda_j]}_{\text{exposure net of wage}} \times \underbrace{\frac{1}{1 + \varepsilon_j \Theta_j}}_{\text{GE dampening}} \times d \ln p_j^E \Big|_{E^D} + \text{const}_s, \quad (13)$$

where

$$\lambda_j \equiv \frac{d \ln w_j}{d \ln p_j^E} = - \frac{\mathcal{A} \bar{\beta}_j^L}{1 + \mathcal{A} \bar{\alpha}_j^L} < 0, \quad \Theta_j \equiv - \frac{d \ln E_j^D}{d \ln p_j^E} = 1 + \mathcal{A} [\bar{\beta}_j^E + \bar{\alpha}_j^E \lambda_j],$$

$$\mathcal{A} \equiv \underbrace{(\sigma - 1)}_{\text{incumbents scale down}} + \underbrace{\frac{1 + (\sigma - 1)\gamma}{\psi - \gamma}}_{\text{firms leave}} = \frac{1 + (\sigma - 1)\psi}{\psi - \gamma}.$$

λ_j and Θ_j have no closed form since they are solved from the same system they enter. Conditional on them, each sector's response is exact. The single-sector version does have a closed form, derived in Appendix B.8.

The general equilibrium location response is the partial-equilibrium one of Proposition 1 modified in two places. The exposure term replaces a sector's own electricity share with its exposure net of the wage relief the price change brings with it: when electricity becomes expensive, electricity-intensive firms leave, labour demand falls, and the wage drops by $\lambda_j < 0$ for every firm in the country. A sector keeps most of its exposure if electricity is a large share of what it buys, and loses it if labour is. The dampening $1/(1 + \varepsilon_j \Theta_j)$ then scales every sector in the country alike: firms that arrive raise electricity demand, which raises the price along the supply curve at rate ε_j , and Θ_j measures how much demand responds to that price in the first place.³⁶

Both country objects ε_j and Θ_j are built from the same activity multiplier $\mathcal{A}$, which converts a change in unit costs into a change in the production a country carries out, which can be split into two margins.³⁷ Incumbents scale down, with elasticity $\sigma - 1$ under CES demand. Firms leave, at the location multiplier $1/(\psi - \gamma)$, each taking its own output and, through agglomeration, a further $(\sigma - 1)\gamma$ of the output of the firms that stay. Both objects are well behaved under the partial-equilibrium condition alone: given $\psi > \gamma$ the dampening lies between zero and one, because the sectors holding most of a country's electricity demand are necessarily more electricity-intensive

³⁶What the regression of Section 4 identifies is therefore not what the counterfactuals of Section 6 deliver: the regression absorbs precisely the responses that λ_j and $\varepsilon_j \Theta_j$ collect. Section 5.5 measures the resulting gap in the calibrated model.

³⁷The two margins are the elasticity of cell revenue to firm mass at given factor prices, $1 + (\sigma - 1)\gamma$, scaled by how far firms move, $1/(\psi - \gamma)$, plus the scale response of firms already in place, $\sigma - 1$. Together they give $\mathcal{A} = [1 + (\sigma - 1)\psi]/(\psi - \gamma)$, the same composite that scales expenditure shares in the productivity decomposition of Section 3.6.

than those holding most of its employment, so the wage offset cannot outweigh the electricity saving in the aggregate (Appendix B.8). General equilibrium therefore damps the location response rather than amplifying it, in every country.

Because the sectoral factor intensities differ, a national price change sorts sectors rather than moving them together. The response is damped monotonically in the labour-to-electricity cost ratio, and reverses sign for sectors with $\alpha_s/\beta_s > 1/|\lambda_j|$. For sufficiently labour-intensive production the wage relief released by the departing electricity-intensive firms outweighs its own electricity saving, and it moves toward the country whose power just became more expensive. A cheap-power country therefore ends up specialised in whatever its electricity price makes it good at (Appendix B, Proposition 5). The response of a country as a whole is a mixture of sectors moving in opposite directions rather than a scaled-up version of any one of them (Section 5.5).

The joint GE dampening force is the third object the quantification needs, and unlike the first two it is not a primitive but a combination of them. The location multiplier $1/(\psi - \gamma)$ of Proposition 1 scales the activity response inside Θ_j . The cost shares β_s and α_s and the demand elasticity σ set how strongly cost changes pass into profits and factor demands. The supply elasticity ε_j of Lemma 2 converts the resulting demand response into a price. The elasticity ε_j and the country's own sector composition are the only country-specific ingredients.

3.6 Productivity effects of price dispersion

The theory can now answer the main question of the paper qualitatively, by decomposing the aggregate productivity impact of electricity price dispersion into four channels. I measure aggregate productivity as the real value added the economy obtains from given quantities of primary inputs, and it has three parts: the expenditure-weighted average of local log productivity, a second-order adjustment for the dispersion of unit costs, and a deduction for the primary resources used up.³⁸ Local productivity itself has two sources, a fundamental $\bar{\Phi}_{js}$ and the agglomeration externality $\gamma \ln M_{js}$. What reallocation changes is the weights across locations: aggregate productivity rises when expenditure shifts toward more productive locations. Factor supplies are fixed at their endowments, so the only primary input whose use moves between two equilibria is the fuel that fossil plants import to generate the electricity those sectors consume.

Productivity channels. A decomposition of TFP into four channels applies to any change in the allocation of production. The three key objects the model delivers enter through two quantities: how much a location's unit cost changes ($\alpha_s \Delta \ln w_j + \beta_s \Delta \ln p_j^E$), and how much production moves

³⁸This is a production-side technology index, not a welfare measure: it excludes the consumer gain when the same reallocation moves sectoral price indices, including love of variety. Appendix B.10 sets that term aside explicitly, and the quantitative model tracks it separately through P_s .

in response ($\Delta \ln M_{js}$).³⁹ The supply elasticity ε_j governs the electricity part of the first (Lemma 2), the location multiplier ($1/(\psi - \gamma)$) and the general equilibrium dampening ($1/(1 + \varepsilon_j \Theta_j)$) govern the second (Proposition 2), and everything else in the decomposition is a weight or a fundamental disciplined outside the firm response.

Lemma 4 (TFP decomposition). *The change in TFP between two allocations is:*

$$\begin{aligned} \Delta \ln TFP \approx & \underbrace{\Delta \left[\sum_s \mu_s \sum_j \omega_{js} \bar{\Phi}_{js} \right]}_{\text{composition}} + \underbrace{\gamma \Delta \left[\sum_s \mu_s \sum_j \omega_{js} \ln M_{js} \right]}_{\text{agglomeration}} \\ & + \underbrace{\sum_s \mu_s \frac{\sigma - 1}{2} \text{Var}_{\omega_{js}} \left(\alpha_s \Delta \ln w_j + \beta_s \Delta \ln p_j^E \right)}_{\text{allocative efficiency}} - \underbrace{\sum_j s_j^F \Delta \ln E_j^F}_{\text{fossil resource use}}, \end{aligned} \quad (14)$$

where μ_s are sectoral expenditure weights, ω_{js} are expenditure shares, α_s and β_s the labour and electricity cost shares, and $s_j^F = (p_j^E - c_0^E)E_j^F / \mathcal{Y}$ is country j 's fossil fuel bill as a share of aggregate value added $\mathcal{Y}$.⁴⁰

The weights are not a normalisation: by Shephard's lemma the sectoral price index satisfies $\partial \ln P_s / \partial \ln c_{js} = \omega_{js}$, so a location that becomes one percent more productive raises aggregate productivity by exactly its share of sectoral expenditure, with $\mu_s \omega_{js}$ the Domar weights of j . Reallocation matters because those weights move.

The first three channels operate through reallocation of expenditure across locations. Expenditure weights change because firms relocate or because firms within a location produce and sell more when costs fall, and the second operates even when firms cannot move (Appendix B.10). Composition asks whether reallocation shifts production toward the countries that are inherently more productive in a sector ($\bar{\Phi}_{js}$). Agglomeration asks how much productive clusters grow, and whether the expenditure that moves goes to clusters that were already large ($\gamma \ln M_{js}$). Allocative efficiency asks how unevenly the reallocation is spread: for given inputs, aggregate output is higher when marginal costs are less dispersed across locations within a sector, because demand

³⁹The bracket in Proposition 1 is exactly the cost change, so in equilibrium $\alpha_s \Delta \ln w_j + \beta_s \Delta \ln p_j^E = -(\psi - \gamma) \Delta \ln M_{js}$. The location multiplier is the exchange rate between the two quantities.

⁴⁰The two location channels are exact; the resource term holds to first order and the allocative term to second, evaluated at fixed weights. Appendix B.10 derives each object in the form the calibrated model evaluates it, with weights and costs revalued at the counterfactual allocation, and shows that the additional endogenous-weight terms are not sign-definite in general. The weights themselves are parameterised in Section 5.2, together with the fundamentals $\bar{\Phi}_{js}$, which are recovered by model inversion.

shifts toward lower-cost producers.⁴¹ ⁴² Composition and agglomeration are linear in the reallocation while allocative efficiency is quadratic in it, so the first two scale with a sector's exposure to cost differences, electricity intensity β_s , and the third with its square. Appendix B.10 derives the weight responses behind all three.

The fossil resource channel is independent of production locations. Electricity is generated domestically, so payments for it are internal transfers, except for the fuel that fossil plants import. A fall in aggregate fossil generation therefore raises real value added directly as a level term, and how much generation a given policy displaces is fixed by the same merit-order supply curves that set ε_j .

Policy choices on generation and integration yield price differences that act on all four channels. Policy-induced dispersion can redirect concentration toward countries whose low prices reflect policy choices rather than natural cost advantages, and changes which generation plants are dispatched to serve the resulting demand. The change that matters for the question of this paper is the removal of that dispersion through market integration.

Electricity market integration. An electricity union removes all electricity price differences including the distortion they were imposing. Under market segmentation, electricity congestion is a national force, where a country that attracts electricity-intensive production pays through a higher national electricity price depending on its supply elasticity ε_j . Integration yields a common electricity congestion force across countries and therefore removes it from the location decision, such that production sorts on fundamentals and wages alone (Proposition 6, Appendix B). Whether this raises productivity depends on why prices differed.

What still reallocates production is the baseline dispersion itself. The union price cancels from every location comparison, so a country gains production where its own baseline price stood above its sector's average, weighted by where that sector's firms were, and not where it stands relative to the union price (equation (B.15), Appendix B.11). A country whose electricity cost rises can therefore still gain if its competitors' costs rise by more. Wage congestion survives, since national labour markets keep clearing, so part of every move is absorbed by the wage in the country receiving it.

Proposition 3 (TFP gain from an electricity union). *Let M_{js}^{EU} denote the firm distribution under a full electricity union that equalises wholesale prices across countries. Appendix B.11 derives the price that clears*

⁴¹The term is a variance, so it is non-negative for any shock, including one that raises dispersion: with CES demand, expenditure reallocates toward whichever locations become relatively cheaper. Counterfactuals therefore differ in the size of the relative cost changes they induce, not in the sign of this term, and for the union it scales with the baseline dispersion removed.

⁴²The composition term is an ? covariance with countries in place of firms; ? extend it to entry and exit. The object as a whole is the Hsieh and Klenow (2009) accounting in space, with electricity prices as the wedge, as in Hsieh and Moretti (2019) and Fajgelbaum et al. (2019). That misallocation enters at second order while the reallocation channels enter at first is the structure of Baqaee and Farhi (2020) for inefficient economies.

it and proves the result. To first order, the union's aggregate-TFP gain is the decomposition of Lemma 4 read as moments of the reallocation it induces, with $\Lambda \equiv 1 + (\sigma - 1)\psi$ the elasticity of a country-sector's expenditure to its firm mass:⁴³

$$\begin{aligned} \Delta \ln TFP^{EU} \approx & \underbrace{\sum_s \mu_s \Lambda \text{Cov}_{\omega_{js}} \left(\Delta^{EU} \ln M_{js}, \bar{\Phi}_{js} \right)}_{\substack{\text{composition} \\ (\pm)}} + \underbrace{\sum_s \mu_s \gamma \left[\mathbb{E}_{\omega_{js}} \Delta^{EU} \ln M_{js} + \Lambda \text{Cov}_{\omega_{js}} \left(\Delta^{EU} \ln M_{js}, \ln M_{js}^0 \right) \right]}_{\substack{\text{agglomeration} \\ (\pm)}} \\ & + \underbrace{\sum_s \mu_s \frac{(\sigma - 1)(\psi - \gamma)^2}{2} \text{Var}_{\omega_{js}} \left(\Delta^{EU} \ln M_{js} \right)}_{\substack{\text{allocative efficiency} \\ (\geq 0)}} - \underbrace{\sum_j s_j^F \Delta^{EU} \ln E_j^F}_{\substack{\text{fossil resource use} \\ (\geq 0 \text{ if } E^F \text{ falls})}}. \end{aligned} \quad (15)$$

The union therefore raises aggregate TFP unless the two location terms are jointly negative and large enough to outweigh the other two.

Each channel is a moment of the reallocation the union induces ($\Delta \ln M_{js}$): composition captures its covariance with fundamental productivities $\bar{\Phi}_{js}$, agglomeration captures its mean together with its covariance with initial cluster size, and allocative efficiency its variance. The signs follow from that reading.

The allocative and resource gains are positive regardless of production reallocation. They are delivered by different generation mixes across countries, resulting in variation in price and supply elasticity ε_j . The allocative gain scales with the dispersion removed by merging the merit orders. The resource gain is the fossil generation displaced across the merged merit order, which shrinks Europe's fuel-import bill.

Price equalisation reallocates production in two competing ways, which is why both location channels can flip sign. Partially, price equalisation unwinds policy residuals, such that countries whose high prices reflect policy choices rather than resource scarcity gain production, while countries that are artificially cheap lose excess activity. A union price also compresses resource-based differences, with a country's natural generation advantage delivered through cross-border trade rather than by attracting production, so naturally cheap locations export part of that advantage and may lose electricity-intensive concentration. Both location channels come out positive when the first dominates: the covariance with $\bar{\Phi}_{js}$ of the countries the union relieves are the more productive ones, and the agglomeration term if the clusters that grow were already the large ones.

⁴³ Λ converts a change in firm mass into a change in expenditure weight: cell expenditure moves as $d \ln(M_{js} \bar{\pi}_{js}) = \Lambda d \ln M_{js}$, which is why the reallocation of firms shifts the weights the productivity index applies. It splits into the output that arrives with the firms, $1 + (\sigma - 1)\gamma$, counting both their own production and the agglomeration gain they confer on the firms already there, and the scale gain of those incumbents as costs fall, $(\sigma - 1)(\psi - \gamma)$, since mass rises only because unit costs fell by $\psi - \gamma$ per unit of $\ln M_{js}$ along the location equation. It is the same composite as the activity multiplier of Proposition 2, $\mathcal{A} = \Lambda / (\psi - \gamma)$: one measures the expenditure response per unit of reallocation, the other per unit of cost change.

Which lever can act on the price dispersion depends on how much of it is policy rather than geography. The supply side already separates the two: low-cost capacity $\bar{E}_j^0$ combines a natural potential θ_j^R that is fixed with installed capacity K_j^R and K_j^N that is chosen. Generation siting can remove only the part of the dispersion that reflects where capacity was built, not the part set by natural potential. Integration closes the price gap whatever its source, by letting cheap power reach demand instead of requiring production to move to it. Section 6 compares both. Aligning generation with natural potential measures the generation policy share of demand-weighted price dispersion, which turns out to be small. Market design rather than generation policy is the lever with larger potential gains.

3.7 Limitations and extensions

A few modelling choices restrict the framework in ways the data cannot settle. Each bears on how large the measured gain is rather than on the mechanism that produces it, and their biases do not all point the same way. Appendix B derives the alternatives, and Section 6.1 (with Appendix G) prices those the quantitative model can evaluate.

On the supply side, generation capital is predetermined and the merit order clears annually, so the model abstracts from renewable intermittency and from any feedback between industrial concentration and capacity build-out.⁴⁴ Appendix G endogenises investment and finds the resulting exposure confined to the fuel-saving channel, leaving the sectoral results unchanged.

On the firm side, three restrictions push in different directions. Common attachment ψ and agglomeration γ make the location multiplier a single scalar, so all cross-sector variation in the firm response runs through electricity intensity β_s . That is what the cross-sector regression identifies, and sector-specific ψ_s or γ_s are observationally equivalent in it (Appendix B). Because electricity-intensive manufacturing clusters through input sharing rather than the knowledge spillovers that generate the largest density gains (Henderson, 2003, Ellison, Glaeser and Kerr, 2010), a common γ if anything overstates both the responsiveness of these sectors and the agglomeration value at stake when they move.

The total firm mass M_s is exogenous, and a Melitz-type free-entry extension (Appendix E) would endogenise it and add a selection margin. The baseline nests that case as the limit of slow adjustment, consistent with the stable aggregate firm counts in the Norwegian experiment in Section 4.

Firms are exposed to electricity only through their own purchases, so in an input–output economy the measured share β_s would understate total exposure. Cobb–Douglas unit elasticity re-

⁴⁴An annual merit order values renewable output at the average rather than at the hours it is produced. The omission matters more the higher renewable penetration goes.

stricts substitution in the same direction by fixing that share in every country.⁴⁵

Firms sell into an integrated European market, so a location's attractiveness depends on costs and not where customers are. With trade costs, distance from the continental core would offset part of a cheap-power advantage, and the naturally cheap countries are the peripheral ones, so the model overstates how much production a cheap-power periphery attracts. The counterfactuals bound how much this can matter: under the union, electricity-intensive production moves toward the countries with the strongest existing clusters rather than out to the periphery, so the penalty trade costs would impose falls on a margin the union barely uses (Section 6.1).⁴⁶

Finally, the analytical results are local, around an interior equilibrium of a fixed point in $JS + 2J$ unknowns, and global uniqueness is not established: multiple equilibria are possible with stronger agglomeration than the calibration allows, and the inelastic residual load, endogenous sectoral demand and cross-border trade of the quantitative model fall outside the analytical guarantee, so the solver verifies stability at each counterfactual.⁴⁷

3.8 What the theory delivers

The role of electricity. In this paper electricity shapes productivity through the spatial allocation of production, and what gives it that role is the supply side. At the firm level, labour and electricity enter profits symmetrically, so only cost shares determine which sectors are more exposed. What differs is the supply side. Labour is in fixed national supply, so wages rise only because entry pushes demand against a vertical supply curve. Electricity has an endogenous supply elasticity $\varepsilon_j \geq 0$, delivered by merit order pricing and the country's electricity generation mix, steepening with fossil utilisation where fossil plants set the price. How much low-cost capacity a country has installed therefore fixes both the level of its electricity price and how far that price rises when new firms arrive. Setting $\varepsilon_j = 0$ recovers the labour-only benchmark of standard economic geography, in which firm distributions are independent of the generation mix. What this paper measures lives in the gap between that benchmark and the supply elasticities countries actually face.

⁴⁵The macroeconomic debate over the Russian gas cut-off established that output losses from energy shocks hinge on the substitution elasticity near zero: a 20% gas supply cut costs 20% of output under Leontief technology but under 3% at an elasticity of 0.05 (Moll, Schularick and Zachmann, 2023, Bachmann et al., 2024). Across horizons, Hassler, Krusell and Olovsson (2021) estimate an elasticity between energy and other inputs of essentially zero at business-cycle frequency but close to unity in the long run, which is the case the union counterfactual approximates by comparing stationary equilibria. Appendix G solves the union across a range of elasticities and leaves the reallocation channels largely unchanged, with lower substitution implying larger gains in the more electricity-intensive sectors.

⁴⁶The calibration absorbs part of the restriction as well: the fundamentals $\bar{\Phi}_{js}$ are inverted to reproduce observed firm locations, so an existing market-access disadvantage of peripheral countries is already in the baseline, recorded as low $\bar{\Phi}_{js}$ rather than as distance. What the restriction leaves open is how responsive those locations are to a cost advantage, not how much production they host today.

⁴⁷Appendix B carries the technical details.

Three parameter groups. Three inputs determine the productivity impact of any policy that shifts electricity supply, and each is disciplined differently. The location multiplier $1/(\psi - \gamma)$ determines how much production moves for a given cost change (Proposition 1). It is the one input not available off the shelf, so Section 4 estimates it and Section 5 splits it into attachment ψ and agglomeration γ . The supply elasticities ε_j fix the price impulse a policy delivers (Lemma 2). They are measured from generation and resource data, with only the fossil curvature ζ calibrated in Section 5. The general equilibrium dampening $1/(1 + \varepsilon_j \Theta_j)$ is not a primitive but is assembled from the other two, together with the demand elasticity σ and the cost shares (Proposition 2). The remaining parameters feeding into it are calibrated in Section 5.

Productivity consequences. All three inputs feed into the four productivity channels of Section 3.6, which convert the reallocation they generate into TFP changes. These additionally require the productivity fundamentals $\bar{\Phi}_{js}$ and the fossil bill share s_j^F (Lemma 4), both recovered in Section 5. Proposition 3 signs the result for the market integration policy this paper evaluates. The allocative and fossil terms are positive, so integration raises productivity unless the reallocation it induces is adversely aligned with inherent productivity advantages or cluster benefits, which the counterfactual in Section 6 computes. Electricity market integration can raise aggregate productivity by influencing industrial geography, making the coordination of energy transition policy a concern for long-run productivity as much as for the green transition itself.

4 Empirical identification: production reallocation

The theory isolates a key object that needs to be pinned down empirically: how strongly production reallocates when electricity is cheaper in one location than in another, netting firms' attachment to particular locations (ψ) against agglomeration effects (γ). It determines the magnitude of the resulting reallocation from electricity price differences that is converted into aggregate productivity effects in Lemma 4. Because electricity matters to a sector in proportion to its electricity cost share (β_s), what identifies the response is a cross-sector gradient: where power becomes expensive, electricity-intensive activity should contract relative to the rest.

The primary evidence uses a quasi-natural experiment from Norway, where the 2021 interconnector rollout and the continental gas shock opened large, persistent price gaps across bidding zones, providing within-country variation driven by pre-existing grid topology. A cross-country European panel then carries the same specification to the multi-country setting the counterfactuals operate on, testing whether the Norwegian response scales up.

4.1 Quasi-natural experiment: bidding-zone price divergence in Norway

Setting. Norway provides a quasi-natural experiment in persistently different electricity prices across bidding zones that share institutions, currency, language, and macroeconomic environment. The country is separated into five bidding zones (figure 5a), in which wholesale electricity prices are equalised. Grid infrastructure was planned and established in the mid-20th century, and the current zone configuration was finalised between 2008 and 2010.


Prices across zones are historically low and very similar, due to similar generation mixes reliant on 90% hydro and residual wind electricity. Transmission capacity between zones is limited, especially along a North-South division with physical transmission bottlenecks through the mountainous interior. Southern zones (NO1, NO2, NO5) have direct connections to continental Europe via HVDC (high-voltage direct current) subsea cables to Denmark since 1976 and the Netherlands since 2008, before new connections to Germany and the United Kingdom entered service in 2021. Northern zones (NO3, NO4) are connected only to southern Norway and other Nordic countries.

Price divergence. Three shocks hit simultaneously in 2021. The new cables to Germany and the United Kingdom exposed southern zones to continental prices just as the European gas price shock drove continental wholesale prices to record highs in the second half of 2021. A drought additionally reduced hydropower output in southern reservoirs, while northern zones stayed largely insulated due to the limited north-south transmission capacity. The new interconnectors were expected to raise Norwegian electricity prices by 4.9 €/MWh when they were commissioned in 2013, due to higher exports to continental markets (Statnett, 2013, p. 57). The impact however far exceeded prior expectations, as shown in figure 5b. Average prices across the Southern zones reached 197 €/MWh in 2022, while prices in the Northern zones stayed close to pre-period levels.

These shocks introduced persistent price gaps across pricing zones. While prices had been almost identical across zones prior to the crisis, Southern prices were 1.9 times higher than prices in the North after the crisis, in 2023–2024. The structural drivers of this price gap, the new interconnectors and the tighter coupling of southern zones to continental prices, are permanent features of the grid rather than temporary conditions.

The price divergence was also perceived as persistent by firms, such that it can plausibly affect location decisions for production. Granziera et al. (2025) show that firms' inflation expectations in the high-price zones rose persistently relative to those of comparable firms in the north over this period, indicating that firms read the divergence as a lasting change in their cost environment

Figure 5: Norwegian electricity market descriptives



Note: Panel a shows electricity market bidding zones across Norway separated by colours. Grey borders indicate municipalities, with thicker lines showing county borders. Municipalities which are not clearly part of one zone (with less than 95% of their area in a single zone) are excluded from the analysis and shaded in grey. Panel b shows annual average wholesale electricity prices by zone, using data from Bloomberg. Panel c shows the median across municipality-specific firm electricity-use indices within each zone, using municipality-level data from SSB Norway. Each municipality's use is normalised to 2019 = 100, and municipalities with nonpositive 2019 use are omitted.

rather than a one-off disturbance.^{48 49}

A shock with a salient persistent effects mimics the effect of market integration the model aims to simulate. The Norwegian experiment shows the effect of a persistent price divergence. I exploit this to estimate the location response of production and simulate the opposite case in the model, where market integration will lead to a persistent equalisation of prices across space.

Data. The analysis draws on three levels of data. At the entity level, annual financial statements from Norway’s tax authority (*Skatteetaten*) cover all incorporated legal entities (AS/ASA) from 2015 to 2024, providing balance sheet and profit and loss items.⁵⁰ Firm postcodes are matched to municipalities, which are assigned to bidding zones by geography. Only the most recent registered postcode is observed, so within-country relocations cannot be tracked directly in these data. Separate ongoing work using the full Norwegian administrative entity register finds very few firm relocations.

At the sector level, I construct electricity cost shares $\tilde{\beta}_s$ from Exiobase 3 (Stadler et al., 2018), a multi-regional input–output database that disaggregates energy purchases by carrier, allowing monetary electricity expenditure to be isolated from gas, steam, and other energy inputs. The exposure measure is electricity purchases over gross output, averaged across European countries over the latest available years, 2010–2011. Exposure is therefore a pre-period characteristic at the sector level that avoids mechanical correlation with the electricity-price variation used for identification. Wage cost shares $\tilde{\alpha}_s$ are measured as personnel costs over production value from Eurostat Structural Business Statistics.⁵¹

At the regional level, firm electricity consumption by municipality is sourced from Statistics Norway (SSB). Municipality wages are proxied using mean monthly earnings of full-time employees working in the municipality from SSB. Annual wholesale electricity prices by zone are from ENTSO-E and Bloomberg.⁵² For aggregate outcomes, I treat each municipality as a local economy and group firms into municipality–sector cells. This preserves differences in industrial composition across places within the same bidding zone. Each cell is assigned its zone’s wholesale electricity price. Appendix C provides balance tests and summary statistics.

⁴⁸They find a persistent differential in 12-month ahead expectations for inflation from 2021Q1 until the end of their sample in 2024Q4 by 1 percentage point for business leaders, larger for longer term expectations. Household expectations do not react differently across regions, reflecting the policy support with which they were insulated from the shock.

⁴⁹To the extent that some firms nonetheless anticipated mean reversion through the public debate about redrawing bidding zones or through policy support, their responses would be muted relative to a fully permanent differential, so the estimated elasticities would be conservative for the long-run reallocation the counterfactuals consider.

⁵⁰The data are sourced from Wahlstrøm (2022) and cover the universe of incorporated entities filing with the Norwegian Tax Administration. Appendix C describes the dataset, cleaning procedures, and sample coverage.

⁵¹Appendix C describes the Exiobase construction, the SBS labour-share construction, the concordance from Exiobase industries to NACE Rev. 2, and the mapping to the harmonised sector classification used throughout.

⁵²See appendix C.5 for details.

Descriptives. At the municipality level, electricity use by firms increased in low-price zones relative to high-price zones after the price divergence. Figure 5c shows the median municipality-specific electricity-use index by pricing zone. The paths develop similarly until 2021. By 2024, the median indices are 104 and 111 in the northern zones NO3 and NO4, compared with 90, 96, and 95 in the southern zones NO1, NO2, and NO5.⁵³ The median describes local changes, while zone totals show a more muted contrast in aggregate electricity use (see appendix figure D.2).

Financial outcomes and firm counts at municipality level across all sectors show muted differences between zones. Municipalities in low-price zones see higher firm operating profits and revenue than those in the south, and there is a slight relative increase in the number of active firms, although none of these differences are statistically significant (see appendix figure D.3). Location aggregate outcomes therefore show a significant effect in terms of electricity use, but no broad economic contraction in terms of real financial.

Sectoral exposure. The model predicts that a common local electricity-price shock should matter more in sectors for which electricity is a larger cost component. Figure 6 examines this by comparing the growth in the number of active firms by sector in high- and low-price zones.⁵⁴

During the period of low and equal electricity prices across zones, growth rates between the 2015–2017 and 2018–2019 averages do not differ significantly between high- and low-price zones, with all sectors close to zero growth difference. Comparing the period after the shocks from 2021–2024 when prices diverged significantly to the full pre-period of 2015–2019 however shows a different pattern. Sectors with higher electricity shares grew less in higher-price regions. The most exposed sectors are subdivisions of manufacturing, for example basic metals (C24), which has an electricity cost share in output of 3.9% and grew 18.5% less in southern high-price zones than in northern low-price zones.


I give each municipality–sector cell equal weight because the headline descriptive question concerns proportional changes across local sector locations. Weighting by initial firm counts instead gives larger cells more influence. Appendix figure D.5 repeats both panels with weights equal to each cell’s pre-period mean active-firm count within its sector and North/South group, giving a post-period slope of -0.64. Thus the headline gradient is clearer across local cells than among larger pre-existing cells. I interpret it as a local-cell pattern, and report the analogous size-weighting sensitivity when introducing the model-aligned regression.

The muted aggregate responses with price divergence therefore mask a compositional shift, as high-price regions lose active firms in electricity-intensive sectors relative to the lower-price regions, even as their overall firm count continues to grow. The gradient shows that the relevant

⁵³ Appendix figure D.1 reports an unweighted municipality-level log-outcome event study, which estimates a different statistic from the median paths.

⁵⁴ At municipality level, appendix figure D.4 shows the corresponding long difference specifications without any sectoral variation.

Figure 6: Difference in active firm count growth



Note: Each dot represents one of 13 sectors, with manufacturing split into the most electricity-intensive divisions C17, C20, C23, and C24, while keeping other sectors aggregated. For municipality–sector cells with positive period-average active-firm counts in both periods, I calculate the log change in counts and average it with equal cell weight within the South (NO1, NO2, NO5) and North (NO3, NO4). The vertical axis is $100[\exp(g_s) - 1]$, where g_s is the South-minus-North difference in those mean log changes. This is the exact percentage difference between the groups' geometric-mean count-growth factors. Panel (a) compares 2018–19 with 2015–17 and contains 3,040 municipality–sector cells. Panel (b) compares 2021–24 with 2015–19 (3,092 cells). The horizontal axis is the predetermined sector electricity cost share expressed in percentage points. The line fits the 13 sector dots with equal sector weight. Its shaded band and heteroskedasticity-robust standard error describe the fit across those dots, rather than uncertainty across the five price zones. The figure describes growth across active local sector cells, not growth in total sector firm counts.

variation is not the aggregate comparison across price exposure alone, but the interaction between the price and sector electricity intensity, which the model maps to structural parameters (Proposition 1). The model-aligned specification below formalises this relationship and estimates the moment needed for the quantitative calibration.

Model-aligned specification. The specification is written to match the model objects directly, identifying the three empirical moments needed for the quantitative calibration: the continuing-firm profit (π_{fsmzt}) response, the firm-mass (M_{smzt}) response, and their combined effect on municipality-sector profits (Π_{smzt}). I estimate, for $Y \in \{\pi, M, \Pi\}$,

$$\log Y_{fsmzt} = \delta_Y \left(\tilde{\beta}_s \log p_{zt}^E \right) + \rho_Y \left(\tilde{\alpha}_s \log w_{mt} \right) + \eta_f + \eta_{zt} + \eta_{st} + \varepsilon_{fsmzt} \quad (16)$$

The dependent variable is operating profit for the intensive-margin moment. For the firm-mass and total-profit moments, I estimate the analogous municipality–sector–year specification, replacing firm fixed effects with location fixed effects at the aggregation level of the outcome. Zone-year fixed effects η_{zt} absorb all common shocks at the bidding-zone level, including the level effect of the price shock, drought conditions, and other zone-wide changes. Sector-year fixed effects η_{st} absorb national sectoral demand and technology shocks. Identification therefore comes from the differential response of sectors with different pre-period electricity intensity $\tilde{\beta}_s$ to the same

zone-year electricity-price variation.

The wage interaction is included because the model implies opposite cross-sector loadings on electricity and labour costs (Section 3.5). However, local wages are themselves equilibrium outcomes of the shock rather than exogenous controls. I therefore report the model moments both with and without the wage interaction $\alpha_s \log w_{mt}$. The specification without wages is the reduced-form electricity-intensity gradient, whereas the specification with wages is the model-aligned decomposition between electricity-cost and labour-market channels. The key diagnostic is whether the electricity coefficient δ is stable when the wage interaction is included.

The model's wage offset also assumes labour cannot move between locations, which describes EU member states better than bidding zones inside one country: Norwegian workers move between zones more freely than across borders, so the offset should be weaker here than the cross-country model implies. The key diagnostic is whether the electricity coefficient δ is stable when the wage interaction is included.

These three margins map into the model as follows:

$$\hat{\delta}_\pi = -\frac{\sigma}{\sigma-1} \frac{(\sigma-1)\psi}{\psi-\gamma}, \quad \hat{\delta}_M = -\frac{\sigma}{\sigma-1} \frac{1}{\psi-\gamma}, \quad \hat{\delta}_\Pi = \hat{\delta}_M + \hat{\delta}_\pi.$$

The common factor $\sigma/(\sigma-1)$ converts the measured expenditure shares $(\tilde{\alpha}_s, \tilde{\beta}_s)$ into the production elasticities of the model (α_s, β_s) . With a constant markup, expenditure on an input is a fraction $(\sigma-1)/\sigma$ of revenue, so a measured shares of gross output $\tilde{\beta}_s = \frac{\sigma}{\sigma-1}\beta_s$ and $\tilde{\alpha}_s = \frac{\sigma-1}{\sigma}\alpha_s$ (Section 3). Regressing on $\tilde{\beta}_s$ and $\tilde{\alpha}_s$ therefore recovers the location multiplier scaled up by $\sigma/(\sigma-1)$.

Identification. The identifying variation is therefore differential exposure to continental prices through pre-existing grid infrastructure, within a common national setting. The identifying assumption is not that southern and northern Norway are identical. They are not: the zones differ in geography, hydrology, industrial composition, and exposure to continental markets. The identifying restriction is that, conditional on the fixed effects in equation (16), remaining shocks do not differentially affect high-electricity-intensity sectors exactly when local electricity prices rise. Pre-period sector electricity intensity fixes exposure before the crisis, ruling out contemporaneous sorting into high- β_s sectors in response to the shock. Standard errors are clustered at the municipality level throughout.

Results. Table 1 reports the estimates. The combined moment $\hat{\delta}_\Pi$ on municipality–sector total profits is negative and significant (-14.51): where local electricity prices rise, profits fall in proportion to a sector's electricity intensity. This is the aggregate object the model's location equation speaks to, and the two margins below decompose it.

Almost all of it is extensive. The coefficient $\hat{\delta}_M$ on active firms is large and significant (-15.82),

so the number of firms in a municipality–sector cell falls sharply in electricity-intensive sectors when local prices rise. To illustrate magnitudes: a 10% electricity price increase implies a $\hat{\delta}_M \times \beta_s \times 0.1$ log-point reduction in firm mass, which for manufacturing as a whole ($\beta_s = 0.017$ of gross output) amounts to roughly 2.7% fewer active firms, and for a high-intensity division such as basic metals ($\beta_s = 0.039$) roughly 6%.

The intensive margin, by contrast, is close to zero. The coefficient $\hat{\delta}_\pi$ on continuing-firm profits is negative (−1.98) but imprecisely estimated: within-firm profit adjustment scales with the electricity-intensity gradient in sign, but is not individually significant. One candidate explanation is contract structure: if firms buffer wholesale-price movements through fixed-price and bilateral agreements, their measured profits would track the wholesale price only with incomplete and delayed pass-through over the sample horizon.

The mechanical sum of the wage-controlled point estimates is $\hat{\delta}_M + \hat{\delta}_\pi = -17.81$, somewhat larger in absolute value than the total-profit estimate, as an internal-consistency check. The economically informative over-identifying restriction is the ratio $\delta_\pi / \delta_M = (\sigma - 1)\psi$, which the imprecision of the intensive margin leaves uninformative in this sample and which Section 5.2 therefore treats as a cross-check rather than an identifying moment. For all three margins, the electricity coefficient is stable when the wage interaction is included, indicating that the cross-sector gradient is not confounded by differential exposure to local labour-market conditions.

Table 1: Norwegian model-aligned moments across adjustment margins

	Firm profit		Active firms		Total profits	
	(1)	(2)	(3)	(4)	(5)	(6)
Log electricity price x electricity share	-1.949 (1.867)	-1.957 (1.867)	-16.011*** (3.414)	-15.656*** (3.254)	-14.071** (6.838)	-13.838** (6.796)
Log earnings x labour share		0.286 (0.267)		9.057*** (1.067)		5.667*** (1.650)
Observations	435678	435678	81767	81767	62036	62036
R^2	0.776	0.776	0.767	0.770	0.474	0.474

Notes: Columns report continuing-firm profits, active firms, and total profits, each with and without the wage interaction. All specifications include the relevant location fixed effects, zone-year fixed effects, and sector-year fixed effects, as in equation (16). Standard errors clustered at the municipality level. Appendix table D.8 shows intensive margin responses across different variables and transformations.

Robustness. Appendix Table D.6 reports the same specification in a border-adjacent municipality sample. These estimates are much less precise, as expected from the smaller and geographically restricted sample, so the all-zone specification remains the preferred calibration moment. Appendix Table D.7 reports the revenue-weighted intensive margin. It raises the magnitude of

the intensive margin response, but does not create sensitivity to the wage interaction, supporting the unweighted specification as the transparent baseline.

The price spike also triggered a targeted policy response. In late 2022 the Norwegian government introduced the *Energitilskuddsordningen*, a temporary support scheme for electricity-intensive firms in high-price zones that subsidised electricity costs above a threshold. A second tier tied support to energy-efficiency investments, which can raise measured fixed assets through mandated upgrades. Take-up was not widespread in the full firm population, so the scheme is unlikely to mechanically drive average responses. Because the subsidy is a post-treatment response to the price divergence, it is excluded from the baseline specification. Appendix C describes the scheme and tracks firm take-up.

The Norwegian experiment delivers the response the rest of the paper scales. A persistent, location-wide price gap moves the number of active firms in proportion to a sector’s electricity intensity, almost entirely through firm mass rather than the scale of incumbents.

4.2 Cross-country evidence

The Norwegian experiment recovers the reallocation elasticities from a single, sharp shock. A complementary cross-country exercise applies the same cross-sector interaction logic to a panel of firms across Europe. This setting does not allow for a clean quasi-experiment with causal interpretation, but tests the applicability of the results from within Norway to a broader geographic scope directly relevant to the policy counterfactuals.

Data. Firm-level outcomes come from Orbis annual financial statements, using gross profit and operating revenue from 2005 onward, with country of registration as the location identifier. Electricity prices are annual average day-ahead wholesale prices from ENTSO-E and Bloomberg, used at the bidding-zone level in the border comparison below and aggregated to national averages with load weights for the firm-level estimation. The wage is country-level compensation per employee from Eurostat national accounts. The sector loadings are the same fixed pre-sample sector characteristics as in the Norwegian estimation to ensure comparability of coefficients. Appendix C.2 provides sample coverage, the concordance to firm-level sector codes, and summary statistics, and Appendix C.5 describes the price construction. A parallel employment check plausibly independent of accounting changes draws on Revelio Labs firm-level employee counts constructed from online worker profiles, reported in Appendix D.3.

From zonal to national borders. Carrying the response elasticities across the two settings requires that the price variation is the same economic object in both: no larger, no differently composed, and no less persistent across borders than within them. I test this in a zone-pair panel of these prices (2015–2025), applying the border-effect regressions of Engel and Rogers (1996) to an-

nual wholesale electricity price levels, where a price gap reflects only transmission constraints and market segmentation. Appendix D.6 details the sample construction, the interconnection-capacity measure, and the inference.

Table 2: Border versus bidding-zone price gaps in Europe

Panel A: Mean absolute log price gaps			Panel B: Border premium		
	Full period	Pairs		(1) Raw	(2) Conditional
Within-country zones	0.156	38			
Border, adjacent	0.169	60	National border	0.220*** (0.033)	0.148*** (0.049)
Border, non-adjacent	0.388	763	NTC / min. load		0.027 (0.040)
			Observations	8473	659
			Zone FE	No	Yes

Notes: Zone-pair panel of annual average day-ahead wholesale electricity prices across European bidding zones, 2015–2025. Panel A reports mean absolute log gaps by pair type, with the number of pairs. Panel B regresses the gap on a national-border indicator: column 1 with year fixed effects on all pairs; column 2 adds interconnection capacity (net transfer capacity over minimum zone load) and symmetric zone fixed effects, adjacent pairs only. Standard errors clustered by zone pair in parentheses. The randomisation-inference p -value permutes the border label across adjacent pairs over 1,000 draws. Appendix D.6 details the construction and reports the Nordic sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

In Europe, cross-border gaps look like zonal gaps: the mean gap across adjacent borders is close to the within-country mean, and the large gaps open at poorly connected pairs, border or not (Table 2, Panel A). Conditioning on interconnection capacity and zone fixed effects (Panel B) cuts the raw border premium by a third. In the Nordics (four multi-zone countries sharing one exchange and market design) the comparison yields similar or smaller gaps and a border premium indistinguishable from zero (Appendix D.6).

A national border thus confronts a firm with the same economic congestion-priced gap as a zonal boundary, whose range the Norwegian variation spans. Event studies of changing market structure further show that integrating two zones closes their price gap, while drawing a new zonal border creates almost none (Appendix Figure D.18). Because the treatment is the same economic object in both settings, any gap between the Norwegian and the cross-country response estimates loads onto the response side rather than onto differences in the underlying price variation.

Specification. To isolate the response margin, I run two designs in the cross-country panel. The first is the levels analogue of the Norwegian equation (16), translated by replacing firm fixed

effects with firm–country fixed effects and zone–year with country–year fixed effects:

$$\log Y_{fsct} = \delta \left(\beta_s \log p_{ct}^E \right) + \rho \left(\alpha_s \log w_{ct} \right) + \eta_{fc} + \eta_{ct} + \eta_{st} + \varepsilon_{fsct}. \quad (17)$$

The second is a long-difference version over horizon h , from $t - 1$ to $t + h$:

$$\Delta^h \log Y_{fsct} = \delta^h \left(\beta_s \Delta^h \log p_{ct}^E \right) + \rho^h \left(\alpha_s \Delta^h \log w_{ct} \right) + \eta_{fc} + \eta_{ct} + \eta_{st} + \varepsilon_{fsct}, \quad (18)$$

with $\Delta^h X_{ct} = X_{c,t+h} - X_{c,t-1}$ and $h \in \{0, \dots, 5\}$, so the windows span one to six years. Identification comes from the differential response of sectors with different electricity cost shares β_s to the same national price movement, using within-firm time variation. The levels design therefore cumulates responses to persistent shifts in relative national prices, while the long differences in (18) isolate horizon-specific price changes, so the two designs bracket persistent against transitory variation.

Identification. Identification rests on the same cross-sector loading as the Norwegian model-aligned specification: country–year fixed effects η_{ct} absorb all national shocks, leaving only the differential response across sectors with different electricity shares to the same country-level price change. Firm–country fixed effects η_{fc} control for persistent firm-level heterogeneity, and horizons $h \in \{0, \dots, 5\}$ capture short- to medium-run adjustment. The fixed-effect structure absorbs a large set of confounders, but residual threats remain. Industrial composition feeds back into national prices through the demand side of the merit order, likely leading to reverse causality, and omitted sector–country shocks may correlate with both electricity intensity and firm outcomes.

Results. The levels profit gradient is negative, significant, and stable to the wage interaction — the same diagnostic that holds in Norway (Table 3, Panel A). A one-log-point rise in the electricity-cost-weighted national price lowers gross profit by $7.95 \beta_s$ log points, and the gradient survives in inverse hyperbolic sine, which retains the loss-making firms the log specification drops. The sign matches the Norwegian intensive-margin moment and confidence intervals overlap, but the cross-country point estimate is considerably larger. To fix magnitudes, a doubling of the national electricity price, on the order of the gap between cheaper and more expensive large European markets (Section 2), implies about $0.69 \times 7.95 \times \beta_s$ lower gross profit: around 6% at the median cost share ($\beta_s \approx 0.011$) and around 21% for a high-intensity sector such as basic metals ($\beta_s \approx 0.039$). Two readings are consistent with this. Contracts reprice to persistent national price differences, whereas hedging dampened the response to the Norwegian crisis spike; or higher frictions on cross-border reallocation shift more of the incidence of persistent price differences onto incumbents' margins.

Operating revenue is flat in levels, implying that national cost differences are not passed into

output prices. A positive nominal revenue response at longer long-difference horizons (Appendix Table D.10) is consistent with cost pass-through raising nominal revenue while gross profit stays flat. A parallel check using Revelio Labs employee counts finds no employment decline within continuing firms, with a gradient that is small, positive, and significant only at the shortest long-difference horizon (Appendix Tables D.11 and D.12), consistent with the muted within-firm quantity adjustment in Norway.

The long-difference profit gradient is instead null at every horizon of one to six years (Table 3, Panel B). The profit gradient lives in the persistent price component that the levels design captures, consistent with the stylised fact that European price differences are long-lived state variables, and with why the Norwegian experiment is the identification setting while the counterfactuals target permanent differences.

Table 3: Cross-country intensive margin responses

Panel A: Levels						
	Log gross profit		IHS gross profit		Log operating revenue	
	(1)	(2)	(3)	(4)	(5)	(6)
Log electricity price $\times$ electricity share	-7.979*** (2.835)	-7.948*** (2.820)	-10.442*** (3.147)	-10.787*** (3.193)	-1.225 (1.392)	-1.218 (1.422)
Log compensation $\times$ labour share	0.241 (0.321)		-2.701 (3.283)		-0.024 (0.298)	
Observations	8990880	8990880	9163916	9163916	14877493	14877493
R^2	0.916	0.916	0.611	0.611	0.916	0.916
Panel B: Long differences of log gross profit						
	(1)	(2)	(3)	(4)	(5)	(6)
	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$
Δ^h log electricity price $\times$ electricity share	-1.691 (1.134)	-1.039 (2.807)	0.975 (1.863)	1.841 (1.299)	0.969 (1.118)	0.435 (2.152)
Δ^h log compensation $\times$ labour share	-0.500 (0.354)	-0.210 (0.242)	0.115 (0.623)	3.022 (2.302)	4.317 (2.810)	3.722 (2.465)
Observations	6860625	5279754	3764759	2918673	2192837	1676627
R^2	0.245	0.397	0.443	0.523	0.590	0.662

Notes: Panel A reports estimates of equation (17) on the Orbis firm panel; the dependent variable is log gross profit (columns 1–2), the inverse hyperbolic sine of gross profit (columns 3–4), which retains firms with non-positive gross profit that the log specification drops, and log operating revenue (columns 5–6). Within each outcome, the first column includes the labour-cost interaction and the second omits it. Panel B reports estimates of equation (18) for log gross profit, with columns corresponding to horizons $h = 0, \dots, 5$, i.e. windows of one to six years, and the labour term included throughout. All specifications include firm–country, country–year, and sector–year fixed effects. The sample is a European firm panel in NACE Rev. 2 sections B, C, F, G, H, I, J, M, and N, excluding oil- and gas-extraction sectors (NACE 06 and 09). Standard errors, clustered by country, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The cross-country panel therefore shows that the Norwegian firm response is not a Norwegian artefact. The same cross-sector loading appears across European borders, and it attaches to persistent price differences rather than transitory ones, which is the variation the counterfactuals operate on.

4.3 Summary of evidence

The evidence delivers three inputs to the quantification. First, the cross-sector structure of location choice in Proposition 1 holds in the data: firm activity contracts where electricity is expensive, in proportion to a sector's electricity cost share, in Norway and again in the European panel. Second, the Norwegian response runs almost entirely through the number of firms rather than the scale of continuing ones, so the spatial reallocation the counterfactuals compute is driven by changes in the number of active firms rather than by the rescaling of incumbents. The extensive margin pins down the composite of location attachment and agglomeration that governs how far firms move. Splitting the composite into its two parts would take the intensive margin, which is both imprecise and plausibly attenuated by hedging, so Section 5.2 anchors agglomeration externally and recovers location attachment residually. Third, the price variation is the same economic object in both settings, since national borders price transmission congestion no differently than zonal boundaries, so the Norwegian estimate can discipline a multi-country model. What differs is the response: the cross-country profit gradient is steeper than the Norwegian intensive margin and attaches to persistent rather than transitory price differences, consistent with contracts repricing to long-lived cost gaps and with cross-border frictions concentrating their incidence in incumbents' margins.

Taken together, the evidence answers the question the theory posed: production does move across borders when electricity prices differ persistently, and it does so on the extensive margin, which sets industrial geography over the horizon a single market would operate on.

5 Model quantification

The theory of Section 3 ties the productivity loss from price dispersion to the allocation of production when a location's costs change. Section 4 causally estimates the key location multiplier $\frac{1}{\psi-\gamma}$ in partial equilibrium. Valuing market integration requires the rest of the economy, and crucially the electricity supply elasticity ε_j when new demand arrives, and how much reallocation is dampened in general equilibrium, captured by Ω_j . This section builds that economy including those key objects and validates it, so that Section 6 can price the missing single market.

This section first adds quantitative extensions to the analytical model, which do not change the economic mechanisms but are required to obtain credible quantitative estimates of the value of market integration. Electricity is traded up to capacity constraints, the supply elasticity moves

with fossil utilisation, households use electricity, and goods prices and income adjust when production moves. Quantification then follows as a calibration disciplined by causal estimates. The key location multiplier is taken from the empirical estimates of Section 4 ($\hat{\delta}_M$) and mapped into structural parameters (Proposition 1). The remaining parameters are taken from data or from external estimates, and the productivity fundamentals are inverted to reproduce the observed allocation. This two-step procedure follows the standard in quantitative spatial and trade economics (Suárez Serrato and Zidar, 2016, Fajgelbaum et al., 2019, Caliendo and Parro, 2015, Arkolakis and Walsh, 2023).

The rest of the section sets out those extensions, the parameterisation, the solution method, and the fit against moments the calibration does not target. The section closes on the model-implied macro elasticity of a country's production to its electricity price, which the model puts at roughly three-fifths of the estimated cross-sector micro elasticity after accounting for general equilibrium adjustments.

5.1 Quantitative model

The quantitative model extends the analytical framework in six directions. Each closes a channel that the analytical model deliberately abstracts from but that matters for the magnitude (though not the sign) of the counterfactual TFP gap.

Cross-border electricity trade. The quantitative model introduces bilateral electricity trade up to the cross-border transmission capacity between connected countries to capture the realistic baseline of heavily constrained trade. Let τ_{ij} denote the available transfer capacity from country i to country j and x_{ij} the scheduled flow. Each national market then clears with imports and exports,

$$S_j(p_j^E; K_j) + \sum_i x_{ij} = E_j^D(p_j^E, M, w, H_j) + \sum_k x_{jk}, \quad 0 \leq x_{ij} \leq \tau_{ij}, \quad (19)$$

where S_j is the country's merit-order supply schedule (Lemma 2) and E_j^D its electricity demand. Market coupling determines prices and flows jointly: electricity moves from low- to high-price countries until prices either equalise or an interconnector binds. The residual gaps are the shadow values of scarce transfer capacity.⁵⁵

A key consequence of trade is that the supply elasticity facing firms becomes network-wide rather than only locally determined. Where interconnectors are uncongested, firm electricity de-

⁵⁵This representation errs toward understating the gains from integration. The model clears an annual market with annual-average transfer capacities, so a link is congested or slack for the year as a whole. In reality congestion binds hour by hour, and an interconnector that is slack on average is often saturated precisely in the hours that set prices. Measured capacities also overstate in comparison to operators' discretionary reliability margins and reservable buffers to loop flows displacing internal congestion onto borders. The baseline network is thus credited with more delivered integration than the hourly market achieves.

mand in one country draws on the merit-order supply of the entire connected component, so the local congestion force ε_j is replaced by an effective elasticity of the integrated market, and the country-specific ε_j re-emerges only on borders where transmission binds.

Residual non-firm electricity demand. The analytical model treats electricity demand as firm demand in order to keep the GE expressions transparent. The quantitative model separates total load into three objects: observed national system demand, which fixes the scale of the electricity market; modelled-sector demand from the thirteen production sectors in the firm-location block; and an exogenous residual.

$$E_j^D = H_j + E_j^{\text{firm}}, \quad E_j^{\text{firm}} = \sum_s M_{js} \frac{\beta_s(\sigma - 1)}{p_j^E} \bar{\pi}_{js},$$

where E_j^{firm} is modelled-sector demand and H_j is the residual between observed system demand and modelled-sector demand, capturing household, public-sector, and non-modelled business consumption as well as network losses and accounting differences. Firm entry moves only the modelled-sector component of total load. H_j is held fixed, which attenuates electricity congestion relative to a model in which every unit of demand is generated by the firms whose locations are endogenous.

Endogenous supply elasticity. In the analytical comparative statics, the electricity supply elasticity ε_j is evaluated at the initial equilibrium and held fixed. The quantitative model updates ε_j as utilisation u_j adjusts in each counterfactual, following equation (7).⁵⁶ The analytical logic suggests this should reinforce the baseline results: adding low-cost capacity simultaneously lowers both the electricity price and the local supply elasticity, so the flatter supply curve would absorb subsequent firm entry with smaller price feedback.

Goods-market general equilibrium. The analytical model holds the sectoral demand shifter B_s constant, so that location comparisons are unaffected by aggregate price movements. The quantitative model closes the goods market explicitly, following Caliendo and Parro (2015) and Arkolakis and Walsh (2023): a representative consumer has Cobb–Douglas preferences across sectors with fixed expenditure shares μ_s and CES preferences over the differentiated varieties within each sector, which are sold in an integrated European market. Sectoral expenditure $Y_s = \mu_s I$ and the within-sector price index P_s are then endogenous, where I is aggregate income: labour income and electricity rents net of the imported-fuel cost of fossil generation, plus a fixed residual for

⁵⁶Relative to the full intertemporal dispatch solve in Arkolakis and Walsh (2023), ε_j is a tractable reduced summary of the local supply-curve slope at the operating point, which keeps the multi-country equilibrium computable while retaining the congestion mechanism.

payments to capital and intermediate inputs. As firm activity reallocates toward lower-cost countries, P_s falls, raising real income and feeding back through B_s into firm demand, factor demands, and the equilibrium wages and electricity prices, and delivering a consumer-price channel absent from the analytical solution. Because the Cobb–Douglas aggregator holds the cross-sector composition of expenditure fixed, this feedback operates within sectors and through aggregate income rather than by reallocating expenditure across sectors. Appendix B records the demand system, the income closure, and the price-index expression.

Aggregate productivity. The counterfactuals report the four-channel decomposition (14) of Section 3.6. The quantitative model changes how it is evaluated rather than what it contains: where (14) is a transparent fixed-weight, first-order benchmark, the calibrated model computes the exact change in aggregate TFP between equilibria, with output weights ω_{js} and marginal costs revalued at the counterfactual allocation and the endogenous-weight covariance terms of Appendix B retained in full.

Emissions accounting. Any change in fossil dispatch ΔE_j^F also impacts direct carbon emissions of generation. To account for different emissions factors, each country’s fossil generation is split across four supply blocks by fuel (gas, coal, other fossil, and bioenergy). Emissions are evaluated through fuel-specific heat-rate schedules at which they generate electricity and their emission factors e^F . Fossil fuel prices p^F vary by fuel type and contain a carbon price depending on their emissions factor, with CO₂ priced at the EUA allowance rate. When counterfactuals change fossil dispatch, they can hence influence generation emissions by reducing total fossil generation proportionally, or through shifts in composition across fuels. Emissions accounting covers direct combustion CO₂ only, with local air pollutants and any interaction with the ETS cap left as extensions.

Each extension changes the magnitude of what integration is worth without changing the mechanism: firms still choose countries on cost, and the four channels of Lemma 4 still convert those choices into aggregate productivity. The extensions add the realism the counterfactuals need, especially electricity trade, a goods market that clears, and a fuel bill that can be priced in carbon.

5.2 Parameterisation

The quantitative model is calibrated to a sample of 23 EU member states and thirteen NACE Rev. 2 sectors. Countries exclude the micro-states Luxembourg, Malta, and Cyprus, which have negligible own generation, and Ireland, whose most meaningful electricity interconnections run to Great Britain outside of the EU. The sectors are the sections B, F, G, H, I, J, M, and N, with manufacturing split into its four most electricity-intensive divisions — paper (C17), chemicals (C20), other

non-metallic mineral products (C23), and basic metals (C24) — and the remainder, to capture the dispersion in electricity intensity (Section 2.3). Together these cover mining, manufacturing, construction, and the main market-service sectors, while excluding electricity and water utilities (D/E), agriculture, finance, real estate, public administration, and other non-market services. The baseline equilibrium is the 2015–2019 average, which predates the pandemic and the energy crisis so that it approximates a stationary equilibrium.⁵⁷

Parameterisation proceeds in five steps: indirect recovery from the empirical estimates of the location multiplier of Section 4, calibration of the electricity supply elasticity, combined with additional firm-side parameters to obtain the general equilibrium dampening, additional parameters for counterfactual exercises, and model inversion to recover the remaining objects. Table 4 summarises each parameter’s role, value, and source.

Location multiplier from estimates (Panel A). The Norwegian extensive margin pins the location multiplier directly. With $\hat{\delta}_M \approx -15.8$, firms’ attachment to their locations net of agglomeration is $\psi - \gamma = 1/|\hat{\delta}_M| \approx 0.063$, and this composite governs how much production moves in every counterfactual below. What the estimates do not deliver is the split of the composite into its two parts, and the two enter the productivity channels differently: agglomeration γ raises the productivity of every firm in a growing cluster, while location attachment ψ only slows the move.⁵⁸

I anchor agglomeration externally, at the selection-purged elasticities of Combes, Duranton and Gobillon (2008) and Greenstone, Hornbeck and Moretti (2010) to $\gamma = 0.04$, and recover attachment $\psi = \gamma + 1/|\hat{\delta}_M| \approx 0.10$. I evaluate sensitivity of the main counterfactual to this choice over a band of $\gamma \in [0, 0.08]$, which carries through to ψ .^{59 60}

⁵⁷The baseline is computed from available non-missing observations subject to a minimum of one year per country-sector cell.

⁵⁸In principle the second margin performs the split. The ratio of the continuing-firm profit response to the firm-count response is $(\sigma - 1)\psi$ (Section 3.5), so with an external demand elasticity it separates attachment from agglomeration. The Norwegian intensive margin is too imprecise to do so: taken at face value the ratio of point estimates implies $(\sigma - 1)\psi \approx 0.13$ and hence a negative agglomeration elasticity γ . I do not read that as evidence against agglomeration. The estimate is imprecise, and the restriction additionally assumes both margins respond to the same price. If continuing firms are partially hedged against the wholesale price through long-term contracts, while entry and exit respond to the expected long-run cost of operating in a zone, the measured ratio is attenuated by the treatment definition rather than by agglomeration. This is consistent with the within-country setting: firms adjust on the entry/exit margin rather than by rescaling continuing operations, and the wage-congestion channel that a wage-based second moment would exploit is muted by labour mobility across Norwegian zones. A wage-based separation is therefore better suited to the cross-country panel, where it is reported as a robustness check.

⁵⁹Reporting ψ conditional on an assumed range for γ parallels Fajgelbaum et al. (2019), who report their migration elasticity conditional on an assumed dispersion of idiosyncratic location preferences when a single moment cannot separate the two.

⁶⁰The cross-country panel offers an independent check: substituting the levels profit gradient into the ratio restriction gives $7.98/15.82 \approx 0.50$, which implies a positive agglomeration elasticity. At a demand elasticity of $\sigma = 4$, the pair $\psi \approx 0.17$ and $\gamma \approx 0.10$, modestly above the swept band but statistically indistinguishable from it, and inside the band for σ toward the upper end of its range.

Table 4: Model parameters

Parameter	Role	Value	Source and construction
<i>Panel A: The location multiplier, from the estimates of Section 4</i>			
$-\frac{1}{\psi-\gamma}$	Location multiplier	−15.8	Norwegian extensive margin, $\hat{\delta}_M$
γ	Agglomeration force	0.04	Greenstone, Hornbeck and Moretti (2010)
ψ	Location attachment	0.10	$\gamma + 1/ \hat{\delta}_M $, given γ
<i>Panel B: The supply elasticity, computed from the generation stack</i>			
K_j^R, K_j^N, K_j^F	Capacity by technology	by country	ENTSO-E, Ember
$\bar{E}_j^0$	Low-cost generation	by country	Ember: total generation minus fossil
u_j	Fossil utilisation	0.12–0.68 (med. 0.42)	Fossil generation/ fossil capacity
c_0^F	Non-fuel fossil <i>mc</i>	3.2–5.8 EUR/MWh	DEA catalogue: variable O&M, fossil-mix weighted
ζ	Fossil supply curvature	2	EU plant cost stack
p_j^E	Baseline wholesale price	33–55 EUR/MWh	ENTSO-E day-ahead, 2015–2019 mean
ε_j	Supply elasticity	0.06–2.5 (med. 0.7)	Computed, eq. (7)
<i>Panel C: General equilibrium dampening, computed by adding firm-side parameters</i>			
β_s	Electricity cost share	0.005–0.039	Exiobase: electricity costs/ gross output
α_s	Labour cost share	0.11–0.36	Eurostat: personnel costs/ production value
σ	Demand elasticity	4	Broda and Weinstein (2006), disagg. median
μ_s	Expenditure weights	by sector	Eurostat gross-output shares
$\bar{L}_j$	Labour endowment	by country	Eurostat employment in model sectors
Ω_j	GE dampening	23–75 (med. 38)	Computed from $1/(\psi - \gamma), \varepsilon_j, \sigma, \beta$
<i>Panel D: Electricity market and counterfactual inputs</i>			
θ_j^R	Renewable potential	by country & source	JRC ENSPRESO; Hoes et al. (2017) for hydro
τ_{ij}	Transfer capacity	by country pair	ENTSO-E NTC: flows plus headroom
H_j	Residual demand	by country	Eurostat, Ember: system load minus modelled load
p^F	Fuel/carbon split of $\kappa_j p^F h_{\min}$	by country	World Bank Pink Sheet, EU ETS
$h_{\min}$	Frontier heat rate	1.72	DEA catalogue: 2015 CCGT net efficiency
<i>Panel E: Fitted to reproduce the baseline allocation</i>			
$\bar{\Phi}_{js}$	Base productivity	by country-sector	Inverted; matches observed location shares
I	Aggregate income	1	Normalised; sets the goods-market scale
κ_j	Fossil-segment wedge	0.76–1.21 (med. 0.99)	Inverted; matches baseline price

Note: The three key forces are in bold. Within each panel rows run in the direction of inference: in Panel A the force is estimated and the parameter inferred from it, in Panels B and C the parameters are measured and the force computed from them. Ranges are across the 23 countries, or across the 13 sectors for β_s and α_s , with the median in brackets.

The electricity supply elasticity (Panel B). The key supply elasticity ε_j is computed via equation (7) from inputs available in data. It is the fuel share of the observed price, $\chi_j = (p_j^E - c_0^F)/p_j^E$, divided by the curvature ζ and multiplied by demand per unit of spare fossil capacity, so a country's slope is steep where little of its price is fixed non-fuel cost and its fossil fleet is already running hard.⁶¹ Generation capacities (K_j^R, K_j^N, K_j^F) and fossil utilisation rates u_j come from ENTSO-E and Ember.⁶² The Pareto shape ζ governing the fossil supply-curve curvature is calibrated by fitting the model's merit-order supply curve to a European fossil cost stack built from plant-level heat rates, fuel prices and variable operating costs. Two variants bracket the estimate: an observed-efficiency stack, which captures heat-rate dispersion within fuels as well as switching between them, and a cross-fuel stack covering the full fleet. They give 2.0 and 2.3, and an external check against the PJM zonal supply curves of Sahraei-Ardakani, Blumsack and Kleit (2015) lands in the same range, so I take $\zeta = 2$ as the baseline and report sensitivity over $\zeta \in \{1.5, 2, 2.5, 3\}$.⁶³

The calibrated elasticities span roughly 0.06 to 2.5, ordered by spare fossil capacity relative to demand: systems whose unused fossil capacity is small next to their total load sit at the steep end, since additional demand must be met further up the merit order, while systems holding ample spare fossil capacity relative to their load sit at the flat end. Appendix F reports each country's baseline price, fossil utilisation, and supply elasticity.

General equilibrium dampening, adding firm-side parameters (Panel C). The electricity cost share β_s is electricity purchases over gross output, constructed from Exiobase 3 for 2010–2011. The labour cost share α_s is personnel costs (employee compensation including employers' social contributions) over production value, from Eurostat structural business statistics over the same window.⁶⁴ Both are gross-output cost shares, so $\alpha_s + \beta_s < 1$, with the remainder accruing to capital and intermediate inputs at common cross-country prices, consistent with the gross-output production technology used in the quantification. Shares are averaged across European countries to avoid contamination from within-country endogenous sorting. Fixing the shares at the sector level retains a unit elasticity of substitution, under which expenditure shares are technological

⁶¹Equation (7) writes the fossil segment as $p^F h_{\min}(1 - u_j)^{-1/\zeta}$, but the calibration does not build the price up from fuel costs. It anchors each country's segment to its own observed baseline price through the wedge κ_j (Panel E), so that product is pinned at $p_j^E - c_0^F$ and the fuel cost p^F and frontier heat rate $h_{\min}$ drop out of ε_j altogether. They re-enter only in the fuel- and carbon-price counterfactuals, where they split the anchored segment into its two movable components (Panel D).

⁶²Low-cost capacity $\bar{E}_j^0$, which fixes where the supply curve kinks, is an observed generation quantity rather than a capacity: total generation net of fossil generation over the baseline window, split into its nuclear and renewable slices by observed shares.

⁶³Appendix C.6 documents the plant register, the fit, and the consequences of the stack being priced without carbon; none of the paper's location channels moves by more than 1% across the band.

⁶⁴The most recent years built on actual national-accounts source data from Exiobase is available for 2010–2011. I match the computation window of the labour share for consistency.

constants.⁶⁵

The CES demand elasticity σ is set to 4, the disaggregated median of Broda and Weinstein (2006), interpreted as the within-sector elasticity of substitution across varieties, with sensitivity reported over $\sigma \in [3, 6]$. Sectoral expenditure weights μ_s , price indices, and output shares come from Eurostat national accounts. The baseline μ_s uses gross-output shares. Labour endowment $\bar{L}_j$ is measured as total employment in the included sectors, with employment in excluded sectors and non-employment held outside the model. Appendix F documents the construction.

The GE dampening Ω_j of Proposition 2 is a combination of the location multiplier of Panel A, the supply elasticities of Panel B, and the parameters above. Evaluated at those values it runs from about 23 to 75 across the sample, rising with the steepness of a country's supply curve. Ω_j is reported here for the order of magnitude that separates the estimated multiplier from what survives general equilibrium, but it is a single-sector composite, and what the multi-sector model actually delivers is the subject of Section 5.5.

Counterfactual external parameters (Panel D). Natural potential by technology is computed using technical capacity from the JRC ENSPRESO database (Ruiz et al., 2019) for solar PV and on- and offshore wind, and gross theoretical potential from Hoes et al. (2017) for hydro, with generation recomputed at destination-country yields.⁶⁶

Bilateral transfer capacity τ_{ij} is the directed net transfer capacity published by ENTSO-E, aggregated from zone pairs to country pairs and averaged over the baseline window.⁶⁷ Counterfactual flows are constrained to these bounds, outside of the integration scenarios. Appendix C documents the contract-type hierarchy, and the effective-bound construction. Borders to countries outside the modelled sample enter as fixed net-import residuals that hold the external trade position constant across counterfactuals, since extra-EU trade accounts for only 4-5% of the total load in the baseline.

Residual electricity demand H_j is calibrated as the difference between observed country-level electricity load and the model-implied electricity use of the sectors included in the firm-location block.

The fuel-price and carbon-price counterfactuals need one more object. Because each country's fossil segment is anchored to its own observed price (Panel B), the level of that segment is already

⁶⁵For example, a country facing threefold higher electricity prices uses proportionally less electricity per unit of output but spends the same share of gross output on it. Measured country-sector shares in fact rise with national price levels within sectors, indicating substitution below unity, since the sorting of electricity-intensive activity toward cheap countries would push the relationship the other way. Cobb–Douglas is the knife-edge between the two substitution regimes, and cross-country variation in true exposure is the dimension the sector-level construction assumes away. The counterfactuals are reported under a CES sensitivity band.

⁶⁶Because baseline low-cost capacity is measured directly from observed generation (Panel B), the natural-potential parameter θ_j^R is not required separately in the baseline: observed renewable generation already embeds each country's realised yield. Appendix H.1 documents the counterfactual re-siting rule.

⁶⁷Where observed flows exceed published capacity, the solver bounds are widened minimally in the direction needed to make the baseline feasible, never using peak flows to manufacture additional counterfactual headroom.

pinned, and what these exercises require is not the level but its composition: how much of it is world fuel cost and how much is the EU ETS carbon charge. That split comes from p^F , the all-in marginal fuel-use cost, which combines the world fuel price with the EUA allowance price times the fuel-specific emissions intensity, evaluated at the frontier heat rate $h_{\min}$. Benchmark fuel prices are World Bank commodity series, weighted into a country composite by each country's observed fossil generation mix, and $h_{\min}$ is the DEA catalogue's 2015 combined-cycle gas efficiency of 58%, so $h_{\min} = 1.72$. The split remains well defined even though the level is calibrated, for the reason given under κ_j below. Appendix B records the explicit decomposition.⁶⁸

Fitted to the baseline allocation (Panel E). Beyond the parameters, several objects are fitted so that the calibrated equilibrium replicates the baseline. Baseline aggregate income I is normalised to one,⁶⁹ which sets the scale of the goods market and leaves every goods-market object reported as a ratio to its baseline value.⁷⁰

Country-sector base productivity $\bar{\Phi}_{js}$ is recovered by inverting the baseline location equation from observed firm shares, wages, and electricity prices, so that the calibrated equilibrium reproduces the observed firm distribution exactly. It is then validated against an external country-sector labour-productivity measure (Section 5.4).

With the shape of the fossil supply curve fixed, the level of each country's fossil segment is anchored to its own baseline point. The data-based composite $p^F h_{\min}$ — the DEA frontier heat rate times the all-in marginal fuel cost — is scaled by a country-specific wedge κ_j so that the supply curve passes exactly through the observed baseline wholesale price at the observed fossil utilisation. The wedge is centred on one (median 0.99, tenth to ninetieth percentile 0.76–1.21), so the cost stack accounts for the level of the fossil segment on average, and κ_j absorbs what it misses country by country: fleet vintage relative to the frontier plant, scarcity rents, and hydro opportunity cost. Because the wedge scales the fuel and carbon components together, it cancels from their ratio, so the fuel share of the price is a data object and the fuel- and carbon-price counterfactuals move each component separately while the baseline level stays pinned.

These fitted objects are held fixed across counterfactuals, so the counterfactual predictions rest on the estimated location multiplier, the externally calibrated parameters, and the supply-curve

⁶⁸The discipline of the fossil supply schedule against observed dispatch and wholesale-price data follows Arkolakis and Walsh (2023).

⁶⁹The normalisation is implemented through a residual income component standing for payments to capital and intermediate inputs, which the model does not solve for: it is set so that baseline income equals baseline expenditure and held fixed thereafter, so income varies across counterfactuals only through wages, electricity rents, and the imported-fuel bill.

⁷⁰The sectoral demand shifters B_s are not normalised: they follow from the goods-market closure given the sectoral price indices and expenditure, and because both are endogenous, B_s is solved jointly with the rest of the equilibrium and moves along every counterfactual path (Appendix B.9). In the analytical model B_s is common to all countries within a sector and so cancels from the logit location shares. In the quantitative model it still cancels conditional on factor prices, but wages and electricity prices depend on it through factor-market clearing, so it does affect the equilibrium firm distribution.

slopes. The main union counterfactual is reported over the calibration’s sensitivity ranges for agglomeration γ , demand elasticity σ , and supply curvature ζ , together with a band that propagates the first-stage standard error of the firm-mass response $\hat{\delta}_M$ through the calibration. The agglomeration parameter γ is the dominant source of parameter uncertainty. Figure F.3 in Appendix F collects the bands.

5.3 Solution method

The equilibrium is a fixed point in firm masses $\{M_{js}\}$, wages $\{w_j\}$, and wholesale electricity prices $\{p_j^E\}$, together with the sectoral price indices, aggregate income, and interconnector flows that close the goods and electricity markets. I iterate three blocks to convergence. Given prices and wages, location shares follow from the logit of Lemma 1 and wages clear each national labour market (8). Given locations and wages, electricity prices clear along the supply schedules of Lemma 2. With trade, this becomes a single market-coupling program: a convex surplus maximisation over firm electricity demand, cross-border flows, and generation by source, subject to nodal balance and transmission and capacity limits. This process mirrors the welfare maximisation the European day-ahead coupling implements. Zonal prices and congestion rents are the duals of the balance and transmission constraints, so prices equalise across uncongested borders and diverge only where a link is saturated. The goods market then updates the price indices and income that feed back into firm and factor demands. The outer state is updated with dampening and safeguarded Anderson acceleration until all market-clearing residuals fall below tolerance, following Caliendo, Dvorkin and Parro (2019) and Fajgelbaum and Redding (2022). Local stability is verified from the spectral radius of the damped update’s Jacobian at the solution. Appendix F sets out the algorithm step by step.

Counterfactuals perturb the relevant primitives — generation capacity, transmission limits, or a move to a single union-wide price — and re-solve the equilibrium in levels, so non-marginal changes such as the full electricity union are handled exactly.⁷¹

5.4 Model fit

The calibrated model is confronted with moments it does not target, following the over-identifying-moment checks of Fajgelbaum et al. (2019). The central check compares the solved productivity shifter $\bar{\Phi}_{js}$ to an external measure the model never sees — country–sector labour productivity from Eurostat structural business statistics. Across the 299 country–sector cells, the

⁷¹The piecewise merit-order schedule and the capacity-constrained trade network do not admit a clean hat-algebra representation, so unlike the exact-hat counterfactuals of Fajgelbaum et al. (2019) I re-solve in levels. The segmented and full-union counterfactuals are the same coupling program with the transmission bounds set to zero or made slack; the program is conditional on the wages and demand shifters of the current outer iteration, and folding labour- and goods-market clearing into one social-surplus program is left to future work.

two align with a correlation of 0.78 (Figure F.1 in Appendix F). The model also reproduces the cross-country distribution of output shares (a correlation of 0.67 against a value-added proxy).

An out-of-sample stress test then confronts the calibrated model with data from outside its estimation window. Holding every parameter and fundamental at its 2015–2019 calibration, I impose the 2021–2023 energy-crisis environment (crisis-period fossil supply costs, interconnector capacities, and residual electricity load) and compare the solved wholesale prices with observed 2021–2023 country averages, which the calibration never uses. The model reproduces the crisis price geography: the correlation of log prices is 0.73, and the modelled price dispersion is 0.18 against 0.19 observed (Figure F.2 in Appendix F).⁷²

A further cross-check, in the spirit of Arkolakis and Walsh (2023), splits the general-equilibrium wage response to an electricity-price shock into a cost-share incidence term and a model-equilibrium correction. For broad cost shocks, the correction largely offsets the partial-incidence term, isolating the contribution of agglomeration, electricity congestion, and the cross-border network. Appendix F collects the full set of untargeted moments with their diagnostic labels and the numerical verification of the solved equilibrium.⁷³

5.5 Macro elasticity in the model

The quantified model turns the estimated micro elasticity into a model-implied macro elasticity to electricity prices. The location multiplier estimated in Section 4 is a within-location, cross-sector object comparing an electricity-intensive sector with a less intensive one in the same place and year, with location–time effects absorbing every general-equilibrium response. Applying that specification to the calibrated model’s own simulated reallocation returns $1/(\psi - \gamma) \approx 15.8$ exactly, so the estimated response enters the quantification undiluted. The decisive object for the counterfactuals instead is a macro elasticity, and the quantitative model delivers it when accounting for the general equilibrium congestion, firm reallocation, and attached output consequences. Appendix Figure H.8 places the four objects on the common scale of the estimated multiplier: the within-location cross-sector response the Norwegian design identifies, the same sector compared across countries, and the country response measured by firm count and by value added. Each is divided by the relevant electricity cost share so that the four are comparable, and a positive value means production leaves a country when its electricity becomes more expensive.

Once a sector is compared across countries, the wage adjusts to production reallocation, and the firm response is damped in proportion to its sector’s labour cost relative to its electricity cost.

⁷²The fit is in relative prices rather than in levels: modelled prices sit below observed ones in every country and by about a quarter on average, because the fossil-segment wedge κ_j is held at its 2015–2019 calibration, so the crisis solve carries baseline scarcity rents rather than the premia that pushed crisis prices above marginal fuel and carbon cost.

⁷³A further validation exercise is in progress: simulate the Norwegian bidding-zone price shock inside the full general-equilibrium model and ask whether it reproduces the reduced-form elasticity $\hat{\delta}$ in magnitude as well as sign, a test the multi-country equilibrium is not mechanically guaranteed to pass. The headline productivity check uses a labour-productivity proxy, to be checked against sector-level productivity from CompNet.

Electricity-intensive divisions are damped least, because electricity is a large share of what they buy. In the most labour-intensive sectors the wage effect dominates and they move towards countries where electricity becomes more expensive. Across sectors the response runs from about 14 to -24 , and lines up almost exactly with each sector's labour-to-electricity cost ratio (Appendix Figure H.9).⁷⁴

The aggregate firm count at country level hence barely responds, since it represents a mixture of sectors that move in opposite directions. It typically flips sign, with a median of -1.6 , because the divisions that contract sharply are a small fraction of firms while the many small labour-intensive firms that expand are most of them.⁷⁵ The relevant quantity for a macro elasticity instead is the response of total production. A 20% rise in the electricity price a country's firms face cuts the value added of its production sectors by a little over 1% and its electricity use by about 22%.⁷⁶ Divided by the electricity cost share as the other rungs are, this is an elasticity of about nine for the typical country, ranging from six to fifteen across the sample, against the estimated micro elasticity 15.8. That is the macro counterpart of the estimated multiplier: around three-fifths of the micro elasticity across sectors within a country for the typical economy, and in a few small ones almost the same. About 85% of that response is firms entering or leaving rather than continuing firms changing scale, so the aggregate value-added effect inherits the extensive-margin character of the Norwegian estimate in Section 4.⁷⁷

With the parameters disciplined and the validation passed, the model can be asked what the missing single market costs: how much productivity Europe could gain from completing the single electricity market once equalisation also flattens genuine advantages.

6 The value of electricity market integration

The pieces needed to price electricity market fragmentation are now in place: Section 2 showed how much of the dispersion survives trade and how loosely deployment tracks natural resources, Section 4 measured how far production actually moves when electricity costs differ, and Section 5

⁷⁴Each sector's response is the estimated multiplier scaled by $1 + (\alpha_s/\beta_s)\lambda_j$, where λ_j is the proportional wage response to an electricity price change in country j , and this reproduces the solved sector responses. Values of λ_j lie in a narrow band between -0.022 and -0.035 and are stable across a twentyfold change in the size of the shock. The response crosses zero at a labour-to-electricity cost ratio of $1/|\lambda_j|$, between about 28 and 46 across countries, and reverses above it. Appendix B.8 derives both the scaling and λ_j itself from multi-sector labour-market clearing.

⁷⁵Measured on the union counterfactual rather than on a price wedge, the country-level slope decomposes into $+15.8$ from the electricity price, -17.1 from the induced wage response, and -0.04 from the shift in the sector log-sum. The near-cancellation of the first two is the whole of the collapse; aggregating across sectors contributes almost nothing on its own.

⁷⁶Value added here is production-sector revenue net of electricity purchases, measured country by country. It is not the economy-wide $\mathcal{V}$ of Lemma 4, which additionally nets out the imported fuel bill; that term has no country-level counterpart, since the fuel is burned where the electricity is generated rather than where it is used. Netting it out at the European level would lower the elasticities reported here by under 3%.

⁷⁷The elasticity is stable in the size of the shock. Computed at either a 1% country-specific price change or at the 20% change the union itself involves, it moves by less than 5%.

placed that response inside a European economy. This section asks what the dispersion costs, where those losses fall, and how much of them market integration can recover.

Three calibrated quantities drive every number that follows. The first is how far firms move for a given cost advantage, estimated from the Norwegian bidding-zone divergence: firm mass responds to electricity costs with an elasticity of about sixteen, scaled by a sector's electricity cost share. The second is how much a country's own price moves when firms arrive or leave, its supply elasticity, which spans roughly 0.06 to 2.5 and decides whether cheap power survives an influx of industry or is priced away. The third is how movement becomes productivity, through the four channels of Lemma 4: where production ends up, how concentrated it is, how evenly costs are spread, and how much fuel Europe burns to supply it.

A full electricity union removes every internal transmission constraint so that a single price clears the European market, removing both policy and natural sources of price dispersion. Under a common price, electricity drops out of relative location incentives, and generation policy then enters as a state of the electricity system rather than as an alternative to integration, and determines the aggregate fuel savings.⁷⁸ What integration is worth is not a constant: it is set by how much electricity firms buy and how much cheap supply there is to share, which I show by comparing gains between today's baseline and 2040 energy projections. I focus on the effects within manufacturing, where the location channels embody the misallocation this paper is about. Location effects concentrate in the sectors that are most electricity-intensive, while the aggregate effect reflects their output weights and adds the gains from fuel saving.

6.1 Full market integration

The full electricity union is the central counterfactual and the upper bound on integration. What follows reports what a completed single market is worth, which sectors and countries the gain falls on, and how it splits across the four productivity channels of Lemma 4.


Union electricity price. A hypothetical full electricity union equalises wholesale prices fully across countries by removing every internal transmission constraint. The common price $\bar{p}^E = 43.7$ EUR/MWh clears the aggregate market

$$\sum_j E_j^D(\bar{p}^E) = \sum_j \bar{E}_j^0 + \sum_j K_j^F \left[1 - \left(\frac{p^F h_{\min}}{\bar{p}^E - c_0^F} \right)^\zeta \right],$$

It collapses wholesale prices from a 36–53 EUR/MWh cross-country band (figure 7) and eliminates €0.9 billion of annual congestion rents in the baseline.

⁷⁸Whereas Arkolakis and Walsh (2023) endogenise the generation mix along a learning-by-doing transition path, the exercises here are comparative-static: I take the mix as given and quantify the misallocation cost of the resulting price dispersion.

Figure 7: Price compression under the full electricity union



Note: Bars show solved baseline wholesale electricity prices by country, sorted (EU-23, 2015–2019 baseline calibration, EUR/MWh); the dashed line is the common clearing price of 43.7 EUR/MWh when all internal transmission constraints are removed, holding generation capacity fixed.

Productivity effects through reallocation. Within manufacturing, productivity rises by 13.4 basis points, and within the four most electricity-intensive divisions — paper, chemicals, non-metallic minerals, and basic metals — by 35.1. The sector- and division-level figures comprise the three location channels of Lemma 4 — composition, agglomeration, and allocative efficiency — with output weights renormalised within the group. They measure the misallocation of production that price dispersion causes, which is the question this paper asks.

The three channels are moments of the firm-mass response across countries under output weights. Composition captures the covariance of changing mass with productivity, and delivers 8.6 of the 13.4 basis points in manufacturing and 22.4 of the 35.1 in the exposed divisions, as output weight reallocates toward countries where the sector has higher inherent productivity. Agglomeration adds 4.7 and 12.6 respectively, as activity reallocates toward countries that already host the largest clusters.⁷⁹ Allocative efficiency is given by the variance of the cost change and is hence second order, adding 0.05 and 0.2 respectively. The composition share is the same at every level of aggregation, so electricity intensity scales the size of the misallocation without changing its mechanism. What intensity moves is the level: from the remaining manufacturing divisions to the four exposed ones, composition and agglomeration each rise roughly four-fold, while allocative efficiency rises by more than an order of magnitude but from a base too small to matter.

The reallocation behind these numbers is concentrated in the same way. Within manufactur-

⁷⁹Total sector mass is fixed, so the average cluster barely changes size. Of the 4.7 bp in manufacturing, 0.7 is clusters growing at unchanged output weights and 3.6 is output weight moving toward locations whose clusters are already large.

Table 5: Productivity gains and firm reallocation under the full electricity union

<i>Panel A. Productivity gain (basis points)</i>								
Group	μ share	Location channels				Fossil	Total	
		Composition	Agglomeration	Allocative	Total			
Economy-wide	1.00	2.54	1.50	0.02	4.05	1.89	5.94	
Manufacturing	0.40	8.64	4.69	0.05	13.38	—	—	
EI divisions	0.07	22.36	12.57	0.22	35.15	—	—	

<i>Panel B. Firm-mass reallocation (%)</i>				
Group	Max $ \Delta M $	Largest gain	Largest loss	
Economy-wide	0.47	CZ +0.47	IT −0.32	
Manufacturing	4.08	GR +3.28	DK −4.08	
EI divisions	11.93	GR +9.10	DK −11.93	

Note: Panel A reports changes in production-side TFP in basis points relative to the baseline. The μ share is each group's share of aggregate gross output. The manufacturing and electricity-intensive (EI) rows renormalise these weights within each indicated group. Composition, agglomeration, and allocative efficiency are the three location channels in Lemma 4, and their sum is reported under "Total." The fossil-resource saving is an economy-wide level term with no sector index and is therefore reported only for the aggregate row. Panel B reports percentage changes in country-level firm mass relative to the baseline. Max $|\Delta M|$ is the largest absolute change across countries within each group. The final two columns identify the countries with the largest gain and loss. EI divisions are paper, chemicals, non-metallic minerals, and basic metals.

ing, country firm mass changes by up to 4.1%, and within the electricity-intensive divisions the movement is three times larger again: firm mass rises by 9.2% in Greece and falls by 12.0% in Denmark. This is the structural counterpart of the electricity-intensity gradient that identifies the firm response in the Norwegian experiment and in the cross-country panel (Section 4) — reallocation concentrates where electricity is a large share of costs. Appendix figure H.10 traces the underlying gradient: reallocation under the union rises continuously with sector electricity intensity, with the four exposed divisions far from the services cluster.

Reinforcement of clusters. The reallocation consolidates existing agglomerations instead of breaking them up. Flows run from locations that hosted electricity-intensive activity for the power (e.g. 12.0% of exposed-division firm mass in Denmark) toward the countries with the strongest fundamentals and existing clusters, often the expensive ones the union relieves. Under segmentation, hosting a large electricity-intensive cluster raised a country's own price, while the union removes that penalty on incumbency. Concentration accordingly rises in all four highly exposed divisions, and no leading producer loses its lead.⁸⁰ Total firm mass is fixed, so integration cannot make clusters larger on net: some grow, some shrink, and average cluster size only moves by under half a percent. But the clusters that do grow tend to be those that were already larger, so the average firm and unit of output sit in a larger cluster, raising aggregate productivity.⁸¹

Aggregate gains. Economy-wide, the same union raises productivity by 5.9 basis points (table 5).⁸² The location channels account for 4.1 of that. Manufacturing contributes a 5.4 bp gain given its 40% output weight, but the rest of the economy loses 1.3 bp as the union pulls it out of the countries it makes attractive for electricity-intensive production.⁸³ The aggregate reallocation is correspondingly muted: no country's total firm mass changes by more than half a percent, with the Czech Republic the largest gainer (+0.47%) and Italy the largest loser (−0.33%). The remaining 1.9 bp are the fossil-resource saving, an economy-wide level term from a 56 TWh fall in fossil dispatch that belongs to no sector and so has no counterpart in the sectoral figures. Appendix H maps both: the aggregate firm-mass changes are visibly small (figure H.2), while the electricity-intensive reallocation traces the price map.

One property of the integrated equilibrium organises everything that follows. Under a sin-

⁸⁰In basic metals, the output Herfindahl moves from 0.171 to 0.185 and the top-three producer share from 63 to 64%.

⁸¹Of the 4.7 basis points within manufacturing, 0.7 is clusters growing at unchanged output weights, 3.6 is output weight moving from smaller clusters to larger ones, and the remaining 0.4 the interaction of the two. In levels, the cluster the average firm sits in is 0.8% larger under the union and the one the average unit of output sits in 1.2% larger. Output concentrates faster than firm count because incumbents in the growing clusters expand as well.

⁸²The object throughout is production-side TFP rather than welfare: consumer-side and love-of-variety gains are not included.

⁸³Services account for the whole of this loss and slightly more: measured in units of services output they lose 2.3 basis points, which is 1.4 bp weighted by their share of total output, while mining and quarrying adds 0.1 bp back.

gle price, electricity is common to every country and drops out of relative location incentives.⁸⁴ The firm allocation under the union is therefore the same whatever generation Europe has built (Proposition 6).

Generation policy is not a substitute. The other lever Europe debates is generation: build capacity where natural advantage in renewable generation is highest, and let the cheap power justify itself. It fails without integration and could at best act on the ninth of exploitable dispersion that generation siting explains: re-siting the existing renewable stock according to natural potential on the segmented grid lowers manufacturing productivity by 6.2 basis points. The cheap power strands behind transmission bottlenecks, usable only by the firms willing to relocate to it. The same re-siting on top of the union leaves the firm allocation unchanged (by the invariance property above) and moves the economy-wide figure by less than a tenth of a basis point, through the fuel bill alone. Generation policy is a cost and emissions lever, but where production locates is decided by market design.

Robustness. The headline aggregate union gain varies between 4.0 and 8.2 basis points across the calibration sensitivity ranges, with the agglomeration parameter γ the dominant source of parameter uncertainty. Figure F.3 in Appendix F collects the ranges.

Beyond the parameter ranges, a handful of specification choices are not settled by the data. First, a lower elasticity of substitution of electricity against labour instead of the imposed unit elasticity lowers the economy-wide gain because inelastic electricity demand shrinks the fuel-bill saving, but raises the manufacturing gain. The others concern the fossil supply side, all affecting the fuel-saving channel instead of production allocation. Assuming that fossil fleets in countries that rely less on fossils are on average more efficient moves the economy-wide gain by nine percent, entirely through the fuel-saving channel, and leaves the firm allocation unchanged.⁸⁵ Allowing different slopes by fuel type leaves the aggregate gain essentially unchanged, as does the placement of the country-level cost wedge, which bears only on how fuel and carbon prices pass through. Allowing the union to draw on renewable output curtailed under segmentation raises the gain slightly. Appendix G reports details.

A few modelling choices bound what any counterfactual can capture. Total firm mass is held fixed, so production is reallocated within Europe but never drawn into or out of it.⁸⁶ Integration itself does not operate on that margin, since it equalises prices without lowering them significantly.

⁸⁴The price level cannot restore them either, because the sectoral price index moves one-for-one with the common electricity cost.

⁸⁵Precisely, this different treatment of the fossil supply curve would assume countries with small fossil shares have a truncated electricity supply curve instead of one drawn from the full Pareto distribution of efficiencies, making their supply curves significantly flatter.

⁸⁶The 2040 build-out of Section 6.3, which does lower the European price level, would in practice also raise the world market share of European firms, a margin the reallocation exercise cannot see.

The framework also treats intermediate inputs as priced identically across countries, so firms are exposed to electricity prices only through their own purchases rather than through electricity-intensive suppliers, which understates exposure in the sectors furthest downstream. Third, the model treats the EU as a closed electricity system with imports from non-EU neighbours held fixed, and relaxing it would attenuate price dispersion on the periphery. Fourth, the model imposes restrictive wage congestion, with labour immobile across countries and supplied inelastically although freely mobile between sectors, so wages absorb the adjustment to firm entry. Non-firm electricity demand, about half of load, is likewise held fixed rather than price-responsive, so adjustments falls on electricity prices more strongly.

Unlike Arkolakis and Walsh (2023), who model the physical grid with quadratic transmission losses and locational marginal pricing, I abstract from line losses and represent cross-border exchange through NTC-bounded market coupling. This abstraction is institutionally aligned with the European zonal market where prices are set at the zone rather than the node and line losses are second-order relative to the price wedges that NTC limits generate when interconnectors bind. Appendix G.5 shows the headline counterfactual moves by 14% when assuming 3% losses across all connectors.

Lastly, the counterfactual equalises wholesale prices, while firms pay retail prices that add network charges and levies on top, and within a sector the dispersion of what firms actually pay across countries is about a third larger than the dispersion of wholesale prices. Re-solving the model with those wedges in place leaves a wholesale union worth about two-thirds of the headline within the electricity-intensive divisions, while removing wholesale and retail dispersion together is worth more than the headline itself, the two acting as complements rather than substitutes. Appendix G discusses this in more detail, together with fuller treatments of the sensitivities above.

6.2 What determines the value of integration

A union does not automatically make electricity cheaper on average, it makes it uniform. Prices fall in the countries above the union clearing price and rise in those below it, so the union delivers a spatial effect rather than a cost effect. This section shows what determines the size of that effect, by discussing the features of the model economy the headline gain is a product of, and how they would have to change to obtain a different result.

How much of the value of a union depends on firms actually moving? Freezing every firm at its baseline location and re-solving the union equalises prices, allowing incumbents in newly-cheap countries to scale up, while no production relocates. This immobile union is worth 2.7 basis

points economy-wide and 2.3 bp within manufacturing.⁸⁷ The aggregate gain is dominated by the fuel saving of the merged merit order (1.7 bp) — the object existing estimates of electricity market integration measure — with an in-place scale adjustment of incumbent firms (0.9 bp).⁸⁸ Reallocation contributes the remaining 3.3 bp economy-wide and over 80% of the manufacturing gain, and it is the margin the Norwegian experiment disciplines. Most of what integration is worth to the production economy runs through the production that moves. What follows is about that margin. The fuel-saving channel follows a simpler mechanism and is discussed with its environmental consequences in Section 7.4.⁸⁹

The reallocation gain from the union is the product of three main factors: how big the cost impulse of a combined supply is, how far firms move for a given cost advantage, and whether that movement is aligned with productivity advantages. First, the cost impulse is determined by a sector's electricity intensity times the price dispersion integration removes, including the wage adjustment. Second, the location multiplier turns that cost difference into firm reallocation, dampened by equilibrium adjustment of prices and wages. Third, alignment asks whether the countries a union makes cheaper are more productive, either inherently as measured by the composition channel, or because they already carry large clusters with agglomeration externalities. Under the union both is true, such that alignment determines the positive sign of the productivity impact scaled by the other two factors.

Cost impulse: electricity exposure $\times$ exploitable dispersion. The cost impulse depends crucially on exposure to electricity prices, which separates the manufacturing number from the economy-wide one. Electricity is about one percent of gross output for the average sector but several times that in the exposed divisions (Section 2), and this gradient sorts production. Each sector responds to integration according to its electricity share set against the labour share through which the equilibrium wage response operates. Electricity-intensive sectors move

⁸⁷The full union delivers 5.94 basis points economy-wide against 2.67 bp without sorting, so reallocation contributes 3.27 bp (55%); within manufacturing, 13.38 bp against 2.29 bp, so reallocation contributes 11.09 bp (83%); within the four exposed divisions, 35.15 bp against 5.09 bp (86%). In the immobile union no cluster changes size, but output weight still shifts toward the countries whose costs fall, which are on average the larger clusters, so the agglomeration channel is 0.3 bp rather than zero (0.6 bp within manufacturing). Allocative efficiency is negligible (0.03 bp). The fossil saving is 1.72 bp against the full union's 1.89 bp, since dispatch differs when firm electricity demand cannot move across countries.

⁸⁸An in-place response might seem at odds with the muted Norwegian intensive margin (Section 4). The Norwegian estimate measures continuing firms' profits against a crisis-period wholesale price under hedging, whereas the scale adjustment here is the long-run response of incumbent output to a permanent cost change under CES demand. The cross-country panel prices this margin (Table 3), which is negative, significant, and steeper than the Norwegian crisis response, consistent with contracts repricing to long-lived cost differences. At 0.9 basis points, the in-place layer is small and dominated by relocation.

⁸⁹The fuel gain scales with the fossil generation a unified merit order displaces, valued at the fuel content of the wholesale price since domestic generation rents are transfers within the economy. It is largest where fossil plants set the price and fuel is a large share of it: 1.9 of the 5.9 basis points in the calm baseline, 3.5 of 9.5 bp during the energy crisis, and growing with the carbon price. It shrinks towards zero as the renewable build-out displaces fossil generation on its own, which is why the union's value becomes entirely spatial in the 2040 economy (Section 6.3).

toward the countries the electricity union relieves and labour-intensive ones move away, which is the endogenous specialisation of Proposition 5 (Appendix B.7). Exposure is also the margin on which the union gain grows, which Section 6.3 takes up.

The other component of the cost impulse is the exploitable price dispersion removed by the union. The price span across countries removed by the union is 36–53 EUR/MWh, but what moves production is the demand-weighted average gap across countries, which is just 10.5%. Dispersion sitting in countries with small total load cannot be arbitrated since there is no quantity to ship and little production to attract.⁹⁰

This exploitable price dispersion is a product of natural advantages, policy choices in generation, and market integration. The policy content of Europe’s dispersion accordingly lies less in the generation policy choices of where renewables were built than in the lack of interconnections. Geography decides which countries are cheap, and market segmentation allows differences to persist as a price gap. Removing the generation policy distortion of renewable capacity allocations not following natural potential only accounts for roughly a ninth of the price dispersion. Reallocating renewable generation capital according to natural potential leaves demand-weighted dispersion of 9.3% against 10.5% observed.⁹¹ This is the quantitative counterpart of Fact 2 (Section 2.2), and the reason the union, which harvests dispersion regardless of its source, dominates any lever that acts on the generation policy component alone.⁹²

Reallocation: location multiplier $\div$ equilibrium dampening. How the cost impulse translates into production reallocation is determined by the location multiplier identified in the empirics and the general equilibrium dampening. A Europe whose firms are less attached to a location (ψ), or whose agglomeration economies are stronger (γ), would gain more from integrating its electricity market. Holding today’s productivity geography fixed and treating attachment as a comparative static, firms twenty percent less attached to their locations would re-sort about 1% of aggregate revenue weight under the union in place of 0.7%, raising the aggregate location channels by almost half, from 4.1 to about 5.9 basis points. Twenty percent stronger agglomeration, which amplifies whatever re-sorting attachment permits, delivers about a third as much.⁹³

⁹⁰The energy crisis makes the distinction concrete: it raised the raw cross-country price span by a factor of 4.7 but exploitable dispersion only by 1.7, and the productivity channels tracked the latter (Section 7).

⁹¹The split rests on the demand weighting, which is the economically relevant metric here: the natural-potential allocation narrows dispersion where demand sits but widens it across small countries, so the unweighted share is negative. Weighted variants deliver 11–13%. Weights are country generation shares, the convention behind the published dispersion figures throughout.

⁹²The most prominent generation policy choice of the period makes the same point. Restoring Germany’s 2010 nuclear fleet, holding everything else fixed, raises European productivity by 0.75 basis points, but 0.69 bp of that is a fuel bill and only 0.06 bp runs through location. Even that residual is close to a pure transfer: Germany’s contribution to the location channels rises by 1.58 bp against –1.52 bp for the other twenty-two countries. Appendix H.6 reports the exercise.

⁹³How far either comparative static can be pushed is bounded, as agglomeration must stay below attachment for the location problem to remain well defined.

This is why the paper estimates the location multiplier rather than importing it. Attachment carries the greater leverage of the two, and it is the one the Norwegian bidding-zone divergence identifies: pinning ψ given γ and propagating the standard error of $\hat{\delta}_M$ through the calibration bounds the aggregate gain within [5.8, 6.3] basis points. Varying agglomeration from zero to double its value bounds the gain between 4.0 and 8.2.⁹⁴ Attachment moves the answer three times as much, and the experiment reduces its contribution to half a basis point, leaving agglomeration as the source of parameter uncertainty with a third of the leverage.

The general equilibrium response dampens production reallocation significantly. As firms move toward cheap power they increase local wages, and that wage response absorbs more than half of the gross reallocation impulse. Wages move only about three percent as much as electricity prices, but labour is roughly a fifth of gross output against electricity's one percent, so the induced wage change is a first-order part of the cost change firms respond to. Across sectors, firms respond according to their electricity intensity net of the wage response, and for roughly half of output weight the wage is the larger factor. Splitting the composition channel into its two costs gives an electricity term of 5.9 basis points against a wage term of -3.4 , netting to the 2.5 total.⁹⁵ Electricity congestion is the other dampening force, as more firm demand pushes up country prices according to the national supply elasticity under segmentation. Under a union, total demand determines the price and electricity congestion no longer discriminates between locations.

Productivity alignment. The third factor determines whether the resulting production allocation gains or loses aggregate productivity. With segmented markets, countries that host electricity-intensive clusters are penalised through higher national electricity prices. This penalty is removed in the union with the electricity congestion force shared. A positive impact of reallocation is not guaranteed, but determined and scaled by the correlation of cost changes with the two productivity components.⁹⁶ The countries that were more expensive in Europe are on average both inherently more productive and the ones already carrying the larger clusters, which is why the composition and agglomeration channels are both positive (Section 6.1).

Comparing magnitudes. Read as an economy-wide number, 6 basis points is the right order of magnitude rather than a disappointment. Electricity is about one percent of what a firm spends, and integration moves that price by roughly a tenth in either direction, so what firm activity moves for is on the order of basis points of their cost base. Production moves a little in response, and it pays off only where it moves somewhere more productive.

⁹⁴When varying external γ on a sensitivity band, estimated ψ automatically adjusts to keep the multiplier stable.

⁹⁵The split uses the location first-order condition $\alpha_s \Delta \ln w_j + \beta_s \Delta \ln p_j^E = -(\psi - \gamma) \Delta \ln M_{js}$, which reproduces the solved composition to within three percent.

⁹⁶For example, the announced 2040 build-out delivers more dispersion at the same exposure and the same multiplier, and costs manufacturing 6.8 basis points instead of a 13.4 bp gain, because capacity built away from load makes those countries cheaper than have neither the fundamentals nor the established clusters (Section 6.3).

The closest analogues in the spatial-misallocation literature sit exactly there, finding gains below a percentage point.⁹⁷ The much larger figures in the electricity-policy debate are different objects rather than different answers.⁹⁸ The relevant question is therefore not why the economy-wide number is small, but how it might change in the future with decarbonisation, which is explored in detail next.

6.3 Integration gains scale going forward

Both factors determining the cost impulse of integration move over the next fifteen years. Exposure rises as industry electrifies, and exploitable dispersion rises as Europe builds a large and unevenly distributed low-cost generation capacity. I price both projections separately and jointly, holding the location multiplier and the factor-market responses at their calibrated values throughout. The question is what a given attachment of firms to locations implies when the electricity system changes around them. The value of completing integration rises from 13.4 basis points of manufacturing productivity today to 24.1 bp in the 2040 economy at current cost shares, and to 34.8 bp once industry has electrified — close to a tripling, and from 35.1 bp to 86.1 bp within the exposed divisions.

Low-cost electricity supply. Exploitable dispersion is bounded by how much low-cost supply exists relative to demand and by where it sits. Adding cheap generation does not mechanically raise the value of integration, as its value depends on whether the new supply lands where demand is or the wires exist to move it towards demand. Under segmentation, capacity built away from demand creates price gaps that cannot be arbitrated, and the cheap power is used only by whatever production is local or willing to move towards it. Under a common price the same capacity is shared, not creating additional dispersion but decreasing the union price.

Countries following their announced national supply trajectories without integration generate an uncoordinated pattern. Taking projected renewable capacity from ENTSO-E’s TYNDP 2024 scenarios and national energy and climate plans at the 2040 horizon, I solve for the equilibrium prices,

⁹⁷Harmonising US state taxes — a wedge several times one percent of costs — raises welfare by 0.6% (Fajgelbaum et al., 2019), and the housing-constraint gains of Hsieh and Moretti (2019) fall to basis points once corrected (Greaney, 2026). Large numbers in this literature require large dispersion on inputs with large cost shares: Hsieh and Klenow (2009) obtain 30–60% from TFPR dispersion of half a log point or more on the entire input bundle.

⁹⁸The European Commission’s estimate of €12.5–40 billion per year from energy market integration bundles price-differential-times-volume transfers with coordinated investment; the market-coupling welfare component is below €4 billion per year, and the cost of abandoning coupling altogether is 0.25–0.5% of wholesale market value (Newbery, Strbac and Viehoff, 2016). Estimates of dispatch-efficiency gains, such as the 5% cost saving from market-based dispatch in Cicala (2022), are measured against the cost base of generation itself — a large share of a small base.

firm distribution, and productivity implied by the announced plans.⁹⁹ The build-out is large and uneven: EU renewable generation rises more than fivefold, with ratios ranging from 2 to over 30 across countries (Appendix figure H.6).¹⁰⁰ While today's exploitable dispersion is mostly determined by demand and natural potential geography behind unconnected borders (Section 6.2), the uncoordinated national choices of the build-out create additional policy-driven dispersion.

The consequences for the 2040 economy are large. Prices fall from a 36–53 to a 15–39 EUR/MWh band, firm electricity demand expands from 1,436 to 2,573 TWh, and modelled fossil generation goes to zero. The economy-wide impact is positive due to large fuel savings, with 22 basis points from displaced fossil, while location channels are negative (–3.9 bp).¹⁰¹ Manufacturing loses 6.8 basis points of productivity, and the exposed divisions lose 18.4 bp. The loss is mostly a loss of concentration, as the new dispersion pulls electricity-intensive production toward whichever countries build out most and partially breaks up established clusters.¹⁰² Both figures already include the planned grid expansion, which softens the loss within the exposed divisions but leaves the manufacturing total essentially unchanged.¹⁰³ While integration concentrates production, an uncoordinated build-out disperses it.

Integration reverses the location impact of the build-out. Integrating the 2040 system is worth 24.1 basis points of manufacturing productivity against 13.4 today, and 63.4 in the most exposed divisions against 35.1. The aggregate location channels also become positive for a combined gain of 26.2 bp, so the economy-wide gain from integration is 8.0 bp.¹⁰⁴ The renewable expansion

⁹⁹Appendix H.2 documents the projected inputs. These are full 2040 scenarios rather than capacity changes in today's economy. Scenarios run on the TYNDP-2040 grid of the planned-integration exercise. Residual demand is rebuilt from the scenario's own sector detail — rail carved out of transport, tertiary out of buildings, losses at the total-load factor — and then held fixed, while firm demand remains price-elastic and responds to the solved equilibrium. Firm electricity expands from 1,436 to 2,573 TWh, so total load reaches 4,704 TWh against a staged anchor of 3,845 TWh evaluated at baseline prices. I also include the carbon price from the same set of projections, up to 147 EUR/tCO₂ in 2040.

¹⁰⁰Any differences reported use the calibrated baseline economy as base, with the averaged 2015-2019 window.

¹⁰¹The fuel saving is identical in every 2040 variant, integrated or not, because the announced build-out is large enough to displace modelled fossil generation almost entirely on its own, leaving nothing for market design to displace. It is therefore a stock that the transition gains, rather than a flow benefit, as beyond 2040 there is no further fuel to displace. What remains after that is the location channel, which is the part integration governs and the part that grows with exposure.

¹⁰²The build-out de-concentrates all four exposed divisions — in basic metals, the output Herfindahl falls from 0.171 to 0.159 on the planned grid, and to 0.150 on today's as production scatters toward whichever countries build fastest.

¹⁰³On today's grid the same build-out costs manufacturing 6.7 basis points and the exposed divisions 20.9, against 6.8 and 18.4 with the planned expansion in place. Transmission acts on both location channels at once: the expansion repairs 4.1 bp of the manufacturing cluster damage, but compressing dispersion also leaves less production moving toward better fundamentals, and the two nearly cancel in the manufacturing total. Within the exposed divisions the repair dominates.

¹⁰⁴The union values are the invariant ones of Section 6.1, since its location effect does not depend on what has been built. The 2040 exercises are solved on two projections of non-firm electricity demand that are otherwise identical — the same TYNDP-2040 grid, the same announced build-out, and the same 2040 carbon price. The demand and price-band figures above and the within-manufacturing union gains use the Distributed Energy path; the productivity effects of the uncoordinated build-out, and the economy-wide increment below, use the earlier scenario-midpoint path. Every comparison reported is internally consistent on its own path, and the choice is not material: measured on the midpoint path the union's 2040 manufacturing gain is 20.2 basis points in place of 24.1, and no sign or ordering changes.

lowers the union price from 43.7 to 20.8 EUR/MWh, exploiting larger low-cost quantity in the cheap tails and demand centres. The build-out multiplies the dispersion at stake, and integration determines whether that dispersion is damaging or harvested.

Electrification of industry. The second margin is exposure. The counterfactuals above hold sector electricity cost shares at their calibration vintage, but industry is expected to electrify substantially by 2040. I take sector-level electrification from the same TYNDP 2024 scenario family the generation and demand projections come from.¹⁰⁵ Electrification is uneven across sectors, so it widens the cost-share gradient rather than shifting it (see Appendix Figure H.1). The ratio of electricity intensity in the exposed divisions to services rises from 4.4 to 5.8, with the largest increase in electricity-intensive chemicals from 3% to 4.9% and most services staying constant.¹⁰⁶


The union gain at projected electricity cost shares rises for manufacturing from 13.4 to 19.4 basis points and the exposed divisions from 35.1 bp to 49.6 bp. Manufacturing's own cost share rises by a factor of 1.32 and its union gain by 1.44. At the aggregate level, projected electrification makes the 2040 price band both higher and wider, 18–47 against 15–39 EUR/MWh, because more electricity-intensive production bids up the price where power is cheap. The union therefore becomes more valuable at 8.0 basis points economy-wide compared to 5.9 bp. Almost all of the increase is spatial, as the location channels rise by the same factor as within manufacturing, while the fossil-saving term increases moderately by 15%.¹⁰⁷

¹⁰⁵I use the Distributed Energy scenario, which is the more electrified of the two published paths, so the figures are an upper bound on this margin. Cost shares are scaled by the projected ratio of electricity intensity of activity, holding relative prices constant, so the exercise moves technology rather than prices. Residual load is rebuilt from the same scenario's sector detail so that industrial electrification is not counted twice. Appendix H.3 documents the construction.

¹⁰⁶Market services other than information and communication are held at their current intensity throughout, because the projection source carries no output denominator for them, so the exercise understates economy-wide electrification. Information and communication (J) is the one service sector the projection reaches, through TYNDP's datacentre demand. TYNDP files datacentres under industry, but they belong to J rather than to residual manufacturing here: the model's residual-manufacturing electricity is built from manufacturing energy balances that carry no datacentre load, and it matches the TYNDP basket only once datacentres are excluded, while TYNDP's datacentre demand alone is the size of the modelled J sector. J's projected intensity falls to 0.73 of its baseline share. Datacentre electricity grows by a factor of 2.4 but the corresponding activity index by 3.4, so electricity per unit of activity declines. The projected datacentre boom is therefore activity growth rather than electrification, and the activity margin lies outside an exercise that holds the economy at its baseline size. Where datacentres locate within Europe is a natural extension.

¹⁰⁷The economy-wide figure is not a clean exposure elasticity. The output-weighted electricity cost share of the modelled economy rises by a factor of 1.19, well below the 1.44 by which its location channels rise, because electrification concentrates in the sectors that relocate most and the location gain loads on the widened gradient rather than on the average. The exercise also moves electricity between blocks as well as raising intensity within them: modelled-sector electricity rises from 1,436 to 1,708 TWh while residual load falls from 1,408 to 1,135 TWh in line with projections, leaving total European load unchanged to within 0.1% but raising the model's mobile share of it from 50 to 60%. That reallocation is the no-double-counting adjustment described above and is intended, but it means part of the aggregate increase is demand entering the priced and mobile block rather than industry becoming more electricity-intensive relative to the rest of the economy. The within-manufacturing figures are free of this, since the within-group normalisation excludes the residual.

Figure 8: Union gain under renewable electrification by 2040



Note: Each bar is the union gain in basis points of production-side TFP within manufacturing, in a different economy. The 2040 build-out uses national projections for renewable generation and carbon prices, electrification includes projected electricity cost shares, and the interaction captures the residual. The last bar combines the different projections into a union gain under the full 2040-projected economy.

Both margins: renewable-powered electrification. Applied jointly, the two margins multiply rather than add. Economy-wide, electrification alone is worth 2.1 basis points and the 2040 build-out alone 3.5 bp, but together they are worth 7.8 bp — an interaction of 2.2 bp, 16% of the total. Exposure raises the stakes in a decarbonised and spatially uneven system, and the renewable transition is worth more when production is electricity-intensive. Integration is what converts the build-out from a location dispersion cost with fuel savings into a common gain in all channels.

The composition of the gain also changes with the transition, such that the economy-wide and manufacturing gains converge as decarbonisation proceeds. At baseline, a third of the union's economy-wide gain is fossil-fuel saving, while in the 2040 economy the build-out drives fossil generation to zero, so the union's entire value is spatial reallocation. The negative effect of the electricity union on services falls similarly, from 2.3 basis points at baseline to 0.4 bp in 2040.

The value of integration is therefore not a fixed quantity to be weighed once against the cost of building interconnectors. Instead the value rises with renewable electrification, most in exactly the sectors European industrial policy is most concerned to keep. The rationale shifts from burning less gas for generation to locating production in the most productive places. The misallocation this paper measures does not fade as Europe decarbonises, it becomes the essence of what integration delivers. Given the long timelines involved in building both transmission and generation until activation, the decision to integrate has to be taken together with the generation policy that in isolation otherwise creates further dispersion.

7 Integration in the policy debate

The counterfactuals so far price integration in terms of the central outcome of this paper: aggregate and sectoral productivity, through the location of production. European electricity policy, however, is prominently debated for other concerns — competitiveness relative to the United States, price stability with exposure to fuel markets, and decarbonisation. This section confronts the results with each in turn: how far Europe’s internal fragmentation stands from the US benchmark and what closing that distance would deliver; whether integration is worth more precisely when energy is expensive; and whether a merged market cuts emissions or merely relocates them. Together they determine where the main integration results stand in the debate that will decide whether interconnection is expanded further.

7.1 How far is Europe from the United States?

Europe’s productivity shortfall relative to the United States is a central concern of European policy, and energy costs figure prominently in the diagnosis: EU firms face electricity prices two to three times American levels (Draghi, 2024). That is a statement about the level of European power prices, and it is not what this paper measures.¹⁰⁸ European electricity is both more expensive on average than American electricity and markets are more fragmented across Europe, and the exercises here price the unevenness in interconnection alone. Integration does not act on the price level: as Section 6.2 showed, a union does not make electricity cheaper but uniform, and the common price of 43.7 EUR/MWh in fact clears slightly above the demand-weighted baseline of 42.7. What follows therefore asks how far Europe’s internal fragmentation sits from the American benchmark and what closing that distance would deliver — not how much of the transatlantic price gap integration could remove, which is none of it.

The United States provides an empirical benchmark for the degree of integration. Relative to load, EU cross-border transfer capacity is lower than transfer capability across US regions, with an annual import-capability ratio of 0.23 against 0.40. It is also far more unevenly distributed, with a cross-country coefficient of variation of 0.88 against 0.40.¹⁰⁹ The US-equivalent integration scenario expands interconnector capacity bounds τ_{ij} , preserving each country’s import-portfolio shares, until the cross-country distribution of capacity relative to load reaches the US level and

¹⁰⁸The report focuses on retail prices, which adds a large wedge between US and EU price levels. In terms of wholesale prices only, US prices on average were about 25% higher than in the EU during the baseline sample of 2015–2019.

¹⁰⁹The US benchmark is constructed from NERC simultaneous import capability (North American Electric Reliability Corporation, 2024) for MISO and PJM, the two largest US wholesale markets, which the NERC assessment divides into seven subregions (MISO West, Central, East, and South; PJM West, East, and South). These subregions are the US analogue of the model’s 23 countries: each is a separately assessed transmission area with an annual load roughly the size of a mid-sized to large European country. Import capability is converted to an annual capacity-to-load equivalent; the EU side sums directional transfer capacities border by border, so the scenario should be read as the model-consistent annual equivalent of the US integration level.

approaches its dispersion.¹¹⁰

US-equivalent integration compresses the cross-country price span from 17.8 to 11.3 EUR/MWh. It raises manufacturing productivity by 3.3 basis points and economy-wide it delivers 2.1, about a third of the full union. It is the evenness of interconnection relieving binding bottlenecks rather than its quantity that carries the gain. Scaling every border uniformly to the US aggregate level without compressing the cross-country distribution delivers no gains.

The same pattern holds for the transmission Europe actually plans to build. Expanding each interconnection τ_{ij} to its projected TYNDP 2040 value, holding generation and demand at baseline, compresses the price span even more, from 17.8 to 7.8 EUR/MWh (Appendix figure H.4). However, the remaining bottlenecks bind harder rather than less: congestion rents rise from €0.9 to €1.0 billion annually as larger flows press on the constraints that survive, concentrated on the French borders with Italy and Spain (Appendix figure H.3). Planned integration consequently raises aggregate productivity, but by less than the US-equivalent integration: the economy-wide figure is 1.5 basis points, and on today's generation mix manufacturing does not gain at all, losing 0.2. Neither benchmark reaches half of what full integration is worth.¹¹¹

7.2 How much productivity can planned connection investment gain?

The productivity gains from electricity-market integration must be weighed against the cost of the transmission capacity required to achieve them. Starting from the current grid, the lowest-cost implementation of the electricity union requires 84.3 GW of additional interconnection capacity at an estimated €120.7 billion in upfront investment. Under a 25-year horizon and a 4 percent discount rate, the resulting benefits have a net present value of €99.2 billion when productivity gains within manufacturing are combined with fuel-cost savings.¹¹² In the projected 2040 economy, completing integration from the planned TYNDP grid requires a further €106.3 billion and generates corresponding net present values of approximately €332 billion and €328 billion.¹¹³

A positive return to complete integration does not imply that the planned portfolio allocates a limited investment budget efficiently. I therefore consider whether the funds required to expand the current grid to the TYNDP 2040 configuration could produce larger productivity gains if allocated differently. The estimated upfront cost of the TYNDP expansion is €194.2 billion, while the union would require XXX. I hold this expenditure fixed and evaluate portfolios by their effect on manufacturing TFP, while requiring them to weakly reduce the model-implied average electricity price and not to lower implied aggregate TFP gains.

¹¹⁰The dispersion cannot be met fully without reducing capacity across some links, so the resulting EU coefficient of variation reaches only 0.61.

¹¹¹The planned grid's return in the 2040 economy sits in the most exposed divisions rather than in manufacturing as a whole: in the build-out economy of Section 6.3, the same expansion repairs 2.5 of the 20.9 bp of damage that uncoordinated build-out inflicts on the exposed divisions, while leaving the manufacturing total roughly unchanged.

¹¹²NPV is €58.0 billion when the broader aggregate productivity measure is used instead.

¹¹³Appendix G.11 presents the construction and assumptions underlying these estimates.

One such portfolio concentrates additional investment on Italy, whose electricity market remains segmented through several binding interconnection constraints. The portfolio adds 2GW of transmission capacity on each Italian border (with Austria, France, Greece, and Slovenia) at the expense of a uniform reduction in planned interconnection investment on all other borders by 9.6%.¹¹⁴ This investment portfolio substantially outperforms the TFP impact of planned integration while adhering to the constraints. Relative to the planned grid, it raises manufacturing TFP by 30.2 basis points while reducing the capacity-weighted electricity-price index by €0.26 per MWh and further increasing aggregate TFP gains.¹¹⁵ The manufacturing gain relative to the TYNDP grid is equivalent to 85.6 percent of the gain from complete price integration. The result therefore shows that the economic allocation of interconnection investment can matter as much as its overall scale: a feasible reallocation of already planned expenditure can preserve the portfolio's aggregate and price benefits while delivering considerably larger productivity gains. Appendix H.4 provides the full construction and sensitivity analysis.

7.3 Does integration insure against an energy crisis?

The counterfactuals above are evaluated at calm 2015–2019 prices. When prices spike, as during the 2021–2023 energy crisis, does integration become more valuable? I confront the calm economy with the fuel-price shock of the crisis, including its carbon costs, holding each policy's interconnection capacity in place as if it had been adopted before the shock hit. The resulting equilibrium reproduces the observed 2021–2023 price geography well out of sample (Section 5.4), with a large effect on price dispersion, raising the cross-country span from 17.8 to 83.6 EUR/MWh (table H.3). Firm responses are lower bounds, since the model assumes a long-run unit elasticity of substitution.

Integration is worth more precisely when energy is expensive. Without additional integration, the shock costs manufacturing 3.4 basis points, so a union would have been worth 16.8 bp of manufacturing productivity during the crisis, compared against 13.4 bp in calm conditions. The composition and agglomeration channels scale with demand-weighted price dispersion, which the crisis raises by roughly 70%. This is far less than the raw dispersion, because a large part of the crisis dispersion sits in a cheap tail with little demand behind it, exactly the kind of dispersion integration cannot arbitrage.

On aggregate, crisis insurance has higher value once fuel savings are included. The fossil-resource channel scales with the price level directly, and the union's fuel saving roughly doubles, from 1.9 basis points during calm to 3.5 bp during the crisis. Together with location channels of

¹¹⁴This is not presented as the globally optimal network. It is a deliberately restricted and readily interpretable counterfactual demonstrating what can be achieved through a targeted reallocation of the planned investment budget.

¹¹⁵Relative to the current grid in the projected 2040 economy, manufacturing TFP rises by 28.9 basis points under the Italy-focused portfolio, compared with a decline of 1.3 basis points under the TYNDP portfolio, while aggregate TFP rises by 11.6 rather than 4.1 basis points.

6.0 bp, integration is worth 9.5 bp economy-wide against 5.9 in calm times.¹¹⁶¹¹⁷

Expanding renewable capacity earlier without integration leads to aggregate productivity gains that are entirely fuel saving, while location channels turn negative and manufacturing loses 0.6 basis points. The divergence is the crisis-state version of the result in Section 6.3: building capacity insures the fuel bill, while only integration insures the geography of production.

7.4 What is the environmental impact of integration?

The main outcome of interest is the gain in productivity through location effects, but any level of integration also carries an environmental impact through aggregate fuel savings. Under a union the carbon price does not touch the location channels: manufacturing gains 13.0 basis points whether allowances trade at the baseline price or at today's, by the same invariance property that makes generation policy a cost lever rather than a location lever (Proposition 6). What the carbon price governs is the fuel bill and the emissions that come with it, so this is a question about the aggregate alone. Improved market coupling lowers aggregate fossil dispatch, but whether lower fossil generation also implies less CO₂ depends on the fuel composition of fossil dispatch, which depends on the carbon price.


At the baseline-period carbon price (11.9 EUR/tCO₂, the 2015–2019 EU ETS average), the full union raises emissions slightly, by 9.4 MtCO₂ (+1.2%) against baseline emissions of 772 MtCO₂ from modelled fossil generation. Fossil dispatch falls by 56 TWh, but the remaining dispatch concentrates in Germany's comparatively carbon-intensive fleet (+97 TWh) while falling in Italy, Spain, and Greece (–142, –66, and –30 TWh), because a low carbon price makes coal the low-cost margin.

A modestly higher carbon price flips the sign: the union delivers a CO₂ dividend above a crossover allowance price of roughly 13 EUR/tCO₂ (figure 9), far below the 75–80 EUR/tCO₂ at which EU allowances currently trade. Evaluated at today's allowance price, the union cuts emissions by 320 MtCO₂ (49% of the corresponding carbon-price baseline) and its aggregate produc-

¹¹⁶Each shock variant is solved against its own matched crisis baseline; increments are differences from the shock-alone equilibrium, which itself sits 3.2 basis points above the calm baseline economy-wide through the revaluation of the fossil-resource channel at crisis fuel prices. The economy-wide insurance ranking is build-out +13.7, union +9.5, US-level interconnection +5.2, planned grid +3.5, re-siting –2.6.

¹¹⁷A full recalibration corroborates the shock-overlay results: re-anchoring the electricity supply curves to 2021–2023 fuel and carbon costs and re-solving the baseline and the full union in that state doubles the union's gain — 19.7 bp in manufacturing against its own crisis baseline, versus 13.4 bp in calm times, 12.0 versus 5.9 bp on aggregate — with the union clearing price at 148 EUR/MWh in place of 43.7. The two exercises differ in what is held fixed, which is why their location channels differ. The overlay keeps the calm baseline economy and changes only fuel and carbon costs, so by Proposition 6 the union's location channels are identical to those in calm times; the recalibration re-anchors the supply curves and hence the baseline economy itself, against which the union is then measured, so the invariance no longer applies. The recalibrated crisis baseline clears the solver's relaxed convergence tolerance rather than the strict standard met by the other headline stems.

Figure 9: The carbon-price crossover: emissions and TFP effects of the union



Note: Change in direct CO₂ emissions from modelled fossil generation (left axis) and aggregate TFP gain (right axis) under the full electricity union, as a function of the EU ETS allowance price. Each point is a solved counterfactual depending on the carbon price, up to the TYNDP 2040 scenario (147 EUR/tCO₂). Each emissions change is measured against the 2015–2019 baseline carbon price (11.9 EUR/tCO₂). The emissions change crosses zero at an allowance price of 13.1 EUR/tCO₂ marked vertically. Emissions are split by fuel type which are affected differently based on their carbon intensity.

tivity gain rises from 5.9 to 14.7 basis points through the fossil-resource-saving channel.¹¹⁸ Integration and carbon pricing are complements: the wires determine where generation can move, the carbon price determines whether it moves toward clean or cheap generation. Appendix figure H.7 maps country emission levels under the baseline and the coordinated scenarios.

These figures are a power-sector accounting. They measure the change in direct combustion CO₂ from modelled fossil generation, holding the supply of EU emission allowances fixed outside the model, and so describe how integration changes the emissions the power sector produces rather than the European total, which the EU Emissions Trading System governs.¹¹⁹ The power-sector figures therefore bound the European effect in magnitude rather than measuring it: how much of it survives depends on how the allowance market absorbs the change, which lies outside the model. The productivity results do not turn on that question.

Taken together, the three exercises locate the productivity value of integration in the policy debate. A union does not close the transatlantic gap, and matching American interconnection delivers less than half of what completing the market is worth. But the union is worth about 60% more when fuel is expensive, and at prevailing carbon prices it halves power-sector emissions. These are benefits alongside the productivity gain rather than substitutes for it.

8 Conclusion

The value of integrating an input market depends on how production moves in response. For electricity in Europe, a full electricity union raises manufacturing productivity by 13 basis points today and 35 bp within the most electricity-intensive divisions. Over 80% percent of that gain arrives through production reallocating towards more productive locations and reinforcing clusters rather than through firms adjusting in place. The value roughly triples by 2040 as industry electrifies and the renewable build-out widens the dispersion at stake.

This exercise demonstrates what has to be measured to value integration of any input market beyond the direct price impact. How much production responds to a cost advantage is identified by the Norwegian bidding-zone divergence. How much of the cost advantage survives being used then depends on how steep the input's supply curve is where new demand arrives. Together, these factors determine the scale of the effect. The sign of the TFP impact beyond cost savings then depends on where production moves, and whether integration makes firms sort by productivity

¹¹⁸That figure applies today's allowance price to the 2015–2019 generation mix, which is considerably more fossil-intensive than the current fleet, so it overstates what the same carbon price would deliver on today's dispatch base.

¹¹⁹Under a strictly binding cap the two differ. Lower power-sector emissions reduce the allowance price and free permits for use elsewhere, so the abatement reappears outside the sector and aggregate emissions are unchanged — the waterbed effect. The Market Stability Reserve punctures the waterbed by cancelling surplus allowances, so that part of any reduction becomes permanent rather than being competed away (Perino, 2018). Which fraction survives depends on the reserve's intake and cancellation parameters over the relevant horizon, so integration's effect on European emissions as a whole lies between the power-sector figure reported here and zero. Quantifying it would require modelling the allowance market and is left to future work.

instead of cost advantages.

That gain is not visible from evaluating prices alone. A model of the supply side of electricity alone might even suggest negative effects of integration, as the demand-weighted average price increases with the union.¹²⁰ But the initial high-price locations are also the most productive, through location advantages and industry agglomeration. Relieving those locations of their input price congestion via integration allows for larger overall productivity gains through increased specialisation and concentration of production across space. Integration hence has a positive TFP impact equivalent to a uniform price decrease of 8% across the union, despite increasing the average price.

The European debate weighs electricity policy mainly by its contributions to decarbonisation and price or supply stability. The production side belongs alongside them, since how Europe builds, connects, and prices its electricity system decides where industry locates and how much the continent produces. Current price dispersion reflects both natural resources and a lack of integration, with uncoordinated plans for future renewable reliance increasing it further. Whether electrification ties production to geography is a policy choice rather than a physical fact.

¹²⁰The higher clearing price is not itself a resource cost. Payments above fuel cost are generation rents transferred within Europe, and the merged merit order burns less fuel than the segmented one.

References

- Allcott, Hunt, Allan Collard-Wexler, and Stephen D O’Connell. 2016. “How do electricity shortages affect industry? Evidence from India.” *American Economic Review*, 106(3): 587–624.
- Arkolakis, Costas, and Conor Walsh. 2023. “Clean growth.” National Bureau of Economic Research Working Paper 31615.
- Arkolakis, Costas, and Conor Walsh. 2024. “The Economic Impacts of Clean Energy.” *Brookings Papers on Economic Activity*, 2024(2): 161–212.
- Arndt, Lukas. 2023. “When Should We Believe Research Using ORBIS Firm Data?” Available at SSRN 4584106.
- Arpaia, Alfonso, Aron Kiss, Balázs Palvölgyi, and Alessandro Turrini. 2016. “Labour mobility and labour market adjustment in the EU.” *IZA Journal of Migration*, 5(1): 21.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. “Dynamic inputs and resource (mis) allocation.” *Journal of Political Economy*, 122(5): 1013–1063.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2019. “(Mis) allocation, market power, and global oil extraction.” *American Economic Review*, 109(4): 1568–1615.
- Bachmann, Rüdiger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick. 2024. “What if? The macroeconomic and distributional effects for Germany of a stop of energy imports from Russia.” *Economica*, 91(364): 1157–1200.
- Bajgar, Matej, Giuseppe Berlingieri, Sara Calligaris, Chiara Criscuolo, and Jonathan Timmis. 2020. “Coverage and representativeness of Orbis data.” OECD.
- Baqaee, David Rezza, and Emmanuel Farhi. 2020. “Productivity and misallocation in general equilibrium.” *The Quarterly Journal of Economics*, 135(1): 105–163.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. 2013. “Cross-country differences in productivity: The role of allocation and selection.” *American Economic Review*, 103(1): 305–334.
- BASF. 2021. “Curtailed ammonia production in Antwerp and Ludwigshafen.” *Trade news release*, Accessed 13 May 2026.
- BASF. 2023. “BASF shows resilience in a challenging market environment and implements measures to strengthen competitiveness.” *Business and financial news release*, Accessed 13 May 2026.
- Bau, Natalie, and Adrien Matray. 2023. “Misallocation and capital market integration: Evidence from India.” *Econometrica*, 91(1): 67–106.
- Beyer, Robert C. M., and Frank Smets. 2015. “Labour market adjustments and migration in Europe and the United States: how different?” *Economic Policy*, 30(84): 643–682.
- Bigio, Saki, and Jennifer La’O. 2020. “Distortions in production networks.” *The Quarterly Journal of Economics*, 135(4): 2187–2253.
- Brandt, Loren, Trevor Tombe, and Xiaodong Zhu. 2013. “Factor market distortions across time, space and sectors in China.” *Review of Economic Dynamics*, 16(1): 39–58.
- Brey, Björn. 2025. “Sparkling knowledge: Early technology adoption, innovation ability and long-run growth.” *Journal of the European Economic Association*, jvaf034.

- Broda, Christian, and David E Weinstein.** 2006. "Globalization and the gains from variety." *The Quarterly Journal of Economics*, 121(2): 541–585.
- Caliendo, Lorenzo, and Fernando Parro.** 2015. "Estimates of the trade and welfare effects of NAFTA." *The Review of Economic Studies*, 82(1): 1–44.
- Caliendo, Lorenzo, Luca David Oromolla, Fernando Parro, and Alessandro Sforza.** 2021. "Goods and factor market integration: A quantitative assessment of the EU enlargement." *Journal of Political Economy*, 129(12): 3491–3545.
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro.** 2019. "Trade and labor market dynamics: General equilibrium analysis of the china trade shock." *Econometrica*, 87(3): 741–835.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber.** 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review*, 101(3): 471–475.
- Choi, Bongseok.** 2020. "Productivity and misallocation of energy resources: Evidence from Korea's Manufacturing Sector." *Resource and Energy Economics*, 61: 101184.
- Cicala, Steve.** 2022. "Imperfect markets versus imperfect regulation in US electricity generation." *American Economic Review*, 112(2): 409–441.
- Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J Wagner.** 2025. "Does pricing carbon mitigate climate change? firm-level evidence from the european union emissions trading system." *Review of Economic Studies*, 92(3): 1625–1660.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon.** 2008. "Spatial wage disparities: Sorting matters!" *Journal of Urban Economics*, 63(2): 723–742.
- Costinot, Arnaud, and Dave Donaldson.** 2016. "How Large Are the Gains from Economic Integration? Theory and Evidence from U.S. Agriculture, 1880–1997." National Bureau of Economic Research 22946.
- Crafts, Nicholas, and Nikolaus Wolf.** 2014. "The location of the UK cotton textiles industry in 1838: a quantitative analysis." *The Journal of Economic History*, 74(4): 1103–1139.
- Davis, Steven J, Cheryl Grim, John Haltiwanger, and Mary Streitwieser.** 2008. "Electricity Pricing to US Manufacturing Plants, 1963–2000." National Bureau of Economic Research.
- Davis, Steven J, Cheryl Grim, John Haltiwanger, and Mary Streitwieser.** 2013. "Electricity unit value prices and purchase quantities: US manufacturing plants, 1963–2000." *Review of Economics and Statistics*, 95(4): 1150–1165.
- Draghi, Mario.** 2024. "The Future of European Competitiveness." *European Commission*, September 2024.
- Ellison, Glenn, and Edward L Glaeser.** 1999. "The geographic concentration of industry: does natural advantage explain agglomeration?" *American Economic Review*, 89(2): 311–316.
- Ellison, Glenn, Edward L Glaeser, and William R Kerr.** 2010. "What causes industry agglomeration? Evidence from coagglomeration patterns." *American Economic Review*, 100(3): 1195–1213.
- Engel, Charles, and John H. Rogers.** 1996. "How Wide Is the Border?" *American Economic Review*, 86(5): 1112–1125.
- European Commission.** 2023. "Assessment of the Potential for Energy Efficiency in Electricity Generation, Transmission and Storage." European Commission, Joint Research Centre SWD(2023) 1 final. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52023SC0001>.

- European Commission.** 2026. “Annual Report on Intra-EU Labour Mobility, 2025 Edition.” Publications Office of the European Union. Directorate-General for Employment, Social Affairs and Inclusion. Table 1, p. 13.
- European Network of Transmission System Operators for Electricity.** 2024a. “4th ENTSO-E Guideline for Cost Benefit Analysis of Grid Development Projects.” ENTSO-E. https://eepublicdownloads.blob.core.windows.net/public-cdn-container/clean-documents/news/2024/entso-e_4th_CBA_Guideline_240409.pdf.
- European Network of Transmission System Operators for Electricity.** 2024b. “Investment Candidates and Cost Assumptions.” *TYNDP 2024 dataset*, Version for opinion. https://eepublicdownloads.blob.core.windows.net/public-cdn-container/tyndp-documents/TYNDP2024/foropinion/investmentCandidates_and_CostAssumptions.xlsx.
- European Network of Transmission System Operators for Electricity.** 2024c. “TYNDP 2024 Cost-Benefit Analysis Implementation Guidelines.” ENTSO-E. Version for opinion. https://eepublicdownloads.blob.core.windows.net/public-cdn-container/tyndp-documents/TYNDP2024/foropinion/CBA_Implementation_Guidelines.pdf.
- European Network of Transmission System Operators for Electricity.** 2025. “Identification of System Needs: Methodology.” ENTSO-E. Version for opinion, 9 April 2025. <https://eepublicdownloads.blob.core.windows.net/public-cdn-container/tyndp-documents/TYNDP2024/foropinion/SystemNeedsMethodology.pdf>.
- Eurostat.** 2026a. “EU population diversity by citizenship and country of birth.” *Statistics Explained, European Commission*, Data extracted February 2026; reference date 1 January 2025.
- Eurostat.** 2026b. “GDP and Main Components (Output, Expenditure and Income).” *Dataset nama_10_gdp, European Commission*, Annual implicit deflator, item B1GQ, unit PD15_EUR, geography EU27_2020; data through 2024, extracted 3 September 2026. https://ec.europa.eu/eurostat/databrowser/view/nama_10_gdp/default/table?lang=en.
- Fajgelbaum, Pablo, and Stephen J Redding.** 2022. “Trade, structural transformation, and development: Evidence from Argentina 1869–1914.” *Journal of Political Economy*, 130(5): 1249–1318.
- Fajgelbaum, Pablo D, Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar.** 2019. “State taxes and spatial misallocation.” *The Review of Economic Studies*, 86(1): 333–376.
- Fiszbein, Martin, Jeanne Lafortune, Ethan G Lewis, and José Tessada.** 2020. “Powering up productivity: the effects of electrification on US manufacturing.” *NBER Working Paper 28076*.
- Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi.** 2023. “Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade.” *The Review of Economic Studies*, 90(1): 331–375.
- Gaubert, Cecile.** 2018. “Firm sorting and agglomeration.” *American Economic Review*, 108(11): 3117–3153.
- Glaeser, Edward L, Sari Pekkala Kerr, and William R Kerr.** 2015. “Entrepreneurship and urban growth: An empirical assessment with historical mines.” *Review of Economics and Statistics*, 97(2): 498–520.
- Gonzales, Luis E, Koichiro Ito, and Mar Reguant.** 2023. “The investment effects of market integration: Evidence from renewable energy expansion in Chile.” *Econometrica*, 91(5): 1659–1693.
- Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez.** 2017. “Capital allocation and productivity in South Europe.” *The Quarterly Journal of Economics*, 132(4): 1915–1967.

- Gotzens, Fabian, Heidi Heinrichs, Jonas Hörsch, and Fabian Hofmann.** 2019. "Performing energy modelling exercises in a transparent way: The issues of replicability and reproducibility." *Energy Strategy Reviews*, 25: 100–111. Reference publication for the powerplantmatching package.
- Granziera, Eleonora, Lovisa Reiche, Nicolò Maffei-Faccioli, Michael Weber, and Tuva Marie Fastbø.** 2025. "Five facts about inflation expectations: Evidence from four different groups of agents." *Norges Bank WP*, 7: 2025.
- Gray, Wayne, Joshua Linn, and Richard Morgenstern.** 2019. "The impacts of lower natural gas prices on jobs in the US manufacturing sector." *The Energy Journal*, 40(5): 169–193.
- Greaney, Brian.** 2026. "Housing constraints and spatial misallocation: Comment." *American Economic Journal: Macroeconomics*, 18(2): 409–428.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti.** 2010. "Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings." *Journal of Political Economy*, 118(3): 536–598.
- Gumber, Anurag, Riccardo Zana, and Bjarne Steffen.** 2024. "A global analysis of renewable energy project commissioning timelines." *Applied Energy*, 358: 122563.
- Gutberlet, Theresa.** 2014. "Mechanization and the spatial distribution of industries in the German Empire, 1875 to 1907." *The Economic History Review*, 67(2): 463–491.
- Hampole, Menaka, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller.** 2025. "Artificial intelligence and the labor market." National Bureau of Economic Research.
- Hassler, John, Per Krusell, and Conny Olovsson.** 2021. "Directed technical change as a response to natural resource scarcity." *Journal of Political Economy*, 129(11): 3039–3072.
- Head, Keith, and Thierry Mayer.** 2014. "Gravity equations: Workhorse, toolkit, and cookbook." In *Handbook of International Economics*. Vol. 4, 131–195. Elsevier.
- Henderson, J Vernon.** 2003. "Marshall's scale economies." *Journal of Urban Economics*, 53(1): 1–28.
- Hoes, Olivier AC, Lourens JJ Meijer, Ruud J Van Der Ent, and Nick C Van De Giesen.** 2017. "Systematic high-resolution assessment of global hydropower potential." *PLoS ONE*, 12(2): e0171844.
- Hsieh, Chang-Tai, and Enrico Moretti.** 2019. "Housing constraints and spatial misallocation." *American Economic Journal: Macroeconomics*, 11(2): 1–39.
- Hsieh, Chang-Tai, and Peter J Klenow.** 2009. "Misallocation and manufacturing TFP in China and India." *The Quarterly Journal of Economics*, 124(4): 1403–1448.
- Intergovernmental Panel on Climate Change.** 2006. "2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 2: Energy." <https://www.ipcc-nggip.iges.or.jp/public/2006gl/vol2.html>.
- International Carbon Action Partnership.** 2026a. "Allowance Price Explorer." <https://icapcarbonaction.com/en/ets-prices>, Accessed 19 June 2026.
- International Carbon Action Partnership.** 2026b. "Terms of Use and Documentation Notes: Allowance Price Explorer." <https://icapcarbonaction.com/en/documentation-allowance-price-explorer>, Accessed 19 June 2026.
- International Energy Agency, and Nuclear Energy Agency.** 2020. "Projected Costs of Generating Electricity: 2020 Edition." OECD Publishing, Paris. Lifetimes p. 36; construction lengths p. 38.

- Joint Research Centre.** 2017. "Covenant of Mayors for Climate and Energy: Default Emission Factors for Local Emission Inventories." European Commission JRC107518.
- Kahn, Matthew E, and Erin T Mansur.** 2013. "Do local energy prices and regulation affect the geographic concentration of employment?" *Journal of Public Economics*, 101: 105–114.
- Kalemli-Özcan, Şebnem, Bent E Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcen Yeşiltaş.** 2024. "How to construct nationally representative firm-level data from the orbis global database: New facts on smes and aggregate implications for industry concentration." *American Economic Journal: Macroeconomics*, 16(2): 353–374.
- Klein, Alexander, and Nicholas Crafts.** 2012. "Making sense of the manufacturing belt: determinants of US industrial location, 1880–1920." *Journal of Economic Geography*, 12(4): 775–807.
- Kline, Patrick, and Enrico Moretti.** 2014a. "Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority." *The Quarterly Journal of Economics*, 129(1): 275–331.
- Kline, Patrick, and Enrico Moretti.** 2014b. "People, places, and public policy: Some simple welfare economics of local economic development programs." *Annual Review of Economics*, 6(1): 629–662.
- Linn, Joshua.** 2008. "Energy prices and the adoption of energy-saving technology." *The Economic Journal*, 118(533): 1986–2012.
- LITGRID AB.** 2024. "Decisions Adopted at the Extraordinary General Meeting of Shareholders of LITGRID AB." 13 September 2024. <https://www.globenewswire.com/de/news-release/2024/09/13/2945975/0/en/Decisions-adopted-at-the-Extraordinary-General-Meeting-of-Shareholders-of-LITGRID-AB.html>.
- Loiacono, Luisa, Leonzio Rizzo, and Carlo Stagnaro.** 2025. "Impact of bidding zone re-configurations on electricity prices: Evidence from Sweden." *Energy Economics*, 141: 108106.
- Manderson, Edward J, and Richard Kneller.** 2020. "Energy endowments and the location of manufacturing firms." *Journal of Environmental Economics and Management*, 101: 102301.
- Marin, Giovanni, and Francesco Vona.** 2021. "The impact of energy prices on socioeconomic and environmental performance: Evidence from French manufacturing establishments, 1997–2015." *European Economic Review*, 135: 103739.
- Melitz, Marc J.** 2003. "The impact of trade on intra-industry reallocations and aggregate industry productivity." *Econometrica*, 71(6): 1695–1725.
- Michaels, Guy.** 2011. "The long term consequences of resource-based specialisation." *The Economic Journal*, 121(551): 31–57.
- Midrigan, Virgiliu, and Daniel Yi Xu.** 2014. "Finance and misallocation: Evidence from plant-level data." *American Economic Review*, 104(2): 422–458.
- Moll, Benjamin.** 2014. "Productivity losses from financial frictions: Can self-financing undo capital misallocation?" *American Economic Review*, 104(10): 3186–3221.
- Moll, Benjamin, Moritz Schularick, and Georg Zachmann.** 2023. "The power of substitution: The great German gas debate in retrospect." *Brookings Papers on Economic Activity*, Fall 2023: 395–455.

- Mundell, Robert A.** 1957. "International trade and factor mobility." *The American Economic Review*, 47(3): 321–335.
- Newbery, David, Goran Strbac, and Ivan Viehoff.** 2016. "The benefits of integrating European electricity markets." *Energy Policy*, 94: 253–263.
- Nordic Regional Coordination Centre.** 2024. "Overview of Settings at Flow-Based Go-Live." Nordic Regional Coordination Centre. <https://nordic-rcc.net/wp-content/uploads/2024/09/05-Overview-of-settings-at-flow-based-go-live.pdf>.
- Nord Pool.** n.d. "Loss Functionality." <https://www.nordpoolgroup.com/en/trading/Day-ahead-trading/loss-functionality/>, Accessed September 2, 2026.
- Norsk Hydro.** 2022. "Slovalco will stop primary aluminium production." *Company news release*, Accessed 13 May 2026.
- North American Electric Reliability Corporation.** 2024. "Interregional Transfer Capability Study (ITCS): Final Report." November 2024. Chapter 4: Total Import Interface Results; Appendix B: Load and Resource Data.
- Open Power System Data.** 2020. "Data Package: Conventional Power Plants." *Open Power System Data platform*, Version 2020-10-01.
- Panhans, Matthew, Lucia Lavric, and Nick Hanley.** 2017. "The effects of electricity costs on firm re-location decisions: insights for the pollution havens hypothesis?" *Environmental and Resource Economics*, 68(4): 893–914.
- Perino, Grischa.** 2018. "New EU ETS Phase 4 rules temporarily puncture waterbed." *Nature Climate Change*, 8(4): 262–264.
- Polskie Sieci Elektroenergetyczne.** 2024. "Finansowa decyzja inwestycyjna na budowę Harmony Link zaakceptowana." 23 December 2024. <https://www.pse.pl/-/finansowa-decyzja-inwestycyjna-na-budowe-harmony-link-zaakceptowana>.
- Rasul, Kajwan, Sarah Schmidt, Edgar G. Hertwich, and Richard Wood.** 2024. "EXIOBASE energy accounts: Improving precision in an open-sourced procedure applicable to any MRIO database." *Journal of Industrial Ecology*, 28(6): 1771–1785.
- Redding, Stephen J.** 2016. "Goods trade, factor mobility and welfare." *Journal of International Economics*, 101: 148–167.
- Restuccia, Diego, and Richard Rogerson.** 2008. "Policy distortions and aggregate productivity with heterogeneous establishments." *Review of Economic Dynamics*, 11(4): 707–720.
- Restuccia, Diego, and Richard Rogerson.** 2017. "The causes and costs of misallocation." *Journal of Economic Perspectives*, 31(3): 151–174.
- Ruiz, Pablo, Wouter Nijs, Denis Tarvydas, Alfredo Sgobbi, Andreas Zucker, Roberto Pilli, Ragnar Jonsson, Andrea Camia, Christian Thiel, Carsten Hoyer-Klick, Sylvain Leduc, Anders Olsson, Luc Beurskens, Stathis Peteves, and Alejandro Soria.** 2019. "ENSPRESO – an open, EU-28 wide, transparent and coherent database of wind, solar and biomass energy potentials." *Energy Strategy Reviews*, 26: 100379.
- Sahraei-Ardakani, Mostafa, Seth Blumsack, and Andrew Kleit.** 2015. "Estimating zonal electricity supply curves in transmission-constrained electricity markets." *Energy*, 80: 10–19.

- Samuelson, Paul A.** 1948. "International trade and the equalisation of factor prices." *The Economic Journal*, 58(230): 163–184.
- Saussay, Aurelien, and Misato Sato.** 2024. "The impact of energy prices on industrial investment location: Evidence from global firm level data." *Journal of Environmental Economics and Management*, 127: 102992.
- Severnini, Edson.** 2023. "The power of hydroelectric dams: historical evidence from the United States over the twentieth century." *The Economic Journal*, 133(649): 420–459.
- Shapiro, Joseph S.** 2025. "Institutions, comparative advantage, and the environment." *Review of Economic Studies*, 92(6): 4152–4193.
- Sraer, David, and David Thesmar.** 2023. "How to use natural experiments to estimate misallocation." *American Economic Review*, 113(4): 906–938.
- Stadler, Konstantin, Richard Wood, Tatyana Bulavskaya, Carl-Johan Södersten, Moana Simas, Sarah Schmidt, Arkaitz Usubiaga, José Acosta-Fernández, Jeroen Neuwahl, Martin Wiebe, Anton Schafartzik, Jochen Mutel, Stefano Faraco, and Arnold Tukker.** 2018. "EXIOBASE 3: Developing a time series of detailed environmentally extended multi-regional input–output tables." *Journal of Industrial Ecology*, 22(3): 502–515.
- Statnett.** 2013. "Kabler til Tyskland og Storbritannia – analyse av samf.øk. nytte ved spothandel." Statnett Analyserapport: underlag til konsesjonssøknad, Oslo.
- Suárez Serrato, Juan Carlos, and Owen Zidar.** 2016. "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms." *American Economic Review*, 106(9): 2582–2624.
- Suárez Serrato, Juan Carlos, and Owen Zidar.** 2023. "Who Benefits from State Corporate Tax Cuts? A Local Labor Market Approach with Heterogeneous Firms: Reply." *American Economic Review*, 113(12): 3401–3410.
- Tan, Ruipeng, Jun Li, Feng Wang, and Jing Yang.** 2025. "Electricity transmission and its effects on resource misallocation and environmental consequences in China: An economic analysis." *Energy Policy*, 198: 114490.
- Tombe, Trevor, and Xiaodong Zhu.** 2019. "Trade, migration, and productivity: A quantitative analysis of China." *American Economic Review*, 109(5): 1843–1872.
- Wahlstrøm, Ranik Raaen.** 2022. "Financial statements of companies in Norway." *arXiv preprint arXiv:2203.12842*.
- Wolverton, Ann, Ronald Shadbegian, and Wayne B Gray.** 2022. "The US manufacturing sector's response to higher electricity prices: evidence from state-level renewable portfolio standards." National Bureau of Economic Research.
- Xu, Mengmeng, and Ruipeng Tan.** 2021. "Removing energy allocation distortion to increase economic output and energy efficiency in China." *Energy Policy*, 150: 112110.
- Yang, Mian, Fuxia Yang, and Chuanwang Sun.** 2018. "Factor market distortion correction, resource reallocation and potential productivity gains: An empirical study on China's heavy industry sector." *Energy Economics*, 69: 270–279.
- Yara International.** 2021. "Yara ammonia curtailment volume update." *Corporate release*, Accessed 13 May 2026.
- Yara International.** 2022. "Yara implements further production curtailments in Europe." *Corporate release*, Accessed 13 May 2026.

Appendix

A Institutional background: European electricity market design

A.1 Wholesale market structure

European electricity wholesale markets are organised into bidding zones within which a single market-clearing price is determined through the merit-order mechanism: generators bid in order of increasing marginal cost, and the price is set by the last unit dispatched to meet demand. In equilibrium, all inframarginal generators earn the same price, and this price equals the marginal cost of the most expensive active plant.

Bidding zones are defined primarily at the national level. Exceptions include Norway (5 zones: NO1–NO5), Sweden (4: SE1–SE4), Italy (6), Denmark (2: DK1–DK2), and the combined German-Luxemburg zone (DE-LU). Spain and Portugal operate as a de facto integrated zone. Within each bidding zone, congestion is managed by redispatch (adjusting generator output after the market clears) rather than by prices, which creates a disconnect between physical flows and price signals.

A.2 Cross-border trade and interconnectors


Electricity trade between zones occurs through physical interconnectors, with capacity allocated through market coupling mechanisms. The Net Transfer Capacity (NTC) of each interconnector — the maximum amount of power that can flow without threatening system security — is published by system operators and changes seasonally. Average cross-border trade represents approximately 10% of consumption for EU member states, though this varies widely: the Nordic region is highly integrated internally, while countries such as Spain and the Iberian peninsula are effectively energy islands. Figure A.1 compares annual interconnection capacity on selected borders with electricity consumption on each side of the border, and then reports the corresponding annual average wholesale price gap.

Market coupling integrates prices across interconnected zones: when the price in zone *A* exceeds that in zone *B* by more than the transmission constraint allows, power flows from *B* to *A* until either prices equalise or the interconnector reaches its capacity limit. This means that countries with abundant cheap generation export to their neighbours, but the benefit is limited by transmission bottlenecks.

A.3 Sources of price dispersion

Price differences across European countries arise from three sources. First, natural advantages: countries with abundant low-cost resources — Norwegian and Alpine hydro, Southern European solar, North Sea wind — face structurally lower supply costs. Figure A.2 shows technical potential

Figure A.1: Cross-border interconnectors and persistent price gaps, selected pairs




Notes: Selected country pairs with a physical interconnector and complete 2019 price and demand data. Panel (a) reports annual transfer potential divided by annual electricity demand on each side of the border. Annual transfer potential is the mean of the two directional country-level NTC values in MW multiplied by 8,760 hours and converted to TWh/year. Panel (b) reports the absolute difference in 2019 average wholesale electricity prices. Sources: ENTSO-E Transparency Platform (NTC), Bloomberg (prices), and Ember (annual electricity demand).

for solar PV, wind, and hydro by country. Potential is concentrated in large countries — France, Spain, Great Britain, Italy — and dominated by solar and wind. The ranking by potential does not match the ranking by prices: Italy has among the highest potential but also the highest prices, while smaller countries with modest potential (Denmark, Sweden) have low prices.

Second, generation mix policy choices: decisions on nuclear (Germany’s phase-out, France’s reliance on its historical fleet, Finland’s Olkiluoto 3 investment), renewable subsidy levels, and permitted coal use directly shape the supply curve. Third, market structure: the segmentation into national pricing zones, combined with limited transmission capacity, prevents cost advantages from fully spreading to neighbours.

Figure A.2: Technical renewable potential by source and country




Notes: The figure reports country-level technical renewable potential by source in TWh/year. Solar is JRC ENSPRESO solar PV technical generation potential, summing roof/facade and ground-mounted PV and excluding CSP. Wind is JRC ENSPRESO onshore plus offshore technical generation potential under the reference large-turbine scenario. Hydro is the current country-level theoretical hydropower-potential measure in the project hydro table and is included separately from ENSPRESO. It keeps only countries with all three source components.

A.4 Retail prices for industrial consumers

Industrial consumers face retail prices that differ from wholesale prices through several layers: network charges (transmission and distribution tariffs), energy taxes, levies for renewable support schemes (e.g., the German EEG surcharge until 2022), and contract structures. Large consumers in some countries can access wholesale prices directly or with favourable long-term contracts, while smaller consumers pay regulated tariffs. The ratio of retail to wholesale industrial prices varies substantially across countries, from close to 1 in Nordic countries to above 2 in some central European countries, and within a sector the demand-weighted dispersion of firm-facing prices across countries is about a third larger than the dispersion of wholesale prices. The main text works with wholesale prices because they capture the underlying fundamental cost differences the theory speaks to. The consequence differs for the two objects the paper builds on, and runs in opposite directions: for estimation, measuring the price at wholesale attenuates the response to what firms actually pay, so the reported elasticities are conservative; for the counterfactual, a wholesale union removes only part of the dispersion firms face, so it overstates what integration alone delivers. Appendix G.10 prices both.

A.5 Effective cost shares by country

Figure A.3: Electricity shares countries



Notes: Country-level composition-predicted electricity cost share, constructed as the gross-output-weighted average of the Exiobase 2010–2011 sector cost shares, using 2015–2019 Eurostat country-sector nominal gross-output shares. Data exclude electricity supply (NACE D35).

A.6 Norwegian bidding zone configuration

Norway's five bidding zones (NO1–NO5) follow the major transmission bottlenecks in the Norwegian high-voltage grid, as established in the 2008–2010 zone reform. The configuration reflects the geography of generation (concentrated in the western fjords and the north) and load (concentrated in the south). Southern zones (NO1, NO2, NO5) are directly connected to continental Europe and other Nordic countries through multiple HVDC cables, as detailed in Section 4. Northern zones (NO3, NO4) are connected only to southern Norway and Sweden, providing the geographic insulation that drives the identification strategy. Wholesale prices within each Norwegian zone are cleared jointly with the Nordic market (Nord Pool), with physical constraints enforced after the fact.

B Model derivations

This appendix records the derivations underlying the theoretical results stated in Section 3. The main text and appendix use the same convention: ψ is the dispersion of idiosyncratic location payoffs measured in units of normalised log profit, i.e. log operating profit divided by $(\sigma - 1)$, so that location shares are $\exp(v_{js}/\psi) / \sum_k \exp(v_{ks}/\psi)$ in the systematic payoff v_{js} of Lemma 1. This yields the compact location multiplier $1/(\psi - \gamma)$, while $(\sigma - 1)$ remains explicit in profits, factor demands, and market clearing. Single-sector results drop the sector subscript s and define $A \equiv \alpha_s(\sigma - 1)$ and $B \equiv \beta_s(\sigma - 1)$, where α_s and β_s are the labour and electricity elasticities of the production function. The residual-input share is written $\kappa_s \equiv 1 - \alpha_s - \beta_s > 0$ wherever it appears, so that $\alpha_s + \beta_s < 1$ throughout: the derivations never impose $\alpha_s = 1 - \beta_s$, and the residual input enters only through the common price p^X , which drops from every location comparison.

B.1 Proof: logit location shares

Firm f in sector s chooses country j^* to maximise normalised log profit $v_{js} + a_{fjs}$, where $v_{js} \equiv \bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E$ is the systematic component and $a_{fjs} \stackrel{iid}{\sim} \text{Type I EV}(0, \psi)$ is the idiosyncratic firm–location match from equation (2). The extreme value theorem gives

$$\Pr(j^* = j) = \frac{\exp(v_{js}/\psi)}{\sum_k \exp(v_{ks}/\psi)}$$

The law of large numbers applied to a continuum of firms gives $m_{js} = \Pr(j^* = j)$, which is Lemma 1. Taking logs, $\ln M_{js} = \ln M_s + v_{js}/\psi - \bar{V}_s$, where $\bar{V}_s \equiv \ln \sum_k \exp(v_{ks}/\psi)$ is the log-sum (inclusive value), common across all countries j within sector s and absorbed by sector-time fixed effects in the empirical specification. $\square$

B.2 Proof: location multiplier

Log-linearise $\ln M_{js} = \ln M_s + v_{js}/\psi - \bar{V}_s$ with $d \ln M_s = 0$ (exogenous total mass). The inclusive-value change $d\bar{V}_s$ is common to all countries within sector s , so it is eliminated by expressing changes in deviation from the sector mean (country differencing); its empirical counterpart is absorbed by sector-time fixed effects. In these deviations:

$$d \ln M_{js} \approx \frac{dv_{js}}{\psi} = \frac{1}{\psi} [\gamma d \ln M_{js} - \alpha_s d \ln w_j - \beta_s d \ln p_j^E]$$

using $d \ln \phi_{js} = \gamma d \ln M_{js}$ from the agglomeration externality, with $d\bar{\Phi}_{js} = 0$ since base productivity is a fixed fundamental. Rearranging:

$$d \ln M_{js}(1 - \gamma/\psi) = \frac{1}{\psi} [-\alpha_s d \ln w_j - \beta_s d \ln p_j^E]$$

which delivers $d \ln M_{js} = \frac{1}{\psi - \gamma} [-\alpha_s d \ln w_j - \beta_s d \ln p_j^E]$. Stability requires $\psi > \gamma$. $\square$

B.3 Mapping to empirical specifications

The location equation (Proposition 1) maps directly to the empirical interaction coefficient. Different dependent variables identify different composites of the structural parameters. With country-time FEs absorbing $d \ln w_j$ and $d \ln p_j^E$, and sector-time FEs absorbing $d \ln B_s$, the identifying variation comes from the interaction of the sector's electricity intensity with the local price change.

The interaction is formed with the measured electricity cost share $\tilde{\beta}_s$, electricity purchases over gross output, rather than with the production elasticity β_s of the model. Under constant markups the two differ by a known factor: expenditure on an input is a fraction $(\sigma - 1)/\sigma$ of revenue, so $\tilde{\beta}_s = \frac{\sigma-1}{\sigma} \beta_s$ and $\tilde{\alpha}_s = \frac{\sigma-1}{\sigma} \alpha_s$. Every coefficient below is therefore the corresponding model composite multiplied by $\sigma/(\sigma - 1)$.

Dependent variable	Level	Margin	Coefficient $\hat{\delta}$ on $\tilde{\beta}_s \times d \ln p_j^E$	Content
Firm counts M_{js}	(j, s)	Extensive	$-\frac{\sigma}{\sigma-1} \frac{1}{\psi-\gamma}$	Prop. 1
Firm profit π_{fjs}	Firm	Intensive	$-\frac{\sigma\psi}{\psi-\gamma}$	Direct + agglom.
Firm employment ℓ_{fjs}	Firm	Intensive	$-\frac{\sigma\psi}{\psi-\gamma}$	Same (w_j absorbed)
Total employment L_{js}	(j, s)	Both	$-\frac{\sigma}{\sigma-1} \frac{1+(\sigma-1)\psi}{\psi-\gamma}$	$\hat{\delta}_M + \hat{\delta}_\pi$

Written in terms of the model elasticity β_s , the same four entries are $-1/(\psi - \gamma)$, $-(\sigma - 1)\psi/(\psi - \gamma)$ twice, and $-[1 + (\sigma - 1)\psi]/(\psi - \gamma)$.

The key relationships are: $\hat{\delta}_\pi/\hat{\delta}_M = (\sigma - 1)\psi$ (intensive/extensive ratio, unaffected by the markup factor because it is common to both), $\hat{\delta}_\Pi = \hat{\delta}_M + \hat{\delta}_\pi$ (aggregate decomposes additively),

and all firm-level intensive-margin measures are equivalent under country-time FEs. In principle these provide over-identifying restrictions: with σ calibrated externally, $\hat{\delta}_\pi/\hat{\delta}_M$ pins down ψ and $\hat{\delta}_M$ then gives γ residually, while $\hat{\delta}_\Pi = \hat{\delta}_M + \hat{\delta}_\pi$ is an exact cross-check. In practice the intensive margin is too imprecise to carry the split, so the calibration of Section 5 anchors γ externally and reads the composite $\psi - \gamma = \frac{\sigma}{\sigma-1}/|\hat{\delta}_M|$ off the extensive margin, retaining the ratio as a consistency check.

Derivation of intensive margin coefficient. For a continuing firm with fixed idiosyncratic draw, $d \ln \pi_{fjs} = (\sigma - 1)\gamma d \ln M_{js} - A d \ln w_j - B d \ln p_j^E + d \ln B_s$. Substituting Proposition 1 for $d \ln M_{js}$ gives

$$\delta_\pi = -\frac{(\sigma - 1)\psi}{\psi - \gamma} = -\left[\frac{(\sigma - 1)\gamma}{\psi - \gamma} + (\sigma - 1)\right]$$

as the coefficient on $\beta_s d \ln p_j^E$, and hence $\hat{\delta}_\pi = \frac{\sigma}{\sigma-1}\delta_\pi = -\sigma\psi/(\psi - \gamma)$ on the measured share. The first term is the indirect agglomeration effect (entry of new firms raises continuing firms' productivity); the second is the direct cost effect through factor prices. The aggregate coefficient follows from $d \ln \Pi_{js} = d \ln M_{js} + d \ln \bar{\pi}_{js}$ and noting that the average profit per firm responds identically to a continuing firm's profit. $\square$

B.4 Electricity supply curve with non-fuel costs

Country j 's fossil fleet consists of a continuum of plants ordered by thermal efficiency. The following assumptions generate the micro-founded supply curve: (i) heat rates, the inverse of thermal efficiency, follow a common Pareto distribution with shape $\zeta > 1$ and lower bound $h_{\min}$, the heat rate of the frontier plant; (ii) plants are equal-sized; (iii) each plant has a linear (Leontief) technology with no capital-fuel substitution; (iv) dispatch is binary (on or off); (v) all plants face a common all-in fuel-use cost p^F (fuel plus carbon; see below). Under these assumptions, the heat rate of the plant at fractional position $k \in [0, 1]$ in the fleet is $h(k) = h_{\min}(1 - k)^{-1/\zeta}$, the quantile function of that distribution.

Under merit-order dispatch, the equilibrium electricity price in the fossil regime ($E_j^D > \bar{E}_j^0$) equals the marginal cost of the last dispatched plant:

$$p_j^E = c_0^F + p^F h_{\min}(1 - u_j)^{-1/\zeta}, \quad u_j \equiv \frac{E_j^D - \bar{E}_j^0}{K_j^F} \in (0, 1) \quad (\text{B.1})$$

where $c_0^F \geq 0$ is a non-fuel component of fossil marginal cost (variable O&M) and u_j is the fossil utilisation rate.

Fuel and carbon costs. The price p^F in (B.1) is the all-in marginal fuel-use cost, combining the world fuel price $p^{F,\text{fuel}}$ with the EU ETS carbon cost:

$$p_j^E = c_0^F + h_{\min} \left(p^{F,\text{fuel}} + \tau^{\text{CO}_2} e^F \right) (1 - u_j)^{-1/\zeta}, \quad p^F \equiv p^{F,\text{fuel}} + \tau^{\text{CO}_2} e^F, \quad (\text{B.2})$$

where τ^{CO_2} is the EU ETS allowance (EUA) price and e^F is the fuel-specific emissions intensity in tonnes of CO_2 per MWh of fuel input. Carbon therefore enters the cost stack used to fit ζ and the structural validation of the supply schedule. Keeping $\tau^{\text{CO}_2} e^F$ as a separable component makes the carbon price a lever for the counterfactuals; ETS price endogeneity and emissions-cap feedback are left as a future extension.

Define the **fuel cost share**

$$\chi_j \equiv \frac{p_j^E - c_0^F}{p_j^E} = \frac{p^F h_{\min} (1 - u_j)^{-1/\zeta}}{p_j^E} \in (0, 1] \quad (\text{B.3})$$

which measures the fraction of the electricity price attributable to fuel costs. When $c_0^F = 0$, $\chi_j = 1$ for all countries. The fuel cost share enters the supply elasticity and the fuel price pass-through.

Supply elasticity. By the chain rule through u_j :

$$\varepsilon_j \equiv \frac{\partial \ln p_j^E}{\partial \ln E_j^D} = \underbrace{\frac{\chi_j}{\zeta(1 - u_j)}}_{\partial \ln p_j^E / \partial u_j} \times \underbrace{\frac{E_j^D}{K_j^F}}_{\partial u_j / \partial \ln E_j^D}$$

The log-price derivative picks up χ_j because only the fuel-cost component responds to changes in utilisation; the non-fuel component c_0^F is fixed. Combining:

$$\varepsilon_j = \frac{\chi_j}{\zeta} \frac{E_j^D}{K_j^F (1 - u_j)} = \frac{\chi_j}{\zeta} \frac{E_j^D}{K_j^F - E_j^F} \quad (\text{B.4})$$

When $c_0^F = 0$: $\varepsilon_j = E_j^D / [\zeta K_j^F (1 - u_j)]$. Three properties are essential: (i) ε_j is increasing in u_j ; (ii) cross-country variation comes through u_j , E_j^D / K_j^F , and χ_j , not from differences in ζ ; (iii) ε_j is endogenous through u_j , evaluated at the initial equilibrium for log-linearisation and updated step-by-step along the counterfactual path in the quantitative model.

Proof of Lemma 2. Three technologies can be at the margin, but the first two deliver the same zero elasticity, which is why the lemma states two cases: the price is then a constant, set by the opportunity value of surplus renewable output in the first case and by the nuclear marginal cost in the second, so domestic demand moves quantities without moving the price. (1) *Renewable margin.*

If $E_j^D < \theta_j^R K_j^R$, zero-marginal-cost renewable capacity is not fully utilised domestically. In a literal autarkic engineering model the marginal cost is zero, but the log-linear location model is evaluated at a positive wholesale price. The relevant price is therefore the opportunity price $p_j > 0$ generated by export options, contracts, curtailment rules, or other market-design floors, and $\varepsilon_j = 0$ locally with respect to domestic demand. (2) *Nuclear margin*. If $\theta_j^R K_j^R \leq E_j^D < \bar{E}_j^0$, nuclear is the marginal technology, $p_j^E = c^N > 0$, and $\varepsilon_j = 0$ locally, since small variations in demand are absorbed by nuclear output without a price change. (3) *Fossil margin*. If $E_j^D \geq \bar{E}_j^0$, demand is served by fossil plants at the margin and equations (B.1)–(B.4) apply, giving $p_j^E = c_0^F + p^F h_{\min}(1 - u_j)^{-1/\zeta}$ and $\varepsilon_j = (\chi_j/\zeta)(E_j^D/[K_j^F(1 - u_j)]) > 0$. The comparative statics in the main text take the fossil-marginal case as the working regime, since fossil dispatch sets the price at the margin in essentially all country-hours observed in the sample period, and the fossil-marginal case is the only one in which domestic supply-side shocks transmit to the wholesale price through a continuously sloped supply curve. $\square$

B.5 GE solution: the single-sector case and the dampening composite

With one sector the labour and electricity weightings of Lemma 3 coincide, $\bar{\alpha}_j^L = \bar{\alpha}_j^E = \alpha$ and $\bar{\beta}_j^L = \bar{\beta}_j^E = \beta$, and the country response collapses to a closed form: the whole of general equilibrium is then one scalar, the dampening composite Ω_j of Proposition 4 below. This subsection derives it, and the impulse cases that the levers of Section 6 deliver.

Proposition 4 (Location response in single-sector general equilibrium). *Consider the single-sector model with endogenous wages and electricity prices. The firm mass response to a direct price impulse $d \ln p_j^E|_{ED}$ is*

$$d \ln M_j = -\frac{\beta}{\psi - \gamma} \frac{1}{\Omega_j} d \ln p_j^E|_{ED}, \quad (\text{B.5})$$

where the general equilibrium dampening $\Omega_j \geq 1$ is

$$\Omega_j \equiv \underbrace{1 + \alpha(\sigma - 1) + \varepsilon_j [1 + (\alpha + \beta)(\sigma - 1)]}_{\text{price response at a given firm distribution}} + \underbrace{\frac{[1 + (\sigma - 1)\gamma] [\alpha + \varepsilon_j(\alpha + \beta)]}{\psi - \gamma}}_{\text{congestion from relocation}}.$$

The dampening collects what resists the move in two parts. The first is the response of factor demand to prices at a given distribution of firms, which absorbs part of the impulse. The second is congestion as firms arrive and their demand raises the local wage and electricity price; it is divided by the net dispersion force $\psi - \gamma$ because the further firms move, the harder they press on local factor markets. Under the PE stability condition $\psi > \gamma$ both terms are positive and $\Omega_j \geq 1$, so general equilibrium always dampens. In the multi-sector model the same two forces appear, but they no longer collapse to one number: the wage offset is sector-specific through $\alpha_s \lambda_j$, while

the dampening $1/(1 + \varepsilon_j \Theta_j)$ is common to the country (Appendix B.8).

The three-equation log-linearised system (suppressing sector subscripts) is:

$$\text{(Location)} \quad d \ln M_j = \frac{1}{\psi - \gamma} \left[-\alpha d \ln w_j - \beta d \ln p_j^E \right] \quad (\text{B.6})$$

$$\text{(Wage)} \quad d \ln w_j = \frac{1}{1 + A} \left[(1 + (\sigma - 1)\gamma) d \ln M_j - B d \ln p_j^E \right] \quad (\text{B.7})$$

$$\text{(Price)} \quad d \ln p_j^E = \frac{\varepsilon_j}{\xi_j} \left[(1 + (\sigma - 1)\gamma) d \ln M_j - A d \ln w_j \right] - \frac{\text{supply-shift}}{\xi_j} \quad (\text{B.8})$$

where $\xi_j \equiv 1 + \varepsilon_j(1 + B)$ and the supply-shift term depends on the shock type (see below); it is the direct price impulse of the main text with the sign reversed, $\text{supply-shift} = -d \ln p_j^E|_{ED}$, so that a favourable shift enters positively. The wage equation follows from log-linearising $M_j \alpha (\sigma - 1) \bar{\pi}_j / w_j = \bar{L}_j$, using $d \ln \bar{\pi}_j = (\sigma - 1)\gamma d \ln M_j - A d \ln w_j - B d \ln p_j^E$. The price equation follows from log-linearising the supply curve (B.1) after substituting the electricity demand response, which includes the direct mass effect, agglomeration-induced profit scale, direct price effect $(-1/p_j^E)$, and profit-scale price effect $(-\beta(\sigma - 1))$ through factor prices), producing the combined denominator ξ_j .

Step 1: eliminate $d \ln w_j$ from (Price). Substituting (B.7) into (B.8) and defining $\Gamma_j \equiv 1 + A + \varepsilon_j(1 + A + B)$:

$$d \ln p_j^E = \frac{\varepsilon_j [1 + (\sigma - 1)\gamma]}{\Gamma_j} d \ln M_j - \frac{\text{supply-shift}(1 + A)}{\Gamma_j}$$

Step 2: back-substitute into (Wage):

$$d \ln w_j = \frac{[1 + (\sigma - 1)\gamma] (1 + \varepsilon_j)}{\Gamma_j} d \ln M_j + \frac{B \text{supply-shift}}{\Gamma_j}$$

Step 3: substitute both into (Location) and use $\alpha(1 + \varepsilon_j) + \beta \varepsilon_j = \alpha + \varepsilon_j(\alpha + \beta)$:

$$d \ln M_j \left[(\psi - \gamma) + \frac{[1 + (\sigma - 1)\gamma] [\alpha + \varepsilon_j(\alpha + \beta)]}{\Gamma_j} \right] = \frac{\beta}{\Gamma_j} \times \text{supply-shift}$$

using $\beta(1 + A) - \alpha B = \beta$ on the supply-shift terms carried by (B.7) and (B.8). With $\alpha + \beta = 1$ the two composites collapse to $\Gamma_j = 1 + A + \varepsilon_j \sigma$ and $\alpha + \varepsilon_j$.

The single-sector linearised equilibrium is locally nonsingular if

$$(\psi - \gamma)\Gamma_j + [1 + (\sigma - 1)\gamma] [\alpha + \varepsilon_j(\alpha + \beta)] \neq 0. \quad (\text{B.9})$$

For the economically relevant monotone comparative statics in Proposition 4, I impose that the left-hand side of (B.9) is positive. The PE condition $\psi > \gamma$ is sufficient because $\Gamma_j > 0$, $\alpha + \varepsilon_j(\alpha + \beta) > 0$.

$\beta) \geq 0$, and $1 + (\sigma - 1)\gamma > 0$, but it is not necessary: sufficiently strong wage and electricity congestion can stabilize the local system even when γ exceeds ψ . When $\psi > \gamma$, the left-hand side of (B.9) equals $(\psi - \gamma)\Omega_j$, with Ω_j the GE dampening composite of Proposition 4.¹²¹

For a generic favourable supply-side shift g_j with direct fixed-demand price effect $d \ln p_j^E|_{E_j^D} = -\iota_j g_j$, the supply-shift is $\iota_j g_j$, where the impulse coefficient ι_j takes the value listed in (B.10) for each lever: ε_j for a proportional capacity expansion, ε_j^0 for a low-cost one, and the fuel cost share χ_j for a fall in the fuel price. The solution is therefore

$$d \ln M_j = \frac{\beta \iota_j}{(\psi - \gamma)\Omega_j} g_j,$$

delivering Proposition 4.

For a *proportional capacity expansion* ($d \ln \bar{E}_j^0 = d \ln K_j^F \equiv d \ln \bar{E}_j$), the supply-shift is $\varepsilon_j d \ln \bar{E}_j$, so $\iota_j = \varepsilon_j$. For a *low-cost capacity expansion* ($d \bar{E}_j^0 > 0$, $d K_j^F = 0$), the supply-shift is $\varepsilon_j^0 d \ln \bar{E}_j^0$ where $\varepsilon_j^0 \equiv \varepsilon_j \bar{E}_j^0 / E_j^D < \varepsilon_j$, delivering the low-cost expansion case:

$$d \ln M_j = \frac{\beta \varepsilon_j^0}{(\psi - \gamma)\Omega_j} d \ln \bar{E}_j^0$$

The dampening composite Ω_j is unchanged across shock types — only the exogenous impulse differs. $\square$

Corollary 1 (Symmetric two-country location share: logistic approximation). *With $J = 2$ countries of equal labour supply ($\bar{L}_1 = \bar{L}_2$) and equal supply elasticities at the symmetric equilibrium ($\varepsilon_1 = \varepsilon_2 = \varepsilon$), the equilibrium location share of country 1 in a single sector satisfies, to first order in the capacity differential around the symmetric point $\bar{E}_1 = \bar{E}_2$,*

$$m_1 \approx \frac{1}{1 + \exp\left(-\frac{\beta \varepsilon}{(\psi - \gamma)\Omega} (\ln \bar{E}_1 - \ln \bar{E}_2)\right)},$$

where $\Gamma \equiv 1 + \alpha(\sigma - 1) + \varepsilon [1 + (\alpha + \beta)(\sigma - 1)]$ and $\Omega \equiv \Gamma + [1 + (\sigma - 1)\gamma] [\alpha + \varepsilon(\alpha + \beta)] / (\psi - \gamma)$ are evaluated at the symmetric equilibrium.

Derivation of Corollary 1. In the symmetric 2-country case ($\bar{L}_1 = \bar{L}_2$, $\varepsilon_1 = \varepsilon_2 = \varepsilon$), the system is expressed in differences $\Delta x \equiv \ln x_1 - \ln x_2$. The differenced log-linearised location equation is structurally identical to the single-country system, giving $\Delta \ln M = (\beta \varepsilon / [(\psi - \gamma)\Omega]) \Delta \ln \bar{E}$, with Ω the common value of Ω_j at the symmetric point. Since $\Delta \ln M = \ln(m_1 / (1 - m_1))$, inverting gives the logistic form. The result is exact in the linearisation but local: away from the symmetric

¹²¹The condition is stated on the unfactored expression because factoring out $\psi - \gamma$ presumes its sign: were γ to exceed ψ , the congestion block of Ω_j would change sign and $\Omega_j \geq 1$ would fail.

point, ε_j and Ω_j vary with the allocation, so the slope $\beta\varepsilon/[(\psi - \gamma)\Omega]$ is not constant and the global share function is logistic only approximately. $\square$

B.6 Supply-shift decomposition

Price impulses by policy lever. The levers differ only in the direct price impulse they deliver:

$$d \ln p_j^E \Big|_{ED} = \begin{cases} \ln \bar{p}^E - \ln p_j^E & \text{(market integration)} \\ -\frac{\varepsilon_j \bar{E}_j^0}{E_j^D} d \ln \bar{E}_j^0 & \text{(low-cost capacity)} \\ \frac{p_j^E - c_0^F}{p_j^E} d \ln p^F & \text{(fuel price)} \end{cases} \quad (\text{B.10})$$

Building low-cost capacity (renewables or nuclear) lowers the price by less than the capacity increase if fossil plants still cover residual demand. A fuel-price change passes through in proportion to the fuel share of the price. Integration closes the gap between the national price and the union price $\bar{p}^E$, derived from a union-wide supply curve.¹²² The latter two impulses are special cases of the supply curve's response to its four exogenous arguments, set out next; the union price $\bar{p}^E$ instead clears a horizontally summed supply curve, derived below.

The supply curve (B.1) responds to four distinct exogenous shocks. Total differentiating with respect to E_j^D , $\bar{E}_j^0$, K_j^F , and p^F :

$$d \ln p_j^E = \varepsilon_j d \ln E_j^D - \varepsilon_j^0 d \ln \bar{E}_j^0 - \varepsilon_j^F d \ln K_j^F + \chi_j d \ln p^F$$

where $\varepsilon_j^0 \equiv \varepsilon_j \bar{E}_j^0 / E_j^D$ (low-cost shift coefficient), $\varepsilon_j^F \equiv (\chi_j / \zeta) u_j / (1 - u_j)$ (fossil shift coefficient), and $\varepsilon_j^0 + \varepsilon_j^F = \varepsilon_j$. The first term is movement *along* the supply curve; the remaining terms are *shifts* of the curve.

Shock	Supply-shift term	Coefficient	Note
Proportional ($d \ln \bar{E}_j$)	$-\varepsilon_j d \ln \bar{E}_j$	ε_j	Preserves mix
Low-cost ($d \ln \bar{E}_j^0$)	$-\varepsilon_j^0 d \ln \bar{E}_j^0$	$\varepsilon_j^0 = \varepsilon_j \bar{E}_j^0 / E_j^D$	Displaces fossil
Fossil ($d \ln K_j^F$)	$-\varepsilon_j^F d \ln K_j^F$	$\varepsilon_j^F = \varepsilon_j E_j^F / E_j^D$	Spreads fleet
Fuel price ($d \ln p^F$)	$+\chi_j d \ln p^F$	$\chi_j = (p_j^E - c_0^F) / p_j^E$	Pass-through ≤ 1

¹²²In the quantitative model version, a further channel operates through market coupling: when transmission between neighbours binds, a country with abundant low-cost capacity can inherit part of the price spike from a fuel-exposed neighbour through cross-border flows. This interaction between domestic supply and cross-border transmission motivates the Norwegian empirical design in Section 4, where bidding zones within the same country face identical fuel prices but differ in their cross-border exposure.

Fuel price pass-through. From (B.1), $\partial p_j^E / \partial p^F = h_{\min}(1 - u_j)^{-1/\zeta} = (p_j^E - c_0^F) / p^F$ at fixed u_j . Converting to log changes:

$$d \ln p_j^E|_{\text{fuel}} = \chi_j d \ln p^F$$

When $c_0^F = 0$, $\chi_j = 1$ and the pass-through is 100% everywhere. When $c_0^F > 0$, $\chi_j < 1$ and varies across countries: countries with higher u_j have higher χ_j (more of the price is fuel-driven), generating heterogeneous exposure to a common fuel shock even before GE responses. In the symmetric 2-country case with $c_0^F = 0$, the fuel shock drops from the differenced system ($\Delta \phi d \ln p^F = 0$): no reallocation. With $c_0^F > 0$ and asymmetric u_j , the initial shift itself is asymmetric ($\phi_1 \neq \phi_2$), introducing a first channel of differential exposure; a second channel operates through the GE demand response when $\varepsilon_1 \neq \varepsilon_2$. Both channels work in the same direction: the country with higher fossil dependence absorbs more of the fuel price shock. This is central to the Norwegian identification.

Proportional expansion. A proportional capacity expansion preserves $\bar{E}_j^0 / \bar{E}_j$ and $K_j^F / \bar{E}_j$ but does not preserve u_j : more capacity serves given demand, so u_j falls. Low-cost expansion reduces u_j and additionally changes ε_j (through χ_j and the composition of capacity). Fossil expansion also reduces u_j but does not change $\bar{E}_j^0$.

B.7 Endogenous comparative advantage

Price differences do not move all sectors alike; they sort them. A country whose power is cheap attracts entry, and entry bids up its wage. Electricity-intensive firms gain more from the cheap power than they lose to the higher wage, while labour-intensive firms are the other way round, so the two sectors move in opposite directions. The country ends up specialised in whatever its electricity price makes it good at, and the pattern is generated by the model rather than assumed through cross-country productivity differences.

Proposition 5 (Endogenous specialisation, two-country case). *Take two countries and two sectors, with common (ψ, γ) and $\psi > \gamma$, and write $\Delta x \equiv \ln x_1 - \ln x_2$. If country 1's productivity advantage is common across sectors ($\Delta \bar{\Phi}_H = \Delta \bar{\Phi}_L$), the relative concentration of the two sectors is*

$$\Delta \ln M_H - \Delta \ln M_L = \frac{-(\alpha_H - \alpha_L) \Delta \ln w - (\beta_H - \beta_L) \Delta \ln p^E}{\psi - \gamma}.$$

If in addition the residual-input share is common across sectors ($\kappa_H = \kappa_L$, so that $\alpha_H - \alpha_L = -(\beta_H - \beta_L)$) and country 1 holds a relative electricity-cost advantage ($\Delta \ln w - \Delta \ln p^E > 0$), the electricity-intensive sector H ($\beta_H > \beta_L$) is more concentrated in country 1:

$$\Delta \ln M_H - \Delta \ln M_L = \frac{\beta_H - \beta_L}{\psi - \gamma} (\Delta \ln w - \Delta \ln p^E) > 0.$$

The logic echoes Heckscher–Ohlin, but the “endowment” driving the pattern is the electricity price, and the supply side showed what builds it. Comparative advantage in electricity-intensive production has a natural component based on natural potential θ_j^R , fixed by geography, and a policy component based on installed low-cost capital K_j^R and K_j^N . When cheap power reflects a country’s natural resources, the sorting sends electricity-intensive production somewhere it genuinely belongs. When it reflects policy, the same sorting is misallocation. Section 3.6 turns that distinction into aggregate productivity. The cross-sector gradient is also the only part of the response the data can see. A price movement common to a location is absorbed by location–time fixed effects, and what survives is the differential across sectors of different electricity intensity, with slope $(\beta_H - \beta_L)/(\psi - \gamma)$ — the object Section 4 estimates.

The restriction that country 1’s productivity advantage is common across sectors is the theoretical counterpart of that identifying assumption: what must be ruled out is a location–sector productivity advantage that lines up with electricity intensity. Dropping it adds $(\Delta\bar{\Phi}_H - \Delta\bar{\Phi}_L)/(\psi - \gamma)$ to the right-hand side, which is indistinguishable from the electricity term in any cross-sector comparison. The quantitative model imposes no such restriction, carrying the full matrix $\bar{\Phi}_{js}$ and computing the multi-sector sorting pattern directly.

Proof of Proposition 5. Consider the restricted two-country, two-sector case in the proposition. Let $\Delta x \equiv \ln x_1 - \ln x_2$. Taking country differences in the PE location equation gives, for each sector s ,

$$\Delta \ln M_s = \frac{1}{\psi - \gamma} \left[-\alpha_s \Delta \ln w - \beta_s \Delta \ln p^E \right].$$

Subtracting the low-electricity-intensity sector L from the high-electricity-intensity sector H yields the general expression

$$\Delta \ln M_H - \Delta \ln M_L = \frac{-(\alpha_H - \alpha_L) \Delta \ln w - (\beta_H - \beta_L) \Delta \ln p^E}{\psi - \gamma}.$$

A sector’s two location-relevant shares can differ across sectors both because electricity substitutes for labour and because the residual input takes a different share, and only the first of these sorts sectors on the electricity price. Imposing a common residual share, $\alpha_s = 1 - \kappa - \beta_s$ with $\kappa_H = \kappa_L = \kappa$, gives $\alpha_H - \alpha_L = -(\beta_H - \beta_L)$ and

$$\Delta \ln M_H - \Delta \ln M_L = \frac{\beta_H - \beta_L}{\psi - \gamma} \left(\Delta \ln w - \Delta \ln p^E \right).$$

Since $\beta_H > \beta_L$, $\psi > \gamma$, and country 1 has a relative electricity-cost advantage $\Delta \ln w - \Delta \ln p^E > 0$, the expression is positive. The same expression shows linear monotonicity in the intensity spread and in the relative electricity-cost advantage. $\square$

B.8 Multi-sector general equilibrium

With $S > 1$ sectors, the system has $J(S + 2)$ unknowns: $\{M_{js}\}_{j,s}$, $\{w_j\}_j$, $\{p_j^E\}_j$. Lemma 3 states it: the location equation holds separately for each (j, s) with that sector's (α_s, β_s) , while each clearing condition sees the average intensity of the factor it clears, labour shares $\omega_{js}^L = L_{js}/\bar{L}_j$ for the wage and electricity shares $\omega_{js}^E = E_{js}/E_j^D$ for the price. This subsection records the two reductions that deliver the country objects λ_j and Θ_j of Proposition 2, and the terms they set aside.

A closed-form for individual $d \ln M_{js}$ does not exist because sectors interact through the shared factor markets: a capacity shock that attracts electricity-intensive firms also raises wages, pushing labour-intensive firms toward the other country. This cross-sector linkage generates the endogenous comparative advantage in Proposition 5. The system has a block-recursive fixed-point structure — given prices, location shares are determined sector-by-sector; given masses, prices are determined country-by-country — which is exploited by the numerical solution procedure described in Section 5.3. In the multi-sector model, local uniqueness and convergence are not implied by the scalar PE condition; they require the Jacobian of the full fixed-point mapping to satisfy the usual nonsingularity or contraction condition, which is checked numerically in the calibrated model.

The wage channel and sector-level responses. The absence of a closed form is a statement about the joint system rather than about the individual response: conditional on the country's wage response, each sector's response is exact. Write $\lambda_j \equiv d \ln w_j / d \ln p_j^E$ for the proportional wage response to an electricity price change in country j . Dividing the location equation of Lemma 3 by the partial-equilibrium response $-\beta_s d \ln p_j^E / (\psi - \gamma)$, the response of sector s is that multiplier scaled by

$$1 + \frac{\alpha_s}{\beta_s} \lambda_j,$$

so one country-level object, applied to each sector's labour-to-electricity cost ratio, delivers the whole cross-section of location responses. Because $\lambda_j < 0$, the response is damped, monotonically in α_s / β_s , and reverses sign for sectors with $\alpha_s / \beta_s > 1/|\lambda_j|$: for sufficiently labour-intensive sectors the wage relief released by the departing electricity-intensive firms exceeds their own electricity saving. This is the multi-sector counterpart of Proposition 5.

Reduction 1: labour clearing pins λ_j . Log-linearised labour demand in cell (j, s) is $d \ln L_{js} = [1 + (\sigma - 1)\gamma] d \ln M_{js} - [1 + \alpha_s(\sigma - 1)] d \ln w_j - \beta_s(\sigma - 1) d \ln p_j^E$. Summing over sectors with the labour shares ω_{js}^L , imposing $\sum_s \omega_{js}^L d \ln L_{js} = 0$, and substituting the location equation for $d \ln \bar{M}_j^L$ gives

$$\lambda_j = -\mathcal{A} \left[\bar{\beta}_j^L + \bar{\alpha}_j^L \lambda_j \right], \quad \text{that is} \quad \lambda_j = -\frac{\mathcal{A} \bar{\beta}_j^L}{1 + \mathcal{A} \bar{\alpha}_j^L},$$

collecting terms with $[1 + (\sigma - 1)\gamma] + (\sigma - 1)(\psi - \gamma) = \mathcal{A}(\psi - \gamma)$, where $\mathcal{A} = [1 + (\sigma - 1)\psi] / (\psi - \gamma)$ is the activity multiplier of Proposition 2. The denominator is Ω_j evaluated at $\varepsilon_j = 0$ with the employment-weighted labour intensity $\bar{\alpha}_j^L$ in place of the single-sector α : with the electricity price held fixed, only wage congestion resists the move. Since $\psi > \gamma$, the expression implies $|\lambda_j| < \bar{\beta}_j^L / \bar{\alpha}_j^L$: the wage offset cannot exceed the country's average electricity-to-labour cost ratio, which is why λ_j is small in the calibrated sample even though the location multiplier is large.

Reduction 2: electricity clearing pins Θ_j and the dampening. Log-linearised electricity demand in cell (j, s) is $d \ln E_{js} = [1 + (\sigma - 1)\gamma] d \ln M_{js} - \alpha_s(\sigma - 1) d \ln w_j - [1 + \beta_s(\sigma - 1)] d \ln p_j^E$. Summing with the electricity shares ω_{js}^E , substituting the location equation and $d \ln w_j = \lambda_j d \ln p_j^E$, and collecting the same activity multiplier gives the country's electricity demand elasticity

$$\Theta_j \equiv -\frac{d \ln E_j^D}{d \ln p_j^E} = \underbrace{1}_{\text{substitution}} + \underbrace{\mathcal{A}}_{\text{activity multiplier}} \underbrace{\sum_s \omega_{js}^E [\beta_s + \alpha_s \lambda_j]}_{\text{average net exposure}} = 1 + \mathcal{A} [\bar{\beta}_j^E + \bar{\alpha}_j^E \lambda_j].$$

The two reductions are therefore the same construction under different weights: each applies the activity multiplier to the country's average net exposure, the bracket that appears sector by sector in Proposition 2, averaged by the factor whose market is being cleared. What differs is the shape of supply. Labour supply is vertical, so the wage moves until the country's labour demand response is zero and no substitution term survives; electricity supply slopes, so Θ_j retains the direct substitution of firms away from electricity at a given level of activity. Combining Θ_j with the supply curve $d \ln p_j^E = \varepsilon_j d \ln E_j^D + d \ln p_j^E|_{E^D}$ gives

$$d \ln p_j^E = \frac{d \ln p_j^E|_{E^D}}{1 + \varepsilon_j \Theta_j},$$

the dampening factor of the proposition: the supply curve turns demand into price at rate ε_j , demand answers the price with elasticity Θ_j , and the loop shrinks the impulse by $1 + \varepsilon_j \Theta_j$.

Stability: the sign of the dampening. The loop is well behaved as long as the factor $1/(1 + \varepsilon_j \Theta_j)$ is positive, which requires $\Theta_j > -1/\varepsilon_j$, and it damps rather than amplifies, $1/(1 + \varepsilon_j \Theta_j) \leq 1$, exactly when $\Theta_j \geq 0$. Both turn on the sign of average net exposure, since $\Theta_j = 1 + \mathcal{A} \bar{n}_j^E$ with $\bar{n}_j^E \equiv \bar{\beta}_j^E + \bar{\alpha}_j^E \lambda_j$ and $\mathcal{A} > 0$ under the partial-equilibrium condition $\psi > \gamma$. Individual sectors can enter that average negatively: those with $\alpha_s/\beta_s > 1/|\lambda_j|$ gain more from the wage relief than they lose on their own electricity bill, and move toward the country whose power became dearer. That is the sorting result of Proposition 5, not an instability, and it does not carry over to the average.

Substituting λ_j from Reduction 1,

$$\bar{n}_j^E = \frac{\bar{\beta}_j^E + \mathcal{A} \left(\bar{\beta}_j^E \bar{\alpha}_j^L - \bar{\alpha}_j^E \bar{\beta}_j^L \right)}{1 + \mathcal{A} \bar{\alpha}_j^L},$$

whose denominator is positive. The bracket in the numerator is non-negative for the model's own weights. Cell electricity and labour demand are $E_{js} = \beta_s(\sigma - 1)R_{js}/p_j^E$ and $L_{js} = \alpha_s(\sigma - 1)R_{js}/w_j$ with $R_{js} = M_{js}\bar{\pi}_{js}$, so $\omega_{js}^E \propto \beta_s R_{js}$ and $\omega_{js}^L \propto \alpha_s R_{js}$, and

$$\bar{\beta}_j^E \bar{\alpha}_j^L - \bar{\alpha}_j^E \bar{\beta}_j^L \propto \left(\sum_s \beta_s^2 R_{js} \right) \left(\sum_s \alpha_s^2 R_{js} \right) - \left(\sum_s \alpha_s \beta_s R_{js} \right)^2 \geq 0$$

by the Cauchy–Schwarz inequality, with equality only if β_s/α_s is common to all sectors. Hence $\bar{n}_j^E > 0$ and $\Theta_j > 1$: in the analytical model the partial-equilibrium condition $\psi > \gamma$ is all that is required, and general equilibrium always damps. The economic content of the inequality is that the sectors holding most of a country's electricity demand are, by construction of the weights, more electricity-intensive than those holding most of its employment, so the wage offset cannot outweigh the electricity saving in the aggregate even where it does sector by sector.

The guarantee is a statement about the analytical model. The quantitative model departs from it in three ways: residual load H_j is inelastic, so the substitution term becomes the modelled share of load $m_j < 1$ rather than one; the sectoral demand shifter B_s is endogenous; and cross-border trade makes the price a country faces depend on its neighbours. The first is innocuous, since $\Theta_j = m_j \left[1 + \mathcal{A} \bar{n}_j^E \right]$ under firm-demand weights and remains positive. The other two enter through λ_j and are not covered by the argument above, so the solver evaluates Θ_j and the local stability of the full fixed point at each counterfactual, and reports them in Section 5.

Two objects are set aside in both reductions, both of which the quantitative model carries. The first is the sector log-sum $d\bar{V}_s$, which is negligible only for countries that are small within a sector. The second is the sectoral demand shifter B_s , held fixed in the analytical model but endogenous in the quantitative one (Appendix B.9). The expression therefore accounts for the magnitude of λ_j , delivering values near -0.03 throughout the calibrated sample, but not for its variation across countries; the values reported in Section 5.5 are solved rather than computed from it.

B.9 The sectoral demand shifter and the goods-market closure

From CES demand with optimal markup pricing, operating profit is $\pi_{fjs} = r_{fjs}/\sigma$, where revenue is

$$r_{fjs} = \left(\frac{\sigma}{\sigma - 1} \right)^{1-\sigma} c_{js}^{1-\sigma} \phi_{js}^{\sigma-1} P_s^{\sigma-1} Y_s$$

with the unit cost of the three-input Cobb–Douglas technology

$$c_{js} = \frac{w_j^{\alpha_s} (p_j^E)^{\beta_s} (p^X)^{\kappa_s}}{\alpha_s^{\alpha_s} \beta_s^{\beta_s} \kappa_s^{\kappa_s}}, \quad \kappa_s \equiv 1 - \alpha_s - \beta_s.$$

Collecting terms that are common across all locations j within sector s , including the residual-input price p^X and the Cobb–Douglas cost constant:

$$B_s \equiv \frac{1}{\sigma} \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} P_s^{\sigma-1} Y_s (\alpha_s^{\alpha_s} \beta_s^{\beta_s} \kappa_s^{\kappa_s})^{\sigma-1} (p^X)^{-\kappa_s(\sigma-1)}$$

Since B_s enters every firm's location value identically and does not vary across countries j , it cancels from all pairwise location comparisons: logit shares $m_{js} = \exp(v_{js}/\psi) / \sum_k \exp(v_{ks}/\psi)$ are invariant to additive constants in v_{js} . In the analytical model, where B_s is held constant, it therefore has no effect on the firm distribution or the comparative statics, even though it is endogenous through P_s and Y_s .

In the quantitative model the goods market is closed explicitly, following Caliendo and Parro (2015) and Arkolakis and Walsh (2023). A representative consumer has Cobb–Douglas preferences across sectors with fixed expenditure shares μ_s and CES preferences over varieties within each sector, sold in an integrated market, so sectoral expenditure is $Y_s = \mu_s I$ with I aggregate income, and the within-sector price index is

$$P_s = \left[\sum_j M_{js} \left(\frac{\sigma}{\sigma-1} \frac{c_{js}}{\bar{\phi}_{js}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad \bar{\phi}_{js} = \exp(\bar{\Phi}_{js} + \gamma \ln M_{js}).$$

Aggregate income I is labour income, plus the producer and congestion rents of electricity generation net of the imported-fuel cost of fossil generation (the resource-cost term in (14)), plus a residual standing for payments to capital and intermediate inputs. The residual load H_j pays the same wholesale price as firms, but its bill is a transfer from households and non-modelled users to generators with no counterpart in the goods-expenditure base, so changes in that bill are netted out of producer rents rather than counted as income. The model solves for labour and electricity but has no capital or materials block, while the expenditure base is the gross output of the modelled sectors, so that residual is set to make baseline income equal baseline expenditure exactly and is held fixed across counterfactuals. Income responds to a counterfactual only through wages, electricity rents, and the fuel bill, and its response is damped relative to a closure in which capital and intermediate income moved with revenue. Baseline income is normalised to one, so all goods-market objects are reported as ratios to their baseline values.

With P_s and Y_s endogenous, B_s varies along any counterfactual path. It still cancels from the within-sector logit shares *conditional on factor prices*, but the equilibrium wages and electric-

ity prices depend on B_s through factor-market clearing, so — unlike in the analytical model — B_s affects the equilibrium firm distribution. The equilibrium is the joint fixed point in $\{M_{js}, w_j, p_j^E, P_s, I\}$, and holding B_s at its baseline value recovers the analytical case and is reported as a special-case comparison.

B.10 Aggregate productivity decomposition: derivation and exact objects

I measure aggregate productivity as real value added per unit of input. With Cobb–Douglas production and CES demand it is convenient to track the production block through the average-productivity object

$$\bar{A}_s \equiv \sum_j \omega_{js} \ln \phi_{js}, \quad \ln \phi_{js} \equiv \bar{\Phi}_{js} + \gamma \ln M_{js}, \quad \omega_{js} \equiv \frac{M_{js} \bar{\pi}_{js}}{\sum_k M_{ks} \bar{\pi}_{ks}},$$

the revenue-weighted mean of location log-productivities $\ln \phi_{js}$, with μ_s sectoral expenditure weights and ω_{js} output weights.¹²³ Real value added departs from average productivity in two ways: CES demand extracts extra output from the dispersion of marginal costs — an allocative-efficiency gain — and electricity is generated from imported fuel — a resource cost. Aggregate TFP therefore has four channels: agglomeration and composition, carried by $\bar{A}_s$; allocative efficiency; and fossil resource use. Each is derived below as the exact object the calibrated model evaluates, with the leading-order form reported in the main text.

Agglomeration and composition. The change in average productivity is exactly the Olley–Pakes within/between split

$$\Delta \bar{A}_s = \underbrace{\sum_j \omega_{js} \gamma \Delta \ln M_{js}}_{\text{within}} + \underbrace{\sum_j \Delta \omega_{js} \ln \phi'_{js}}_{\text{between}},$$

with $\ln \phi'_{js}$ evaluated at the counterfactual allocation and $\Delta \ln \phi_{js} = \gamma \Delta \ln M_{js}$ since $\bar{\Phi}_{js}$ is a fixed primitive. The main text does not read its two channels off this split. The Olley–Pakes terms divide the change by *what moved* — levels at fixed weights, then weights at fixed levels — and leave an interaction $\sum_j \Delta \omega_{js} \gamma \Delta \ln M_{js}$ that belongs to neither. Because $\bar{\Phi}_{js}$ is fixed and $\ln \phi_{js}$ has only two components, dividing instead by *where the productivity comes from* is exact and exhaustive, $\Delta \bar{A}_s = \Delta [\sum_j \omega_{js} \bar{\Phi}_{js}] + \gamma \Delta [\sum_j \omega_{js} \ln M_{js}]$, with no residual. This is the decomposition of (14). Composition, the first term, vanishes only in the special case of fundamentals common across countries, since then $\sum_j \Delta \omega_{js} \bar{\Phi}_s = \bar{\Phi}_s \sum_j \Delta \omega_{js} = 0$. Agglomeration, the second, contains all three Olley–Pakes pieces: clusters growing at fixed weights, output weight moving toward already-large clusters,

¹²³The base productivity $\bar{\Phi}_{js}$ is the model’s country–sector productivity shifter, recovered by inverting the baseline so that the calibrated equilibrium reproduces observed firm locations exactly (Section 5); the composition term uses it in full. An external structural-business-statistics measure is used only to validate $\bar{\Phi}_{js}$, not to construct it; decomposing it into a measured-productivity component and a residual is a planned refinement.

and their interaction. That is the correct grouping because they are not separable events — the location first-order condition gives $\Delta\omega_{js} = \omega_{js} [1 + (\sigma - 1)\psi] (\Delta \ln M_{js} - \sum_k \omega_{ks} \Delta \ln M_{ks})$, so output weight moves only because firm mass moves. Writing $\Delta_{js} \equiv \ln M'_{js} - \ln M_{js}$, $\Lambda_\gamma \equiv 1 + (\sigma - 1)\gamma$, with weights moving along $\omega_{js}(t) = \omega_{js} e^{\Lambda_\gamma t \Delta_{js}} / \sum_k \omega_{ks} e^{\Lambda_\gamma t \Delta_{ks}}$ and moments taken under initial weights,

$$\begin{aligned} \Delta \bar{A}_s &= \gamma E_s[\Delta_{js}] + \Lambda_\gamma \text{Cov}_s(\ln \phi_{js}, \Delta_{js}) + \gamma \Lambda_\gamma \text{Var}_s(\Delta_{js}) \\ &\quad + \frac{\Lambda_\gamma^2}{2} E_s[(\ln \phi_{js} - E_s \ln \phi_{js})(\Delta_{js} - E_s \Delta_{js})^2] + O(\|\Delta\|^3). \end{aligned}$$

The first term is the fixed-weight within expression; the rest is the weight revaluation and is not sign-definite — the covariance term depends on whether reallocation favours initially high-productivity or high-agglomeration locations, and the cubic term on whether large reallocations fall where $\ln \phi_{js}$ is above or below average. A non-negative correction obtains only under restrictions, e.g. $\ln \phi_{js}$ initially equal within a sector with zero covariance, leaving $\gamma \Lambda_\gamma \text{Var}_s(\Delta_{js}) \geq 0$. The calibrated model evaluates $\Delta \bar{A}_s$ exactly; aggregate agglomeration and composition are $\sum_s \mu_s \Delta \bar{A}_s$.

Allocative efficiency. Real value added exceeds average productivity by the output CES demand extracts from cost dispersion. The minimum cost of one unit of the sector- s bundle is the marginal-cost index

$$\tilde{P}_s = \left[\sum_j M_{js} \phi_{js}^{\sigma-1} c_{js}^{1-\sigma} \right]^{1/(1-\sigma)}, \quad c_{js} \propto w_j^{\alpha_s} (p_j^E)^{\beta_s},$$

where the omitted factor of proportionality collects the Cobb–Douglas cost constant and the common residual-input price $(p^X)^{\kappa_s}$, which are identical across countries and so shift the index by a level without entering any cross-country moment below. Real output per unit cost is then $-\ln \tilde{P}_s = \frac{1}{\sigma-1} \ln \sum_j e^{\nu_{js}}$, $\nu_{js} \equiv \Lambda_\gamma \ln M_{js} - (\sigma - 1) \ln c_{js}$, with revenue weights $\omega_{js} = e^{\nu_{js}} / \sum_k e^{\nu_{ks}}$. Between any two equilibria the exact change is the revenue-weighted CES mean

$$\Delta(-\ln \tilde{P}_s) = \frac{1}{\sigma-1} \ln \left(\sum_j \omega_{js} \exp \{ \Lambda_\gamma \Delta \ln M_{js} - (\sigma - 1) \Delta \ln c_{js} \} \right),$$

base productivity entering only through the initial weights. The log-sum-exp has gradient ω_{js} and Hessian $\text{diag}(\omega) - \omega\omega'$, so to second order it equals $\frac{\Lambda_\gamma}{\sigma-1} E_\omega[\Delta \ln M_{js}] - E_\omega[\Delta \ln c_{js}] + \frac{1}{2(\sigma-1)} \text{Var}_\omega(\Lambda_\gamma \Delta \ln M_{js} - (\sigma - 1) \Delta \ln c_{js})$. The agglomeration content of the first term $(\Lambda_\gamma / (\sigma - 1) = \gamma + 1 / (\sigma - 1))$ coincides with the within term above; isolating the cost-dispersion part of the variance — dropping the mass term and the mass–cost covariance — leaves the allocative-efficiency term

$$\frac{\sigma-1}{2} \text{Var}_\omega(\Delta \ln c_{js}),$$

the curvature $(\sigma - 1)$ fixing the coefficient in (14). The leftover love-of-variety term $\frac{1}{\sigma-1}E_\omega[\Delta \ln M_{js}]$ and mean factor-cost change $-E_\omega[\Delta \ln c_{js}]$ are set aside next.

Why the love-of-variety and factor-cost terms are set aside. The two leftover terms are deliberately kept out of the main-text decomposition. The first is the love-of-variety term $\frac{1}{\sigma-1}E_\omega[\Delta \ln M_{js}]$. With the total sector mass fixed ($\sum_j M_{js} = M_s$), it reflects the revenue-weighted reallocation of varieties across countries and operates through the sectoral price index P_s — the consumer-welfare channel noted in the discussion of B_s above. It raises consumer real income but not output per unit of input, so it belongs to welfare rather than to the production-side TFP measure that is the object of this paper; the quantitative model tracks it separately through P_s . The second is the mean factor-cost change $-E_\omega[\Delta \ln c_{js}] = -E_\omega[\alpha_s \Delta \ln w_j + \beta_s \Delta \ln p_j^E]$. Its labour component, and the part of its electricity component that reflects domestic generation rents, are within-economy transfers that net against the input-price deflator and leave real value added unchanged. The residual — the imported-fuel component of $\Delta \ln p_j^E$ — is a genuine resource saving: lower fossil dispatch reduces fuel imports and raises real value added. This is a level (resource-use) channel rather than a spatial-reallocation channel, so it is kept out of the reallocation decomposition (14) but retained as the fossil-use term derived next, which enters the union and comprehensive-coordination results (Proposition 3) and the quantitative decomposition of Section 5.

The fossil-use term. This term follows from value-added accounting. Consolidating production firms with the domestic generation sector, all electricity payments between them are internal transfers and cancel; the only cross-border cost is imported fuel, so aggregate real value added $\mathcal{Y}$ is production-sector real output minus the fuel bill $\sum_j F_j$, with $F_j = p^F Q_j^F$. For the merit-order fleet of Section 3 — heat rate $h(k) = h_{\min}(1 - k)^{-1/\zeta}$, capacity K_j^F , utilisation u_j , fossil generation $E_j^F = u_j K_j^F$ — the total heat input is

$$Q_j^F = K_j^F \int_0^{u_j} h(k) dk = K_j^F h_{\min} \frac{\zeta}{\zeta - 1} [1 - (1 - u_j)^{(\zeta-1)/\zeta}].$$

Holding the world fuel price p^F fixed, the change in the fuel bill is $\Delta F_j = p^F h(u_j) \Delta E_j^F$, and the marginal plant's fuel cost equals the wholesale price net of the non-fuel component, $p^F h(u_j) = p_j^E - c_0^F = \chi_j p_j^E$, by merit-order pricing. Dividing by $\mathcal{Y}$ gives the fossil-use term

$$-\sum_j \frac{\Delta F_j}{\mathcal{Y}} = -\sum_j s_j^F \Delta \ln E_j^F, \quad s_j^F \equiv \frac{(p_j^E - c_0^F) E_j^F}{\mathcal{Y}} = \frac{\chi_j p_j^E E_j^F}{\mathcal{Y}}.$$

This is the imported-fuel residual of the mean factor-cost term: of each unit firms pay for electricity, the domestic generation rents and the non-fuel cost c_0^F are transfers that net out, leaving only the fuel share χ_j as a real economy-wide cost. The term values displaced generation at the

marginal fossil plant's fuel cost $p_j^E - c_0^F$, consistent with the merit-order pricing and the first-order approximation used throughout; the exact resource saving integrates the cheaper fuel of infra-marginal plants through Q_j^F above and coincides with the marginal expression to first order. The weight s_j^F is largest in fossil-dependent, high-price systems with high χ_j , so the resource saving from displacing fossil is concentrated there.

Output weights and the composition term. The composition term in (14) requires the response of the output weights. Country j 's share of sector- s output is $\omega_{js} = M_{js}\bar{\pi}_{js} / \sum_k M_{ks}\bar{\pi}_{ks}$, since revenue is a constant multiple σ of operating profit. Differentiating average profit, $d \ln \bar{\pi}_{js} = (\sigma - 1)\gamma d \ln M_{js} - A_s d \ln w_j - B_s d \ln p_j^E$, and substituting the location equation of Proposition 1 in the form $-\alpha_s d \ln w_j - \beta_s d \ln p_j^E = (\psi - \gamma) d \ln M_{js}$ gives

$$d \ln \bar{\pi}_{js} = (\sigma - 1) [\gamma + (\psi - \gamma)] d \ln M_{js} = (\sigma - 1)\psi d \ln M_{js},$$

so cell revenue moves as $d \ln (M_{js}\bar{\pi}_{js}) = [1 + (\sigma - 1)\psi] d \ln M_{js}$. Because $\bar{\Phi}_{js}$ is a fixed fundamental, it does not contribute to this differential. Hence

$$\begin{aligned} d\omega_{js} &= \omega_{js} [1 + (\sigma - 1)\psi] (d \ln M_{js} - \sum_k \omega_{ks} d \ln M_{ks}), \\ \sum_j d\omega_{js} \bar{\Phi}_{js} &= [1 + (\sigma - 1)\psi] \text{Cov}_{\omega_{js}} (d \ln M_{js}, \bar{\Phi}_{js}), \end{aligned}$$

using $\sum_j \omega_j (x_j - \bar{x}_\omega) y_j = \text{Cov}_\omega(x, y)$. The scale factor $1 + (\sigma - 1)\psi$ is exactly the ratio δ_Π / δ_M of the total-profit to the firm-count moment in Appendix B.3: output weight responds both to the number of firms and to the scale of each.

The same expression shows that the reallocation channels do not require firms to move. The weight response is the product $\Lambda d \ln M_{js}$, and its two parts, $[1 + (\sigma - 1)\gamma]$ and $(\sigma - 1)(\psi - \gamma)$, are the output that arrives with the firms and the scale gain of the firms already in place as their costs fall. Hold the firms still by letting attachment grow without bound: $d \ln M_{js} \rightarrow 0$ as $\psi \rightarrow \infty$, but Λ grows at the same rate, and along the location equation the product converges to

$$\Lambda d \ln M_{js} \longrightarrow -(\sigma - 1) d \ln c_{js},$$

the response of a location's expenditure share to its own unit cost at a fixed set of producers. Composition, the weight-shift part of agglomeration, and allocative efficiency all survive that limit; only the in-place agglomeration term $\gamma \mathbb{E}_{\omega_{js}} d \ln M_{js}$ vanishes, since no cluster grows. What the three channels require is that relative costs change, not that production relocates; firm mobility governs how large the weight movement is, not whether it happens.

Electricity-price specialisations. The forms in (B.11) and (B.12) follow by substitution into (14). For (B.11), the partial-equilibrium response $\Delta \ln M_{js} = -\beta_s \Delta \ln p_j^E / (\psi - \gamma)$ enters the composition and agglomeration terms (with $\Delta \ln p_j^E$ net of the sector mean, since $\sum_j M_{js} = M_s$ implies only relative price changes reallocate), and at fixed wages $\Delta \ln c_{js} = \beta_s \Delta \ln p_j^E$ enters the allocative term, giving, with the expenditure scale $\Lambda \equiv 1 + (\sigma - 1)\psi = \mathcal{A}(\psi - \gamma)$, the coefficients $-\Lambda\beta_s/(\psi - \gamma)$ on composition, $-\gamma\beta_s/(\psi - \gamma)$ on the mean and $-\gamma[1 + (\sigma - 1)\psi]\beta_s/(\psi - \gamma)$ on the covariance in agglomeration, and $\frac{(\sigma-1)\beta_s^2}{2}$ on the allocative term. For (B.12), the location first-order condition gives $\Delta \ln c_{js} = -(\psi - \gamma)\Delta \ln M_{js}$ — the bracket in Proposition 1 is exactly $-\Delta \ln c_{js}$ — while the low-cost expansion case of Proposition 2 gives $\hat{m}_j \equiv \Delta \ln M_j = (\beta_j^0 / [(\psi - \gamma)\Omega_j]) \Delta \ln \bar{E}_j^0$. Substituting both, the allocative term becomes $\frac{(\sigma-1)(\psi-\gamma)^2}{2} \text{Var}_{\omega_j}(\hat{m}_j)$, the agglomeration term $\gamma\{\mathbb{E}_{\omega_j}[\hat{m}_j] + \Lambda \text{Cov}_{\omega_j}(\hat{m}_j, \ln M_j)\}$, and the composition term $\Lambda \text{Cov}_{\omega_j}(\hat{m}_j, \Phi_j)$, so the three spatial channels are the first three output-weighted moments of the same firm-mass response.

$$\begin{aligned} \Delta \ln \text{TFP} \approx & - \sum_s \mu_s \frac{\Lambda \beta_s}{\psi - \gamma} \text{Cov}_{\omega_{js}}(\Delta \ln p_j^E, \Phi_{js}) \\ & - \sum_s \mu_s \frac{\gamma \beta_s}{\psi - \gamma} \left[\sum_j \omega_{js} \Delta \ln p_j^E + \Lambda \text{Cov}_{\omega_{js}}(\Delta \ln p_j^E, \ln M_{js}) \right] \\ & + \sum_s \mu_s \frac{(\sigma - 1) \beta_s^2}{2} \text{Var}_{\omega_{js}}(\Delta \ln p_j^E) - \sum_j s_j^F \Delta \ln E_j^F. \end{aligned} \quad (\text{B.11})$$

$$\begin{aligned} \Delta \ln \text{TFP}^{NP} \approx & \underbrace{\Lambda \text{Cov}_{\omega_j}(\hat{m}_j, \Phi_j)}_{\text{composition}} + \underbrace{\gamma \left\{ \mathbb{E}_{\omega_j}[\hat{m}_j] + \Lambda \text{Cov}_{\omega_j}(\hat{m}_j, \ln M_j) \right\}}_{\text{agglomeration}} \\ & + \underbrace{\frac{(\sigma - 1)(\psi - \gamma)^2}{2} \text{Var}_{\omega_j}(\hat{m}_j)}_{\text{allocative efficiency}} - \underbrace{\sum_j s_j^F \Delta \ln E_j^F}_{\text{fossil resource use}}, \end{aligned} \quad (\text{B.12})$$

where $\mathbb{E}_{\omega_j}$, Cov_{ω_j} and Var_{ω_j} denote moments across countries under the output weights ω_j .

B.11 Electricity union: market-clearing uniform price

Under the electricity union a uniform price $\bar{p}^E$ clears the integrated market, obtained by summing the country supply curves of Section 3 horizontally and intersecting the result with total demand. With heterogeneous fleets the pooled curve is piecewise and the union price is solved numerically along with the rest of the equilibrium. The symmetric case gives the intuition in closed form: when fleets share K_j^F and ζ , the pooled fleet has the same functional form as a national one, $\bar{p}^E = c_0^F + p^F h_{\min}(1 - \bar{u})^{-1/\zeta}$ at the pooled utilisation rate $\bar{u} = (\sum_j E_j^F) / (\sum_j K_j^F)$, and the price change

country j experiences is

$$\ln \bar{p}^E - \ln p_j^E \approx \frac{\chi_j}{\zeta} [\ln(1 - u_j) - \ln(1 - \bar{u})],$$

negative for countries running their fleets harder than the pool and positive for the rest.

What reallocates under a union. The union changes every country's price at once, so the comparative statics of Proposition 2, which trace a shift in one country's supply curve, are not the right instrument. The logit inverts in levels instead, and the result is exact in prices and wages. Collecting $\ln M_{js}$ from both sides of the log share, since it appears on the right through the agglomeration term,

$$\ln M_{js} = \frac{1}{\psi - \gamma} [\bar{\Phi}_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E] + \frac{\psi}{\psi - \gamma} (\ln M_s - \bar{V}_s), \quad \bar{V}_s \equiv \ln \sum_k \exp(v_{ks}/\psi), \quad (\text{B.13})$$

where the agglomeration feedback has folded into the location multiplier exactly as in Proposition 1. Differencing between the union allocation and the baseline, with $\bar{\Phi}_{js}$ and M_s fixed,

$$\Delta \ln M_{js} = -\frac{1}{\psi - \gamma} [\alpha_s \Delta \ln w_j + \beta_s \Delta \ln p_j^E] - \frac{\psi}{\psi - \gamma} \Delta \bar{V}_s, \quad \Delta \ln p_j^E = \ln \bar{p}^E - \ln p_j^{E,0}. \quad (\text{B.14})$$

The union price then cancels. Under a common price the electricity term of the systematic payoff carries no country subscript, so it leaves the inclusive value as a factor: writing $v'_{js} = \tilde{v}'_{js} - \beta_s \ln \bar{p}^E$ with $\tilde{v}'_{js} \equiv \bar{\Phi}_{js} + \gamma \ln M'_{js} - \alpha_s \ln w'_j$,

$$\bar{V}'_s = \ln \sum_k \exp\left(\frac{\tilde{v}'_{ks} - \beta_s \ln \bar{p}^E}{\psi}\right) = -\frac{\beta_s}{\psi} \ln \bar{p}^E + \ln \sum_k \exp(\tilde{v}'_{ks}/\psi).$$

The only $\bar{p}^E$ inside $\Delta \bar{V}_s$ is therefore $-\frac{\beta_s}{\psi} \ln \bar{p}^E$, which meets an equal and opposite copy carried by $\Delta \ln p_j^E$ in (B.14): the two cancel, and $\bar{p}^E$ disappears from the reallocation. Nothing of the kind happens at baseline, where $\ln p_k^{E,0}$ carries a country subscript and stays inside $\bar{V}_s^0$. That asymmetry is the whole experiment.

What remains is the sector normalisation, and it is the one step that is not exact. Since $\partial \bar{V}_s / \partial v_{ks} = m_{ks}/\psi$, a differential change satisfies $d\bar{V}_s = \frac{1}{\psi} \sum_k m_{ks} dv_{ks}$, and fixed sector mass, $\sum_k m_{ks} d\ln M_{ks} = 0$, removes the agglomeration term from it. For a discrete change the same expression holds at shares m_{ks}^* somewhere between the baseline and the union allocation; evaluating them at the baseline gives

$$\Delta \ln M_{js} = \frac{1}{\psi - \gamma} \left[\beta_s \left(\ln p_j^{E,0} - \overline{\ln p^{E,0}_s} \right) - \alpha_s \left(\Delta \ln w_j - \overline{\Delta \ln w_s} \right) \right], \quad \bar{x}_s \equiv \sum_k m_{ks}^0 x_k. \quad (\text{B.15})$$

Three things follow. Only baseline dispersion reallocates production: a country gains where its own price stood above its sector's baseline average, weighted by where that sector's firms were, and not where it stands relative to the union price, so a country whose electricity cost rises can still gain if its competitors' rise by more. The supply elasticities ε_j still matter, but only through the level of $\bar{p}^E$ that clears the pooled market, not through national dampening, which the union has disarmed by making the electricity price common. And wage congestion survives: national labour markets still clear, so $\Delta \ln w_j$ is solved jointly with (B.15) and absorbs part of the reallocation, which is why the second term appears in the same normalised form as the first.¹²⁴

Proposition 3 is then the decomposition of Lemma 4 evaluated at the reallocation (B.15), with the union price and wages that clear the integrated market. The allocative and fossil terms are signed, as a variance and as a resource saving when fossil dispatch falls; the two location terms are covariances and take the sign of the alignment between the reallocation and, respectively, the fundamentals $\bar{\Phi}_{js}$ and the initial cluster sizes $\ln M_{js}^0$.

Location invariance under a union. A single union-wide price does more than compress dispersion: it removes electricity from the location problem altogether.

Proposition 6 (Location invariance under a union). *Let every country clear at a common wholesale price $\bar{p}^E$. Then:*

- (i) *the term $\beta_s \ln \bar{p}^E$ does not vary across countries, so it cancels from the logit of Lemma 1: generation capacity reaches the firm distribution only through the level of $\bar{p}^E$;*
- (ii) *with endogenous sectoral price indices, the firm distribution is independent of $\bar{p}^E$ as well, and hence of the generation allocation entirely.*

Proof. Part (i) is immediate: $v_{js} = \bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln \bar{p}^E$, and a term constant across j within sector s cancels from the share $\exp(v_{js}/\psi) / \sum_k \exp(v_{ks}/\psi)$, exactly as the demand shifter B_s does. For part (ii), firm revenue relative to baseline carries the explicit factor $(\bar{p}^E)^{-\beta_s(\sigma-1)}$. Under a common electricity price every producer in sector s has marginal cost proportional to $(\bar{p}^E)^{\beta_s}$, so the sectoral price index satisfies $P_s \propto (\bar{p}^E)^{\beta_s}$ and the demand shifter $B_s = I P_s^{\sigma-1}$ contributes $(\bar{p}^E)^{\beta_s(\sigma-1)}$. The two factors cancel, leaving sectoral labour demand and hence equilibrium wages independent of $\bar{p}^E$; with wages unchanged, the logit in (i) is unchanged. $\square$

Part (ii) relies on the goods-market closure of Section 5.1, where P_s and I are endogenous; it does not hold in the analytical model, which fixes B_s . The invariance is also stated within a given calibration: it covers counterfactual capacity stocks and fuel- and carbon-price environments, but

¹²⁴Equations (B.13)–(B.14) and the cancellation are exact in prices and wages; only (B.15) approximates, by evaluating the inclusive-value weights at baseline shares rather than at the intermediate allocation, so its error grows with how far the shares themselves move. The expression is a decomposition rather than a closed form, since $\Delta \ln w_j$ is an equilibrium object.

not recalibrations that re-anchor the baseline supply curves, where the baseline economy itself differs. Table H.3 bears this out — every scenario clearing at a union-wide price reports identical composition, agglomeration, and allocative terms and identical $\max|\Delta M|$, across a six-fold range of union prices — and in the solver output the three location channels agree to ten significant figures between the union and the union combined with generation re-siting, with the entire difference falling on the fossil-resource term.

A country's baseline price itself separates into the part set by its natural potential and the part set by where capacity was built, which is the distinction Section 3.6 draws between dispersion the union can undo and dispersion generation policy created. Splitting $\ln p_j^{E,0}$ that way in (B.15) splits the reallocation, and hence each location channel, into a policy-correction component and a natural-potential compression component; Section 6 measures the two by re-siting generation with and without a union rather than from this expansion.

B.12 Alternative parameterisation: a sector-specific location multiplier

The baseline model assumes common ψ and common γ , so the location multiplier $1/(\psi - \gamma)$ is a single scalar and all cross-sector variation in the firm response runs through electricity intensity β_s . Two extensions make the multiplier sector-specific: sector-specific agglomeration γ_s (with common ψ), or sector-specific Gumbel dispersion ψ_s (with common γ), still measured in units of normalised log profit. Under the ψ_s variant the systematic location value $v_{js} = \bar{\Phi}_{js} + \gamma \ln M_{js} - \alpha_s \ln w_j - \beta_s \ln p_j^E$ is scaled by ψ_s rather than ψ inside the logit, and the PE location equation becomes

$$d \ln M_{js} = \frac{1}{\psi_s - \gamma} \left[-\alpha_s d \ln w_j - \beta_s d \ln p_j^E \right]$$

with sector-specific multiplier $1/(\psi_s - \gamma)$; the γ_s variant gives $1/(\psi - \gamma_s)$ symmetrically. The GE results carry through with the sector-specific multiplier in place of the common one: all equilibrium objects have identical functional form.

Observational equivalence. The primary cross-sector regression identifies the coefficient on $\beta_s \ln p^E$, equal to $-1/(\psi - \gamma)$ under the baseline and $-1/(\psi_s - \gamma)$ or $-1/(\psi - \gamma_s)$ under the extensions. A single cross-sector slope cannot distinguish a common multiplier from a sector-specific one, nor separate ψ from γ , so the baseline imposes the most parsimonious structure consistent with the regression.

Discriminating test. Whether $\hat{\delta}_\pi / \hat{\delta}_M = (\sigma - 1)\psi$ is constant across sectors discriminates the common- ψ baseline from the ψ_s extension: under common ψ the ratio is constant; under sector-specific ψ_s it varies. The intensive margin is too imprecise in the Norwegian experiment to run this test informatively, which is the empirical reason the baseline imposes common ψ and common γ

and treats the sector-specific multiplier as an extension. The urban economics literature provides some evidence for sector-varying agglomeration (e.g., knowledge spillovers stronger in high-tech than food processing), which a richer dataset could identify through this ratio.

C Data

This appendix documents the datasets used in the empirical and quantitative analysis. Table C.1 summarises the primary sources, their coverage, and the variables they provide. The remainder of the section describes cleaning and sample construction choices in detail, with particular attention to selection and measurement issues where they affect interpretation of the results.

Table C.1: Primary data sources

Dataset	Variables	Coverage	Use in paper
Skatteetaten (Norwegian Tax Administration)	Entity-level financials: revenue, profit, fixed assets	Norway, 2015–2024	Norwegian model-aligned moments and descriptive checks
SSB Statbank	Net electricity consumption by municipality and broad sector	Norway, 2010–2024	Imputation of firm-level electricity use
Orbis (Bureau van Dijk)	Firm financial statements, ownership structure, NACE industry	27 European countries, 2005–2023	Cross-country long differences
Revelio Labs	Firm-level headcount from online worker profiles	27 European countries, 2008–2024	Cross-country employment long differences
Eurostat NA, PEFA, energy balances	Sectoral electricity consumption, hours worked, and real value added	EU countries, 1990–2024	Physical electricity and labour intensities; imputation and robustness classifications
Exiobase 3; Eurostat SBS	Sectoral electricity purchases, gross output, personnel costs, production value	EU countries, 2010–2011	Harmonised gross-output cost shares β_s and α_s for model-aligned empirics and quantification
ENTSO-E, Bloomberg, ECB	Day-ahead wholesale electricity prices by bidding zone; actual load by zone; exchange rates to the euro	European zones and countries, 2000–2025	Harmonised zone- and country-level wholesale prices
ENTSO-E, Ember, Eurostat	Installed capacity, generation by source, utilisation, bilateral transmission capacity	European countries, 2014–2024	Generation mix, K_j^F , $\bar{E}_j^0$, u_j , NTC matrix
Eurostat SBS	Country–sector labour productivity	EU countries	Validation of base productivity Φ_{js}

Notes: Coverage reflects the sample used in this paper; the underlying datasets are in several cases broader. Full cleaning procedures are described in the text below.

C.1 Norwegian administrative data

The primary dataset for the Norwegian quasi-experiment is legal entity-level financial statements from the Norwegian Tax Administration (*Skatteetaten*), covering all active incorporated firms (AS/ASA) from 2015 to 2024. The data has been kindly provided by Ranik Raaen Wahlstrøm and

is described in Wahlstrøm (2022). It provides annual entity-level outcomes: operating revenue, operating profit, and fixed assets. Firms are identified at the legal entity level, allowing multi-entity corporate groups to be tracked. I exclude finance, insurance, real estate (NACE sections K, L), public administration (O), and holding (00) entities from the sample. Firms are matched to municipalities and municipalities to bidding zones via a spatial join on official boundaries.

Bidding zone and municipality crosswalk. Norwegian Elspot bidding zone (NO1–NO5) polygons are taken from the NVE ArcGIS REST service. Municipality polygons come from Geonorge/Kartverket, pre- and post-2020 versions. Spatial joins use EPSG:25833 (UTM 33N) for accurate area calculations, with a primary-zone assignment by largest area overlap and a flag for municipalities straddling multiple zones ($> 1\%$ area in more than one zone). Norwegian municipality codes changed substantially over the sample period (Trøndelag merger 2018, kommunereformen 2020, partial reversal 2024). A year-specific crosswalk is built from the SSB KLASSE API (`classifications/131/changes`), mapping 965 historical codes to 357 current municipalities with 100% coverage. The analysis uses harmonised municipality codes, assigning historical municipality codes to small groups that stay consistent across all reforms. Svalbard is excluded, since it operates a separate isolated power grid.

Imputation of firm-level electricity use. Firm-level electricity consumption is imputed by combining municipality-level consumption from SSB Statbank (tables 10314 and 14490, in GWh by municipality, broad industry $n \in \{\text{primary, secondary, tertiary}\}$, and year) with sector-level electricity intensity from Eurostat. Within each municipality–SSB sector–year cell, total consumption is allocated to NACE 2-digit sectors in proportion to $\iota_s R_{smt}$, where ι_s is Norway-specific Eurostat electricity intensity (EU average as fallback) and R_{smt} is total firm revenue in that sector–municipality–year cell. Within each NACE sector cell, consumption is then allocated to firms in proportion to revenue. By construction, imputed firm consumption aggregates exactly to observed municipality–sector totals; the procedure preserves geographic and temporal variation from the SSB data while using Eurostat intensities to introduce cross-sector variation.

Physical sector intensities. The physical electricity-intensity measure is constructed from four Eurostat sources combined in priority order: disaggregated industry statistics (`nrg_d_indq_n`) for NACE C16–C33 and F; disaggregated services statistics (`nrg_d_serq_n`) for NACE G–S; physical energy flow accounts (`env_ac_pefasu`) for 44 NACE 2-digit sectors from 2008; and complete energy balances (`nrg_bal_c`) as a deeper historical fallback. Each source fills gaps without overwriting higher-priority values. Real gross value added by NACE division comes from Eurostat National Accounts (`nama_10_a64`) in chain-linked volumes, giving $EI_{s,c,t} = \text{Electricity}_{s,c,t} / \text{Real GVA}_{s,c,t}$. The physical labour-intensity analogue uses hours worked by all employed persons from Eurostat

National Accounts (nama_10_a64_e), giving $LI_{s,c,t} = \text{Hours Worked}_{s,c,t} / \text{Real GVA}_{s,c,t}$. Norway-specific values are used where available; missing Norwegian sectors fall back to EU cross-country averages. These physical measures are used for the firm electricity-use imputation, descriptive sector rankings, and robustness exercises, but they are not the baseline α_s and β_s cost-share inputs in the model-aligned Norwegian specification.

Gross-output factor cost shares. The baseline monetary cost-share measures used in the model-aligned Norwegian regressions and in the quantitative model are harmonised to the same gross-output denominator. The electricity cost share is $\beta_s = \text{electricity purchases} / \text{gross output}$, constructed from Exiobase 3 (Stadler et al., 2018). Exiobase is a multi-regional input–output database covering 163 industries across 44 countries plus five rest-of-world regions; unlike standard national accounts use tables, it disaggregates energy purchases, allowing monetary electricity expenditure to be isolated from gas, steam, and other energy carriers. For each industry column s in country j and year t , electricity purchases are computed as the column sum of all intermediate-use rows corresponding to electricity supply activities (all generation sources, transmission, and distribution) in the industry-by-industry transaction matrix Z , and divided by gross output. Where the harmonised NACE Rev. 2 cell has no direct Exiobase industry because Exiobase follows the older NACE Rev. 1.1/ISIC Rev. 3 classification, I fill the cell using an explicit parent- or sibling-sector fallback within Exiobase and retain fallback indicators in the replication files. The labour cost share is $\alpha_s = \text{personnel costs} / \text{production value}$ from Eurostat Structural Business Statistics (sbs_na_sca_r2), where personnel costs (V13310) are employee compensation inclusive of employers’ social contributions and production value (V12120) is the gross-output denominator. Both shares are constructed over 2010–2011, the common window for which Exiobase 3 is built from actual national-accounts source data rather than post-2011 nowcasts (Rasul et al., 2024). Cells are pooled across European countries before constructing the fixed sector characteristic, so the exposure measure is predetermined with respect to the Norwegian 2021–2024 electricity-price shock.

The same source definitions are used at two aggregation levels. For the Norwegian empirical exercise, shares are produced at the harmonised NACE level used in the firm panel, so the interaction variables are $\beta_s \log p_{zt}^E$ and $\alpha_s \log w_{mt}$ at the same sectoral resolution as the municipality–sector cells. For the quantitative model, the same gross-output shares are constructed at the model’s thirteen-sector resolution: section-level for B, F, G, H, I, J, M, and N, two-digit for the four electricity-intensive manufacturing divisions (C17, C20, C23, C24), with the manufacturing residual anchored on the section-C totals so that weighted adding-up across divisions is exact. Because both shares are measured on gross output, $\alpha_s + \beta_s < 1$, with the remainder accruing to capital and intermediate inputs at common cross-country prices, consistent with the gross-output production technology used in the quantification (Section 5); the analytical model instead normalises $\alpha_s + \beta_s = 1$. An Exiobase-based labour share (compensation of employees over gross

output) and the older value-added electricity cost share are retained as internal consistency and robustness checks, but they no longer define the baseline model-aligned empirical exposure.

Electricity prices. Zone-level day-ahead wholesale prices for the Norwegian analysis are taken from the harmonised zone-month panel described in Section C.5. Prices are measured in EUR/MWh in the canonical panel and converted to NOK/MWh at contemporaneous exchange rates where the Norwegian analysis requires local-currency units.

Electricity subsidy. The *Energitilskuddsordningen* was a temporary government support scheme for electricity-intensive businesses introduced by the Norwegian Ministry of Climate and Environment in November 2022 (Forskrift 2022-11-21 nr. 1994) and administered by Enova SF with a total budget of NOK 2.8 billion. To qualify, firms had to demonstrate electricity costs of at least 3% of revenue in the first half of 2022 and operate in a bidding zone where the average Elspot price in Q3 2022 was at least double the Q3 2021 average; in practice this price-zone condition was met only in the southern zones (NO1, NO2, NO5) affected by the continental gas-price pass-through, and not in the northern zones (NO3, NO4) where hydropower prices remained low, giving the scheme a geographic structure that coincides with the treatment assignment used in the empirical analysis. The scheme operated on two support tiers. Tier 1 (*Støttrinn 1*) required completion of an energy audit and provided up to 25% of electricity costs above a threshold of 70 øre/kWh for Q4 2022, capped at NOK 3.5 million per firm. Tier 2 (*Støttrinn 2*) additionally required firms to commit to implementing specific energy efficiency investments from a predefined eligible list; the subsidy rate on above-threshold electricity costs rose to 45%, and firms were eligible for investment support of up to 50% of investment costs, subject to a combined ceiling of NOK 5 million per corporate group. In total, 3,186 applications were approved and NOK 2,780 million was disbursed; approximately 2,137 approved projects included investment support (Tier 2), with the remainder (roughly 1,050 firms) qualifying under Tier 1 only. Firm-level grant data, including recipient organisasjonsnummer and grant amounts, are published by Enova and matched to the Wahlstrøm financial panel to track subsidy take-up.

C.2 Orbis firm financial data

The cross-country analysis uses Bureau van Dijk's Orbis Historical database, which aggregates firm financial statements from national registries across Europe. The working panel covers 27 European countries from 2000 to 2024. The initial cleaning excludes NACE sections D, E, L, O, P, Q, and U; the regressions apply the narrower model-aligned sector restriction described below.

Estimation sample. The estimation sample for the cross-country regressions in Section 4 is a European firm panel — excluding a small number of United States, Luxembourg, Malta, and

Serbia entries — restricted to NACE Rev. 2 sections B, C, F, G, H, I, J, M, and N. The level specifications contain roughly 9.0 million firm-year observations for log gross profit, 9.2 million for inverse-hyperbolic-sine gross profit, and 14.9 million for log operating revenue. Direct oil- and gas-extraction sectors (NACE 06 and 09) are excluded to mirror the Norwegian sample.

Known limitations. Orbis coverage varies substantially across countries and improves over time. Small firms are systematically underrepresented, so the firm-size distribution is not fully representative (Bajgar et al., 2020); large firms have better coverage and fewer missing values (Arndt, 2023). Apparent entry and exit can reflect coverage changes rather than true dynamics (Bajgar et al., 2020). Consolidated parent accounts can double-count subsidiary activity (Kalemli-Özcan et al., 2024). Employment information is frequently missing. BVD firm identifiers can change over time due to legal restructuring.

Cleaning procedure. I retain unconsolidated accounts covering 12 months and ending in December. I drop firm-years only when total assets, operating revenue, sales, and employment are all missing. To avoid conditioning the historical sample on survival through the November 2025 Orbis extraction, I do not retrospectively remove the accounts of firms subsequently recorded as inactive, dissolved, or in liquidation. When Orbis reports the date on which an inactive legal status took effect, I remove only financial statements closing after that date. Firms recorded as inactive but without a status date remain in the baseline sample. Excluding these undated inactive firms leaves the electricity-price coefficient at a similar magnitude, so this choice does not materially affect the result. I drop a firm if any year contains negative total assets, tangible fixed assets, employment, or sales; employment above 2 million; or employment-to-assets, employment-to-sales, or sales-to-assets ratios above the within-country 99.9th percentile. Duplicate firm–country–year observations are resolved by selecting the record with non-missing and highest total assets, while retaining the maximum employment count reported across duplicates. Firms are assigned their primary four-digit NACE Rev. 2 code. BVD identifier changes are reconciled iteratively using the BVD crosswalk. All monetary variables are denominated in euros.

Multinationals. Corporate groups are identified via the GUO50 (the entity at the top of the ownership chain with at least 50% control). Firms with a common GUO50 across countries are linked into multinational groups. Firms are reclassified as domestic when the GUO is itself a holding company (NACE 6420). This allows the intensive-margin cross-country regressions to interpret the coefficient on an electricity-intensity interaction as within-group reallocation across subsidiaries.

C.3 Revelio Labs employment data

Employment counts for the cross-country analysis come from Revelio Labs, which constructs firm-level headcount series from online worker profiles (primarily LinkedIn, Glassdoor, Indeed); these profile-based counts better capture actual employment geography than accounting registrations alone. The sample scope is the largest ca. 10,000 global companies, with monthly employment by country from 2008 onwards. The working sample covers 27 European countries, 2009–2024, and contains roughly 8,300 firms; 99% of firms in the sample have employees in more than one country, reflecting the large-multinational sampling frame, and within-country 69% of firms are observed in every year. Matching to Orbis is performed on LEI where available (LEI coverage is incomplete before 2018) and through fuzzy name and country matching as a fallback.

Bias and weighting. Revelio’s online sources generate bias by occupation (professional occupations are overrepresented), location (urban bias), and firm size (large firms overrepresented), and there are reporting lags for job changes. The analysis uses Revelio’s cross-sectional sampling weights, which probabilistically classify each profile to ISCO occupation codes, stratify by occupation and country, and reweight against ILO occupational benchmarks. A reporting-lag correction is applied at monthly frequency before annual aggregation. For US firms, Hampole et al. (2025) document a 0.76 correlation of log employment between Revelio and Compustat, which provides the best available external validation. Country fixed effects are included in every specification to absorb level differences in platform penetration.

Cleaning. Sectors 22, 53, 61, 62, and 92 (utilities, real estate, education, health, public administration) are excluded to match the Orbis sample. Minimum employee count is 2, minimum observation length is 2 consecutive years, and observations with implausible year-on-year growth (>300%) or missing headcount are dropped. Single-year gaps are linearly interpolated; longer gaps are left missing. RCIDs identify firms across countries; a 1% workforce threshold determines whether a firm is recorded as operating in a country, and ultimate parent RCIDs are used for group-level aggregation.

C.4 Sectoral data and NAICS–NACE concordance

The sectoral panel of electricity intensity is merged to Orbis firms directly via NACE Rev. 2 and to Revelio firms via a NAICS–NACE concordance. Six-digit NAICS codes are mapped to NACE 2-digit divisions using UN Statistics Division correspondence tables linking NAICS to ISIC Rev. 4, which is identical to NACE Rev. 2 at the 2-digit level by EU regulation. Four NAICS vintages (2007, 2012, 2017, 2022) are processed separately and applied to firm-years according to the vintage in force. Many-to-many mappings are resolved by (i) retaining non-partial correspondences where available and (ii) assigning the NACE division receiving the plurality of ISIC class-level

links otherwise. Across vintages, roughly 85% of six-digit NAICS codes map to a unique NACE division, and the final match assigns electricity intensity to over 98% of six-digit NAICS codes.

C.5 Electricity price data

Wholesale electricity prices are constructed as a harmonised panel at bidding-zone and country level, at both monthly and annual frequency. The objective is to maximise country and time coverage while using a transparent source hierarchy and a single currency convention. The final price variable is measured in EUR/MWh.

Primary source: ENTSO-E day-ahead prices. The benchmark source is the ENTSO-E Transparency Platform.¹²⁵ I query day-ahead prices at the bidding-zone level using the Transparency Platform API data item Day-ahead Prices (document type A44, data item 12.1.D), with the bidding-zone EIC code used as both the input and output domain. The API series start in 2015, so ENTSO-E is the preferred source from 2015 onward where coverage is available. Prices are downloaded at their native settlement resolution, including hourly and 15-minute observations, and are first cleaned at the interval level before aggregation. Monthly and annual zone prices are computed as resolution-weighted averages of interval prices.

Currency conversion. ENTSO-E reports some historical prices in local currency rather than EUR, notably BGN, GBP, PLN, and RON. These observations are converted before aggregation using the ECB euro foreign exchange reference rates.¹²⁶ The ECB file reports rates as one euro equal to x units of foreign currency. Therefore a local-currency electricity price is converted as

$$p_{z\tau}^{EUR} = \frac{p_{z\tau}^{LCU}}{FX_t^{LCU/EUR}},$$

where τ indexes the settlement interval date. Since ECB reference rates are published only on business days, the script carries the previous available ECB rate forward over weekends and holidays.

Fallback sources and source priority. Bloomberg day-ahead spot prices provide the main fallback source, especially before ENTSO-E API coverage begins and for remaining post-2015 gaps. The Bloomberg data are already harmonised in EUR/MWh. At the country level, I also retain an Ember-cleaned ENTSO-E country-level EUR monthly file as an intermediate fallback. The source hierarchy is: (i) ENTSO-E API zone prices, including ECB-converted local-currency observations;

¹²⁵<https://transparency.entsoe.eu>

¹²⁶https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/index.en.html

(ii) Ember country-level EUR prices, for country panels only; and (iii) Bloomberg EUR prices. Annual fallback values from Bloomberg and the ENTSO-E legacy file are retained only when within-year coverage is sufficient: at least 200 daily Bloomberg observations or at least 10 monthly legacy observations.

Aggregation across bidding zones. For single-zone countries, the country price is the corresponding zone price. For multi-zone countries, country prices are constructed from zone prices using load weights. The relevant multi-zone countries are Denmark (DK1–DK2), Italy (IT_NORD, IT_CNOR, IT_CSUD, IT_SUD, IT_CALA, IT_SICI, IT_SARD), Norway (NO1–NO5), and Sweden (SE1–SE4). Zone weights are built from ENTSO-E Actual Total Load data (document type A65, process type A16, data item 6.1.A), downloaded by bidding zone. Load observations are converted from MW to MWh using their reported interval length, and the country-level monthly or annual price is

$$p_{ct}^{EUR} = \sum_{z \in c} \omega_{zt} p_{zt}^{EUR}, \quad \omega_{zt} = \frac{L_{zt}}{\sum_{z' \in c} L_{z't}}.$$

The procedure requires at least 95% of country load weight to be represented by available zone prices before accepting a load-weighted country value; no accepted country aggregation falls below full weight coverage. Italy’s Calabria zone is treated as starting in 2021, consistent with the ENTSO-E price and load coverage; before 2021 Italy is aggregated over the active zones.

C.6 Electricity supply-side primitives

The quantitative model in Section 5 requires, for each country j , the installed capacity of fossil, nuclear, and renewable generation (K_j^F, K_j^N, K_j^R), annual generation by source ($E_j^D, E_j^F, E_j^{\text{clean}}$), the implied fossil utilisation rate u_j , and the bilateral net transfer capacity matrix $\{\text{NTC}_{jk}\}$. These objects feed the supply curve in equation (B.1) and the market-clearing step of the solver. The baseline sample is the 2015–2019 average throughout — capacity, generation, prices, the bilateral NTC matrix, and baseline physical flows alike — with the 2019 single-year baseline retained as a robustness reference.

Installed generation capacity. Country-level installed net electrical capacity by technology is taken from Eurostat `nrg_inf_epc` (“Infrastructure – electricity – annual data”). The raw file reports capacity in MW by SIEC (Standard International Energy Product Classification) code and plant type (main activity producers and autoproducers combined). I use the following mapping to the model’s three technology bins: *fossil* (K_j^F) is SIEC “Combustible fuels”, which aggregates coal, oil, natural gas, and biomass/waste; *nuclear* (K_j^N) is isolated from the SIEC “Nuclear and other” bundle by subtracting the “Other not elsewhere classified” residual, floored at zero; *renewable* (K_j^R) is the sum of hydro (including pure hydro and reservoir, but not pumped storage

pumping capacity), wind (onshore and offshore), solar (photovoltaic and thermal), geothermal, tide/wave/ocean. Biomass is bundled with fossil capacity because Eurostat’s “Combustible fuels” SIEC does not separate it cleanly across all countries; this choice is documented in the derivation script and affects u_j only marginally for most countries. Pumped storage pumping capacity is not subtracted from hydro, since it contributes to dispatchable renewable generation at the annual frequency used here.

Generation and fossil utilisation. Annual generation by source is taken from Ember’s “Yearly Electricity Data” (previously published by Ember/Sandbag), which harmonises ENTSO-E, IEA, and national statistics into a consistent country panel with Nuclear, Fossil, and Clean aggregates in TWh. I use Ember rather than raw ENTSO-E for the generation layer because Ember backfills the small gaps and reconciles reporting-boundary differences across Nordic and non-ENTSO-E countries. Total electricity demand E_j^D is taken as total generation (net imports are handled separately via the NTC matrix in the quantitative solver). Fossil generation E_j^F is the “Fossil” aggregate, and clean generation $E_j^{\text{clean}} = E_j^N + E_j^R$ is split into nuclear and renewable components. The fossil utilisation rate is defined as $u_j = E_j^F / (K_j^F \cdot 8760)$ converted to consistent units (capacity in MW multiplied by 8760 h/yr gives TWh/year), and clipped to $[0.01, 0.99]$ before use in ε_j to avoid corner-case instability in the rare country-years where u_j approaches the boundary. The fuel cost share is $\chi_j = (p_j^E - c_0^F) / p_j^E$, where c_0^F is a technology-specific non-fuel variable cost (gas 4.50, coal 2.95, other fossil 6.00, bioenergy 4.39 EUR/MWh, from the Danish Energy Agency technology-catalogue variable operations-and-maintenance benchmarks) and the frontier heat rate is $h_{\min} = 1.7241$ (MWh of fuel input per MWh of electricity, from the Danish Energy Agency combined-cycle gas efficiency of 58%). The supply elasticity $\varepsilon_j = (\chi_j / \zeta) \cdot E_j^D / [K_j^F (1 - u_j)]$ is computed for $\zeta \in \{1.5, 2, 2.5, 3\}$, with $\zeta = 2$ as the baseline calibration; setting $c_0^F = 0$ (so $\chi_j = 1$) recovers the simpler $(1/\zeta) \cdot E_j^D / [K_j^F (1 - u_j)]$ benchmark. In the solved baseline the DEA non-fuel costs give fuel cost shares $\chi_j \in [0.84, 0.93]$ and, at $\zeta = 2$, supply elasticities $\varepsilon_j \in [0.06, 2.45]$ over fossil utilisation $u_j \in [0.03, 0.70]$ across the twenty-three fossil-marginal countries, reported country by country in Appendix Table F.1. What governs ε_j is spare fossil capacity relative to demand rather than the size of the fossil fleet in absolute terms or utilisation on its own: France holds the smallest fossil fleet relative to its own generation in the sample and has the steepest supply curve despite a moderate utilisation rate, while Lithuania holds fossil capacity seven times its own generation, nearly all of it idle behind imports, and has the flattest.

The fossil supply-curve curvature ζ . The Pareto shape ζ is calibrated by fitting the model’s merit-order supply curve, $p_j^E = c_0^F + p^F h_{\min} (1 - u_j)^{-1/\zeta}$, to a European fossil cost stack built from plant-level data. The plant register is `powerplantmatching` (Gotzens et al., 2019), which reconciles the major European and global plant registries into a single matched fleet. Reported plant efficiencies, which are what identify the curvature, are far sparser than capacity: 510 plants in

the register carry one, and 95% of those are matched to the Open Power System Data conventional-power-plant list (Open Power System Data, 2020). Each plant’s marginal cost is its fuel price divided by its observed efficiency plus a technology-specific variable operations-and-maintenance benchmark, and plants are ordered into a capacity-weighted merit order over which ζ is recovered by capacity-weighted least squares with c_0^F held fixed. Fuel prices are the 2015–2019 window means of the World Bank commodity series. I take the fossil marginal cost from world fuel prices passed through the heat-rate distribution, rather than solving a Hotelling extraction path as in Arkolakis and Walsh (2023), treating Europe as a price-taker in global fossil-fuel markets; the discipline of the schedule against observed dispatch and wholesale-price data follows the same paper.

Two variants bracket the estimate from opposite directions. The observed-efficiency stack uses the 339 active EU fossil plants with a reported efficiency, covering 27% of EU fossil capacity, and captures both within-fuel heat-rate dispersion and between-fuel switching; it gives $\zeta = 2.0$ ($R^2 = 0.92$). The cross-fuel stack applies one capacity-weighted efficiency per fuel to the full fleet, so it covers all capacity but draws curvature from fuel-switching alone; it gives $\zeta = 2.3$ ($R^2 = 0.93$). An external check against the PJM zonal supply curves of Sahraei-Ardakani, Blumsack and Kleit (2015) delivers ζ in the same low-single-digit range. I take $\zeta = 2$ as the baseline and report sensitivity over $\zeta \in \{1.5, 2, 2.5, 3\}$.

The stack the fit is run on prices fuel but not carbon, while p^F in the supply curve is the all-in fuel-plus-carbon cost, so the calibration of ζ is carbon-exclusive. Adding the EU ETS cost over the same window raises the fitted curvature to 2.45 on the observed-efficiency stack and 2.60 on the cross-fuel stack, because the carbon adder is roughly twice as large per unit of fuel for coal as for gas, so it compresses the coal-to-gas step and moves the convexity of the merit order out into the oil tail. Both values sit inside the sensitivity band, and the PJM comparison is itself carbon-free, so the carbon-exclusive fit is the like-for-like one. The choice is immaterial for what this paper measures: across the swept band the three location channels move by less than 1%, and the manufacturing gain of Section 6 runs from 13.6 to 13.1 basis points as ζ goes from 1.5 to 3, because ζ governs the shape of the merit order rather than where firms locate. What it does move is the aggregate fuel saving, which rises from 1.4 to 2.5 bp over the same band.

Carbon cost in variable fuel prices. The carbon component of the variable fossil fuel price is constructed as $\tau_t^{CO_2} e_f^F$, where $\tau_t^{CO_2}$ is the EU ETS allowance price in EUR/tCO₂ and e_f^F is the source-specific emissions factor in tCO₂/MWh_{th}. I take $\tau_t^{CO_2}$ from the ICAP Allowance Price Explorer, using the public EU ETS secondary-market price series (International Carbon Action Partnership, 2026a). ICAP documents the underlying EU ETS spot data as EEX-provided up to 2018 and ICE-provided from 2019 (International Carbon Action Partnership, 2026b). Annual secondary-market means over the baseline window are 7.68, 5.35, 5.83, 15.82, and 24.88 EUR/tCO₂ for 2015–2019, so

the baseline carbon-price input is 11.91 EUR/tCO₂, the unweighted mean of annual averages. I keep primary auction-price averages as a diagnostic but use the secondary price for the marginal-cost wedge faced by generators. Fuel-specific emissions factors come from the IPCC default stationary-combustion factors, cross-checked against the JRC Covenant of Mayors defaults (Intergovernmental Panel on Climate Change, 2006, Joint Research Centre, 2017): natural gas 0.202, hard coal 0.341, lignite 0.364, gas/diesel oil 0.267, and heavy fuel oil 0.279 tCO₂/MWh_{th}. The model aggregation maps gas to natural gas, coal to hard coal (with lignite retained for the plant-level cost-stack fit and sensitivity checks), and other fossil to gas/diesel oil; bioenergy receives a zero EU ETS carbon wedge. The all-in variable fuel price used in the fossil cost component is therefore $p_{ft}^F + \tau_t^{CO_2} e_f^F$ before aggregation to the country fossil mix.

Transmission capacity (NTC). Bilateral transmission capacity between bidding zones is downloaded from the ENTSO-E Transparency Platform API, document type A61 (“Forecasted Transfer Capacity”) with process type A01 (day-ahead), at hourly frequency over 2014–2024. Continental European borders publish day-ahead NTC at daily granularity under contract type `contract.MarketAgreement.Type = A01`. Nordic and Baltic borders (NO, SE, FI, DK, EE, LV, LT) do not publish under A01; ENTSO-E instead serves these borders at monthly (A03) or yearly (A04) granularity, reflecting the Nord Pool long-term allocation mechanism. The download script accordingly tries the three contract types in order (A01 → A03 → A04) for each border-year and records the successful type as a column in the raw output. Each border is queried in both directions because NTC is intrinsically asymmetric (TSOs allow different capacity $A \rightarrow B$ vs $B \rightarrow A$). Zone-pair values are aggregated to country-pair directional NTC by summing over all zone-pair borders connecting the same country pair (e.g., NO→SE is the sum over NO1→SE3, NO3→SE2, NO4→SE1), and then averaged over the baseline years 2015–2019 to produce the $J \times J$ bilateral NTC matrix used by the solver. Self-links and country pairs with no common border are coded as missing.

Three interconnectors entered service part-way through the baseline window: NordBalt (LT–SE, February 2016), LitPol Link (LT–PL, December 2015), and COBRACable (DK–NL, September 2019). For the years before commissioning, ENTSO-E publishes a year-ahead nameplate forecast rather than an operating capacity — the mean, median, and maximum coincide at the rated value — so those observations would enter the average as if the link were fully available before it physically existed. I set the pre-commissioning years to zero and retain them in the five-year mean rather than dropping them, so that a border live for part of the window enters the baseline at its average available capacity over the window rather than at its operating capacity. This matches the zero-inclusive treatment of the price and demand averages over the same window, and is conservative for the counterfactuals: it credits the baseline network with less delivered integration, so the gain from further integration is not understated. The adjustment applies only to newly com-

missioned borders and not to capacity expansions on pre-existing ones; the France–Spain INELFE upgrade of October 2015, for instance, retains its observed 2015 value, which is genuine partial-year operation of a long-standing link. For robustness, an undirected version reporting the min, mean, and max across the two directions is also produced.

Baseline physical flows and effective transfer bounds. The quantitative trade block is anchored to observed baseline cross-border flows rather than assuming either zero baseline trade or full capacity utilisation. Directed baseline physical flows F_{jk}^0 are taken from ENTSO-E physical-flow data (document type A11) averaged over 2015–2019, the same window as the NTC matrix, and with the same commissioning treatment applied to the three links that entered service part-way through it.

I also apply the commissioning mask to ALEGrO (BE–DE), which was the first direct electricity interconnector between Belgium and Germany and entered commercial operation only on 18 November 2020.¹²⁷ ENTSO-E A11 nevertheless contains positive BE–DE observations in 2018–2019, averaging 69.9 MW from Belgium to Germany and 63.2 MW in the reverse direction over the reported years. Because these observations predate the physical completion and testing of the first direct interconnector, they cannot represent direct physical transfers across an operational BE–DE link. I retain the observations in the raw-data archive but set both directions to zero throughout 2015–2019. The resulting treatment keeps baseline flows consistent with the zero NTC and prevents the pre-service records from creating a spurious direct edge in the model network.¹²⁸

Because A11 reports physical rather than commercially scheduled flows, observed flows occasionally exceed the directed A61 NTC bounds; I therefore export effective solver bounds that minimally contain the observed baseline flow. For an ordered border (j, k) with signed baseline flow F_{jk}^0 (positive from j to k) and raw directed capacities $\bar{T}_{jk}, \bar{T}_{kj}$, the effective bounds are

$$\underline{F}_{jk}^{\text{eff}} = \min\{-\bar{T}_{kj}, F_{jk}^0, 0\}, \quad \bar{F}_{jk}^{\text{eff}} = \max\{\bar{T}_{jk}, F_{jk}^0, 0\}.$$

This widens a raw bound only in the direction needed to make the observed baseline feasible, and never uses peak or high-percentile flows to manufacture additional counterfactual headroom. Counterfactual flows in Section 6 are constrained to the directed NTC limits, or their effective extension, around this baseline position. Ireland is a special case: its only observed baseline interconnector in this period is with Great Britain, which lies outside the EU-24 sample, so Ireland

¹²⁷Elia reported in September 2020 that construction had been completed and energisation testing was about to begin; see https://www.elia.be/-/media/project/elia/shared/documents/press-releases/2020/20200916_cp-alegro-final_en_publication.pdf. Amprion and Elia report commercial commissioning on 18 November 2020; see https://www.amprion.net/Press/Press-Detail-Page_37505.html.

¹²⁸This treatment does not rule out German electricity passing through Belgium as transit or loop flow before 2020. Such flows travelled through neighbouring networks, principally via the Netherlands and France, and belong to the corresponding physical borders rather than to a direct BE–DE link.

is treated as having no direct intra-EU-24 transmission link in the baseline; an explicit EU+GB extension is left for future work.

Post-2022 flow-based coupling and future capacity. From June 2022 onwards, the Core Capacity Calculation Region (13 Central European TSOs) switched from NTC to flow-based market coupling, in which the relevant capacity object is a set of zone-to-zone ATCs published by JAO rather than the bilateral NTCs published by ENTSO-E. For the 2015–2019 baseline window this is immaterial – ENTSO-E A61 data is complete for the pre-flow-based period – but for post-2022 robustness and for the counterfactual extension window the JAO Core CCR published ATC download is used as a supplementary source. For forward-looking counterfactuals (transmission expansion scenarios in Section 5), bilateral capacities from ENTSO-E’s TYNDP 2024 Reference Grid scenario provide the expansion trajectory through 2030, 2035, and 2040; these are used only in counterfactual runs and never in the baseline calibration.

Hydro natural potential. The natural-potential comparison in Section 6 requires a measure of each country’s hydropower endowment that reflects geography and hydrology rather than past deployment decisions. I take gross theoretical hydropower potential from Hoes et al. (2017), who derive site-level potential globally at high resolution from runoff and elevation drop, independent of whether a location has been developed. Each point carries an annual gross theoretical potential, which I aggregate to bidding zones — assigning points to zone polygons in an equal-area projection — and then to countries, expressed in TWh/year. This geography-first object is the natural-potential input θ_j^{hydro} used in the counterfactuals; it sums to roughly 1,700 TWh/year across the European countries in the sample. The JRC European Hydropower Database, which records installed plant capacity and generation, is used only to validate the spatial pattern and to measure observed deployment, not as the potential measure itself. Consistent with the renewable-source typology, theoretical, installed, and untapped hydro are kept as separate objects, and gross theoretical potential is not relabelled as technically or economically feasible potential.


D Additional empirical results

D.1 Norwegian analysis: descriptives

This section provides descriptive evidence for the Norwegian firm-level sample used in the main analysis. Table D.1 reports summary statistics for the full sample of active firms, the sub-sample of continuing firms used in the intensive margin analysis, and municipality-sector-year aggregates. Table D.2 compares pre-period (2015–2019) means of key financial outcomes across high- and low-price municipalities using simple t -tests, while Tables D.3 and D.4 provide more detailed balance tests—including distributional statistics and municipality-clustered standard errors—for


each ring definition and the full zone sample; the two treatment groups are comparable on many observables but not perfectly balanced across all specifications. Figure D.6 shows the number of active firms by bidding zone over the sample period. Figure D.7 and Table D.5 document that sector and electricity intensity compositions are similar across treatment groups. Figures D.8 and D.9 plot raw entry and exit rates over time by treatment group; visual inspection reveals no diverging pre-trends. Finally, figure D.10 shows how the key identifying variation—the interaction of log electricity prices and pre-period sectoral electricity intensity—evolves over time and across treatment groups and intensity terciles.

Figure D.1: Event study of firm electricity use at municipality level




Note: This figure shows the result of an event study of firm electricity use at municipality level. It compares municipalities in high-price and low-price zones, controlling for both municipality and year fixed effects. Bars show 95% confidence intervals using municipality-clustered standard errors.

Figure D.2: Firm electricity use at zone level




Note: This figure shows firm electricity use for by zone, normalised to 2019 for comparison, aggregating municipality-level data from SSB Norway.

Figure D.3: Descriptives: median municipality outcomes by zone




Note: Electricity use, operating revenue, and active firms are indexed to 2019 within each municipality before taking the zone median. Operating profit is the median of signed municipality totals, indexed to the zone's 2019 median; municipalities with losses are included. Data are from SSB Norway and firm accounts.

Figure D.4: Long difference municipality outcomes by high- vs. low-price zone



Note: The figure shows the result of a long-difference specification at municipality level, comparing municipalities in high- ad low-price zones. The vertical axis is the equal-cell mean log change in period-average active-firm counts in southern high-price municipalities minus that in northern low-price municipalities, multiplied by 100. The horizontal axis is the sector's predetermined electricity cost share in percentage points. The periods are 2015–19 and 2021–24, omitting 2020. Sector support varies: C17 has five northern and 25 southern cells, C19 has one northern and six southern cells.

Figure D.5: Difference in active firm count growth, weighted

Note: Each dot represents one of 13 sectors, with manufacturing split into C17, C20, C23, and C24. For municipality–sector cells with positive period-average active-firm counts in both periods, I calculate the log change in counts and average it separately within the South (NO1, NO2, NO5) and North (NO3, NO4), weighting each cell by its mean active-firm count in that panel’s pre-period. The vertical axis is $100[\exp(g_s) - 1]$, where g_s is the South-minus-North difference in those weighted mean log changes. Panel (a) compares 2018–19 with 2015–17 (3,040 cells); panel (b) compares 2021–24 with 2015–19 (3,092 cells). The horizontal axis is the predetermined sector electricity cost share in percentage points. The line fits the 13 sector dots with equal sector weight. Its shaded band and heteroskedasticity-robust standard error describe the fit across those dots, rather than uncertainty across the five price zones. These are weighted local-cell growth rates, not changes in total sector firm counts.

Table D.1: Descriptive statistics for company financial data, 2015–2024**(a) Full sample of active firms**

	N	Mean	SD	p25	Median	p75
Operating revenue (NOK 1000)	1,640,698	39,235.8	2,897,059.9	591.9	2,400.2	8,993.6
Sales (NOK 1000)	1,640,698	37,330.5	2,791,810.6	409.9	2,091.9	8,251.9
Value added (NOK 1000)	1,640,698	19,314.9	1,509,336.1	413.5	1,647.7	5,366.2
Operating profit (NOK 1000)	1,640,698	4,524.5	838,478.1	-39.4	87.3	586.8
Fixed assets (NOK 1000)	1,640,698	39,806.8	7,267,458.5	12.6	207.5	1,373.3
Total assets (NOK 1000)	1,640,698	60,017.8	8,306,965.3	470.2	1,754.3	6,480.6
Cost of goods (NOK 1000)	1,640,698	19,920.9	1,663,572.9	0.0	270.1	2,509.5
Labor expenses (NOK 1000)	1,640,698	6,307.9	340,749.4	15.3	773.1	2,687.4
Equity (NOK 1000)	1,640,698	22,750.4	3,111,747.5	48.4	431.1	2,182.5
Total debt (NOK 1000)	1,640,698	37,234.2	5,456,169.8	286.4	1,031.3	3,815.2
Cash (NOK 1000)	1,640,698	4,177.3	392,292.3	75.1	330.0	1,240.6
Firm age (years)	1,640,698	11.3	11.9	3.0	7.0	16.0
Elec. price (NOK/MWh)	1,640,698	534.3	499.1	248.0	388.7	606.5
Electricity cost share (%)	1,639,240	0.9	0.5	0.6	0.7	1.1

(b) Sample of continuing active firms for intensive margin analysis

	N	Mean	SD	p25	Median	p75
Operating revenue (NOK 1000)	733,539	71,911.9	4,316,197.8	1,415.4	4,746.6	16,301.6
Sales (NOK 1000)	733,539	68,580.1	4,158,498.3	1,069.8	4,192.8	15,101.9
Value added (NOK 1000)	733,539	35,139.7	2,250,965.4	1,024.8	3,071.8	9,217.0
Operating profit (NOK 1000)	733,539	9,265.5	1,250,550.4	-9.6	222.2	1,063.2
Fixed assets (NOK 1000)	733,539	62,854.9	4,794,803.0	63.1	490.4	2,706.8
Total assets (NOK 1000)	733,539	98,724.2	7,014,447.5	1,144.0	3,578.7	11,270.0
Cost of goods (NOK 1000)	733,539	36,772.2	2,469,960.1	12.0	805.1	5,089.2
Labor expenses (NOK 1000)	733,539	10,864.1	505,791.6	349.3	1,478.7	4,575.0
Equity (NOK 1000)	733,539	36,192.3	1,856,772.1	248.3	1,189.7	4,428.0
Total debt (NOK 1000)	733,539	62,457.5	5,498,883.6	570.4	1,833.1	6,220.5
Cash (NOK 1000)	733,539	7,252.4	583,626.3	163.8	644.1	2,132.7
Firm age (years)	733,539	17.7	12.8	9.0	15.0	24.0

Table D.2: Simple balance table using mean values**(a) Ring 0**

	LOW	HIGH	Diff	SE	p-val
Operating revenue (NOK 1000)	12308.8	10005.9	2302.91	(1245.14)	0.064
Value added (NOK 1000)	5518.5	5460.3	58.21	(359.01)	0.871
Operating profit (NOK 1000)	466.7	594.9	-128.24	(155.68)	0.410
Fixed assets (NOK 1000)	4681.9	3697.2	984.63	(444.13)	0.027
Total assets (NOK 1000)	9644.9	8688.1	956.78	(871.65)	0.272
Cost of goods (NOK 1000)	6790.4	4545.7	2244.70	(971.79)	0.021
Labor expenses (NOK 1000)	2563.2	2417.4	145.76	(145.58)	0.317
Firm age (years)	11.1	10.8	0.25	(0.21)	0.239
Elec. price (NOK/MWh)	315.1	305.3	9.81	(1.73)	0.000
Electricity cost share (%)	1.0	0.9	0.05	(0.01)	0.000

(b) Ring 01

	LOW	HIGH	Diff	SE	p-val
Operating revenue (NOK 1000)	13322.5	9898.0	3424.45	(612.95)	0.000
Value added (NOK 1000)	6433.8	5213.4	1220.49	(210.50)	0.000
Operating profit (NOK 1000)	616.1	571.8	44.32	(64.69)	0.493
Fixed assets (NOK 1000)	4716.0	3286.3	1429.71	(295.88)	0.000
Total assets (NOK 1000)	10368.5	7632.4	2736.13	(455.40)	0.000
Cost of goods (NOK 1000)	6888.6	4684.7	2203.96	(487.61)	0.000
Labor expenses (NOK 1000)	2960.8	2373.3	587.47	(85.51)	0.000
Firm age (years)	11.7	10.9	0.85	(0.13)	0.000
Elec. price (NOK/MWh)	314.5	305.1	9.44	(1.04)	0.000
Electricity cost share (%)	0.9	0.9	0.05	(0.01)	0.000

(c) Ring 012


	LOW	HIGH	Diff	SE	p-val
Operating revenue (NOK 1000)	18484.4	25367.2	-6882.78	(1153.49)	0.000
Value added (NOK 1000)	8742.3	12357.4	-3615.18	(520.51)	0.000
Operating profit (NOK 1000)	1064.2	1355.7	-291.46	(235.44)	0.216
Fixed assets (NOK 1000)	7643.5	18197.2	-10553.68	(1418.54)	0.000
Total assets (NOK 1000)	15420.5	29224.0	-13803.55	(1852.36)	0.000
Cost of goods (NOK 1000)	9742.2	13009.8	-3267.60	(877.06)	0.000
Labor expenses (NOK 1000)	4095.8	4750.0	-654.21	(141.90)	0.000
Firm age (years)	11.3	11.1	0.20	(0.06)	0.002
Elec. price (NOK/MWh)	313.9	303.0	10.98	(0.48)	0.000
Electricity cost share (%)	0.9	0.9	0.04	(0.00)	0.000

(d) All zones

	LOW	HIGH	Diff	SE	p-val
Operating revenue (NOK 1000)	18347.5	37199.9	-18852.36	(1717.98)	0.000
Value added (NOK 1000)	8646.0	17856.7	-9210.70	(734.52)	0.000
Operating profit (NOK 1000)	1053.0	3022.3	-1969.33	(395.53)	0.000
Fixed assets (NOK 1000)	9468.5	37195.7	-27727.23	(3264.18)	0.000
Total assets (NOK 1000)	17816.6	54355.0	-36538.45	(4036.81)	0.000
Cost of goods (NOK 1000)	9701.5	19343.2	-9641.67	(1379.25)	0.000
Labor expenses (NOK 1000)	3838.7	6161.6	-2322.87	(91.91)	0.000
Firm age (years)	11.2	10.9	0.29	(0.03)	0.000
Elec. price (NOK/MWh)	307.4	304.9	2.53	(0.24)	0.000
Electricity cost share (%)	1.0	0.8	0.14	(0.00)	0.000

Notes: Balance tests using the panel of Norwegian companies for the pre-sample in 2015–2019 for border-adjacent rings. Monetary values converted to NOK 1,000.

Figure D.6: Active firms by bidding zone



Note: Number of active firms in the sample by bidding zone over time.

Table D.3: Balance tests of key firm outcomes in high- vs. low-price municipalities - I**(a) Ring 0**

	LOW (N = 5,110)				HIGH (N = 5,026)				Diff	p-val
	Mean	p25	Median	p75	Mean	p25	Median	p75		
Revenue	12,308.8	695.4	2,462.4	7,849.6	10,005.9	649.9	2,404.9	8,021.4	-2,302.9	0.514
Value added	5,518.5	480.2	1,579.0	4,830.0	5,460.3	450.0	1,606.3	4,815.7	-58.2	0.964
Op. profit	466.7	-22.7	104.9	492.1	594.9	-29.0	90.5	518.1	128.2	0.618
Fixed assets	4,681.9	51.6	392.1	2,033.4	3,697.2	47.0	409.1	2,141.0	-984.6	0.518
Total assets	9,644.9	573.8	1,981.3	5,862.6	8,688.1	563.0	1,898.1	6,147.4	-956.8	0.757
COG	6,790.4	17.6	401.8	2,190.6	4,545.7	10.3	396.3	2,465.7	-2,244.7	0.340
Labor expenses	2,563.2	62.0	736.7	2,279.7	2,417.4	37.2	717.8	2,381.2	-145.8	0.771
Firm age	11.1	3.0	8.0	17.0	10.8	3.0	7.0	17.0	-0.2	0.810
Elec. price	315.1	272.0	278.7	381.6	305.3	248.0	273.1	388.7	-9.8***	0.000
Electricity cost share (%)	0.96	0.51	0.88	1.09	0.92	0.48	0.86	1.09	-0.05	0.336

(b) Ring 01

	LOW (N = 16,191)				HIGH (N = 12,830)				Diff	p-val
	Mean	p25	Median	p75	Mean	p25	Median	p75		
Revenue	13,322.5	693.1	2,611.9	8,535.5	9,898.0	709.4	2,543.0	7,839.5	-3,424.5**	0.046
Value added	6,433.8	484.0	1,647.9	5,137.6	5,213.4	507.5	1,637.0	4,718.8	-1,220.5*	0.065
Op. profit	616.1	-24.1	103.1	519.1	571.8	-23.0	99.9	535.0	-44.3	0.659
Fixed assets	4,716.0	42.8	342.9	1,857.1	3,286.3	42.3	356.9	1,838.4	-1,429.7	0.147
Total assets	10,368.5	579.8	1,915.5	6,134.4	7,632.4	581.2	1,960.2	6,014.8	-2,736.1	0.119
COG	6,888.6	6.9	394.6	2,484.4	4,684.7	10.3	444.6	2,475.5	-2,204.0*	0.074
Labor expenses	2,960.8	58.6	773.2	2,539.7	2,373.3	61.3	745.4	2,303.8	-587.5**	0.028
Firm age	11.7	3.0	8.0	18.0	10.9	3.0	7.0	17.0	-0.8*	0.084
Elec. price	314.5	272.0	278.7	381.6	305.1	234.1	273.1	388.7	-9.4***	0.000
Electricity cost share (%)	0.94	0.48	0.88	1.09	0.89	0.47	0.79	1.09	-0.05*	0.072

Notes: Balance tests using the panel of Norwegian companies for 2015–2019. Geographic restrictions applied according to subtitles. Monetary values converted to NOK 1,000. The final two columns test the difference in means and apply municipality-level clustering to test significance.

Table D.4: Balance tests of key firm outcomes in high- vs. low-price municipalities - II**(a) Ring 012**


	LOW (N = 61,355)				HIGH (N = 70,600)				Diff	p-val
	Mean	p25	Median	p75	Mean	p25	Median	p75		
Revenue	18,484.4	636.6	2,503.3	8,806.3	25,367.2	640.5	2,562.0	9,087.7	6,882.8	0.242
Value added	8,742.3	444.5	1,673.2	5,322.3	12,357.4	460.5	1,722.4	5,380.8	3,615.2	0.139
Op. profit	1,064.2	-33.2	86.8	518.7	1,355.7	-40.9	86.6	562.9	291.5	0.490
Fixed assets	7,643.5	22.3	237.9	1,391.4	18,197.2	17.8	232.0	1,464.3	10,553.7*	0.097
Total assets	15,420.5	508.9	1,767.7	5,994.3	29,224.0	506.1	1,856.2	6,340.1	13,803.5	0.136
COG	9,742.2	0.4	335.8	2,538.2	13,009.8	0.0	318.6	2,586.0	3,267.6	0.345
Labor expenses	4,095.8	37.8	810.9	2,698.9	4,750.0	35.9	812.4	2,707.5	654.2	0.465
Firm age	11.3	3.0	8.0	17.0	11.1	3.0	7.0	16.0	-0.2	0.506
Elec. price	313.9	272.0	278.7	381.6	303.0	234.1	271.2	388.7	-11.0***	0.000
Electricity cost share (%)	0.89	0.57	0.74	1.09	0.85	0.51	0.70	1.09	-0.04	0.139

(b) All zones

	LOW (N = 179,800)				HIGH (N = 566,144)				Diff	p-val
	Mean	p25	Median	p75	Mean	p25	Median	p75		
Revenue	18,347.5	645.1	2,473.6	8,654.0	37,199.9	599.4	2,399.2	8,919.1	18,852.4***	0.005
Value added	8,646.0	471.2	1,690.0	5,253.6	17,856.7	410.1	1,597.6	5,156.6	9,210.7***	0.006
Op. profit	1,053.0	-36.4	86.8	512.4	3,022.3	-39.5	84.1	560.7	1,969.3*	0.080
Fixed assets	9,468.5	32.8	299.5	1,759.7	37,195.7	10.0	182.9	1,203.7	27,727.2**	0.035
Total assets	17,816.6	523.6	1,843.8	6,281.1	54,355.0	453.9	1,661.6	6,085.7	36,538.4**	0.027
COG	9,701.5	0.1	293.7	2,413.2	19,343.2	0.0	308.2	2,629.9	9,641.7***	0.005
Labor expenses	3,838.7	57.8	811.4	2,664.5	6,161.6	17.3	746.7	2,602.2	2,322.9**	0.027
Firm age	11.2	3.0	8.0	17.0	10.9	3.0	7.0	16.0	-0.3	0.135
Elec. price	307.4	236.2	278.7	381.6	304.9	236.5	273.1	388.7	-2.5*	0.083
Electricity cost share (%)	0.97	0.57	0.88	1.10	0.83	0.57	0.70	0.93	-0.14***	0.000

Notes: Balance tests using the panel of Norwegian companies for 2015–2019. Geographic restrictions applied according to subtitles. Monetary values converted to NOK 1,000. The final two columns test the difference in means and apply municipality-level clustering to test significance.

Figure D.7: Sector composition balance




Note: Scatter graphs showing sector shares in high- vs. low-price municipalities in the respective samples. Sectors are coloured by terciles according to their pre-crisis average electricity intensity.

Table D.5: Electricity intensity tercile shares by treatment group

Electricity tercile	LOW		HIGH	
	N	Share (%)	N	Share (%)
Low EI	87,732	48.8	334,408	59.1
Medium EI	60,593	33.7	171,762	30.4
High EI	31,377	17.5	59,351	10.5
χ^2 test: 8484.44, $p = 0.000$				


Notes: Firms and shares in the sample by EI tercile (defined at sector level) by treatment group.

Figure D.8: Entry rate by price region




Note: Entry rate by high and low-price region restricted by the sub-title geographies. Computed as share of new entrants out of total number of firms.

Figure D.9: Exit rate by price region




Note: Exit rate by high and low-price region restricted by the sub-title geographies. Computed as share of firms reporting their last year out of total number of firms. Series ends in 2023 because data for 2025 is not yet available, restricting the ability to observe operational exit in 2024.

Figure D.10: Variation in the interaction of log electricity price and pre-period sectoral electricity intensity




Note: Average zonal log electricity price multiplied by pre-period sectoral electricity intensity, split by electricity zones into high and low groups, and terciles of electricity intensity.

Figure D.11: Pre-post exposure gradient (zone–sector level)




Note: Unit of observation is a zone–sector cell. Pre-period is 2017–2019 and post period is 2022–2024. For each outcome, the figure plots the partial-regression relationship from $\Delta \log Y_{zs} = \theta_z + \kappa_s + \delta[\beta_s \Delta \log p_z^E] + u_{zs}$, where θ_z are zone effects and κ_s are sector effects. Both axes are residualised with respect to θ_z and κ_s . Dots are weighted by the corresponding pre-period outcome level: active firms in panel (a) and revenue in panel (b).

Figure D.12: Share of active firms by electricity intensity tercile



Note: Shares are computed within each treatment group-year among sectors with non-missing electricity cost-share tercile. Intensity terciles for sectors are determined in a sector cross-section across European countries.

Figure D.13: Average municipality-sector intensity tercile cell trajectories by price group



Note: Series plot the total index within high- and low-price zones by electricity cost share tercile, normalised to 2019 = 100. Sample excludes any zone-straddling municipalities. The number of active firms and total revenue is constructed from company financial statements.

D.2 Norwegian analysis: model-aligned robustness

Table D.6: Norwegian model-aligned moments, R01 sample

	Firm profit		Active firms		Total profits	
	(1)	(2)	(3)	(4)	(5)	(6)
Log electricity price x electricity share, fallback-filled	8.948 (7.801)	8.842 (7.759)	-7.473 (6.451)	-7.656 (6.417)	-3.385 (17.449)	-3.649 (17.400)
Log earnings x labour share, fallback-filled		1.182 (1.292)		3.804* (2.095)		4.222* (2.419)
Observations	23401	23401	12320	12320	9088	9088
R^2	0.747	0.747	0.670	0.671	0.432	0.432

Notes: Columns (1)–(2) estimate log operating profits among continuing firms with positive operating profit. Columns (3)–(4) estimate log active firm counts at the municipality–sector–year level. Columns (5)–(6) estimate log total operating profits at the municipality–sector–year level. Odd columns include only the electricity price $\times$ electricity share interaction; even columns also include the wage $\times$ labour share interaction. The sample is restricted to municipalities in border ring 0+1. All specifications include firm or municipality fixed effects as appropriate, zone-year fixed effects, and sector-year fixed effects. Standard errors are clustered by municipality.

Table D.7: Firm-level intensive-margin moments, unweighted and weighted specifications

	Unweighted		P95-weighted	
	(1)	(2)	(3)	(4)
	No wage	Wage	No wage	Wage
Log electricity price x electricity share	-0.778 (0.578)	-0.780 (0.578)	-1.530* (0.873)	-1.550* (0.875)
Log earnings x labour share		0.105 (0.158)		0.840*** (0.219)
Observations	440493	440493	440493	440493
R^2	0.776	0.776	0.814	0.814

Notes: Dependent variable is log operating profit among continuing firms with positive operating profit. Columns (1)–(2) are unweighted. Columns (3)–(4) use analytic weights equal to firm pre-period average revenue, winsorized at the 95th percentile and normalized to mean one. Columns (1) and (3) omit the wage interaction; columns (2) and (4) include it. The no-wage columns are restricted to the same non-missing wage-interaction sample as the wage columns. All specifications include firm, zone-year, and sector-year fixed effects. Standard errors are clustered by municipality.

D.3 Cross-country analysis: additional results

This appendix collects the cross-country estimates referenced in Section 4: the long-difference operating-revenue gradient and the Revelio Labs employment check. All specifications use the

Table D.8: Norwegian model-aligned moments across adjustment margins

	Revenue		Profit, IHS		Profit, log	
	(1)	(2)	(3)	(4)	(5)	(6)
	Unweighted	Weighted	Unweighted	Weighted	Unweighted	Weighted
Log electricity price x electricity share	-0.121 (0.327)	-0.924** (0.452)	-1.698 (4.568)	-10.197 (7.766)	-0.780 (0.578)	-1.550* (0.875)
Log earnings x labour share	0.243 (0.151)	0.287 (0.189)	5.150 (3.436)	5.751 (5.378)	0.105 (0.158)	0.840*** (0.219)
Observations	600055	600055	609628	609628	440493	440493
R^2	0.911	0.927	0.372	0.401	0.776	0.814

Notes: Columns report continuing-firm outcomes, each with and without the wage interaction. All specifications include the relevant location fixed effects, zone-year fixed effects, and sector-year fixed effects, as in equation (16). Standard errors clustered at the municipality level.

Table D.9: Norwegian aggregate moments across total-outcome transformations

	Total revenue		Total profit, IHS		Total profit, log	
	(1)	(2)	(3)	(4)	(5)	(6)
	No wage	Wage	No wage	Wage	No wage	Wage
Log electricity price x electricity share	-12.147* (7.220)	-11.700 (7.095)	3.032 (25.974)	4.333 (25.711)	-14.071** (6.838)	-13.838** (6.796)
Log earnings x labour share		11.382*** (1.824)		33.134*** (7.375)		5.667*** (1.650)
Observations	81767	81767	81767	81767	62036	62036
R^2	0.525	0.526	0.121	0.121	0.474	0.474

Notes: Columns report municipality–sector–year aggregate outcomes, each with and without the wage interaction. The outcomes are total operating revenue, inverse-hyperbolic-sine total operating profits, and log total operating profits. All specifications use the preferred model-aligned cost shares and include municipality fixed effects, zone-year fixed effects, and sector-year fixed effects, as in equation (16). Standard errors clustered at the municipality level.

harmonised sector loadings β_s (Exiobase electricity cost shares) and α_s (Eurostat SBS labour shares) and the same fixed-effect and clustering structure as the main cross-country tables.

Table D.10: Cross-country long differences: operating revenue

	(1)	(2)	(3)	(4)	(5)	(6)
	h=0	h=1	h=2	h=3	h=4	h=5
$\Delta^h \log \text{electricity price} \times \text{electricity share}$	-0.818 (0.618)	0.957 (0.979)	2.595** (1.068)	2.638*** (0.839)	3.737*** (0.971)	4.538** (1.621)
$\Delta^h \log \text{compensation} \times \text{labour share}$	0.334 (0.302)	0.691 (0.415)	0.629 (0.611)	0.370 (0.721)	-0.133 (0.871)	-0.385 (1.032)
Observations	12145974	10029111	8373925	6969616	5768816	4756979
R^2	0.228	0.330	0.415	0.494	0.556	0.602

Notes: Estimates of equation (18) for log operating revenue on the Orbis firm panel. The dependent variable is the long difference $\Delta^h \log Y = \log Y_{t+h} - \log Y_{t-1}$; columns correspond to horizons $h = 0, \dots, 5$, i.e. windows of one to six years. The reported coefficient is on $\beta_s \Delta^h \log p_{ct}^E$, the sector electricity cost share β_s times the long difference in the log national wholesale price; the labour term $\alpha_s \Delta^h \log w_{ct}$ is also included. All specifications include firm-country, country-year, and sector-year fixed effects. Sample as in Table 3. Standard errors, clustered by country, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.11: Cross-country levels: employment (Revelio Labs)

	(1)	(2)
Log electricity price $\times$ electricity share	0.325 (0.558)	0.210 (0.563)
Log compensation $\times$ labour share	0.212* (0.112)	
Observations	586864	586864
R^2	0.980	0.980

Notes: Estimates of equation (17) with log employee count as the dependent variable, on the Revelio Labs firm-country-year panel of employee counts constructed from online worker profiles. The reported coefficient is on the electricity term $\beta_s \log p_{ct}^E$, the interaction of sector electricity cost share β_s (Exiobase) with the log national wholesale price. Column 1 includes the labour-cost interaction $\alpha_s \log w_{ct}$ and column 2 omits it. All specifications include firm-country, country-year, and sector-year fixed effects. Standard errors, clustered by country, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.12: Cross-country long differences: employment (Revelio Labs)

	(1) h=0	(2) h=1	(3) h=2	(4) h=3	(5) h=4	(6) h=5
$\Delta^h \log \text{electricity price} \times \text{electricity share}$	0.386** (0.184)	0.452 (0.296)	0.533 (0.430)	0.515 (0.516)	0.901 (0.782)	1.043 (0.907)
$\Delta^h \log \text{compensation} \times \text{labour share}$	-0.021 (0.088)	-0.003 (0.073)	0.042 (0.097)	0.099 (0.147)	0.080 (0.192)	0.081 (0.190)
Observations	539461	490578	417356	375665	333485	295567
R^2	0.210	0.308	0.358	0.449	0.537	0.628

Notes: Estimates of equation (18) with the long difference in log employee count as the dependent variable, on the Revelio Labs firm-country-year panel. The dependent variable is $\Delta^h \log Y = \log Y_{t+h} - \log Y_{t-1}$; columns correspond to horizons $h = 0, \dots, 5$, i.e. windows of one to six years. The reported coefficient is on $\beta_s \Delta^h \log p_{ct}^E$, the sector electricity cost share β_s (Exiobase) times the long difference in the log national wholesale price; the labour term $\alpha_s \Delta^h \log w_{ct}$ is also included. All specifications include firm-country, country-year, and sector-year fixed effects. Standard errors, clustered by country, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


D.4 Cross-country analysis: German–Austrian bidding-zone split

As an illustration of the cross-country approach, this subsection studies the October 2018 separation of Austria from the common German–Austrian bidding zone. Before October 2018, the two countries shared a wholesale market area; after the split, Austria and Germany–Luxembourg formed separate price zones. The exercise uses the corrected monthly wholesale-price splice and the Eurostat non-household electricity-price series, and then matches Austria and Germany within harmonised SBS sectors. The results are intended as an appendix illustration of the price mapping and its reduced form, rather than as headline causal evidence: treatment is a single country-level market-design change, and the split may coincide with other country-specific adjustments.

The wholesale-price first stage is sharp. By construction, the wholesale log price gap is zero while the countries share a bidding zone; it reaches 0.040 log points in 2018, 0.055 in 2019, and 0.108 in 2022, before narrowing to 0.038 in 2024. A post-October indicator regression gives a wholesale gap of 0.069 log points (SE = 0.014). The downstream Eurostat first stage is less clean. Prices for the 500–1,999 MWh annual-consumption band show a pooled post-2018 Austria–Germany gap of 0.167 log points (SE = 0.042) when excluding all taxes, 0.169 (SE = 0.097) when excluding recoverable taxes and VAT, and 0.227 (SE = 0.101) including taxes. The second-half of 2018 is entered separately as a transition period; its estimates are small and imprecise. The Eurostat event profiles also contain pre-split movement and much of their widening occurs during and after the energy crisis. I therefore treat them as descriptive evidence on downstream industrial prices, rather than as a sharp pass-through design in their own right.

For the SBS exercise, let b_s denote the fixed model electricity share of sector s , P_{ct} the country-

Figure D.14: German–Austrian split: price first stages



Note: Event-study estimates of the Austria–Germany price gap relative to 2016. Panel A reports the wholesale spot-price gap; Panel B reports Eurostat non-household prices for the 500–1,999 MWh annual-consumption band under three tax concepts. The vertical line marks the October 2018 bidding-zone split; points show estimates and bars show 95% Newey–West confidence intervals.

year wholesale price, and $D_{ct} = 1\{c = \text{AT}, t \geq 2019\}$. The preferred specification uses country–sector, country–year, and sector–year fixed effects, and controls for the interaction between the national wage and the model labour share:

$$b_s \log P_{ct} = \pi(b_s D_{ct}) + \text{FE}_{cs} + \text{FE}_{ct} + \text{FE}_{st} + X_{ct,s}\gamma + u_{cst}, \quad (\text{D.1})$$

$$\log M_{cst} = \theta(b_s D_{ct}) + \text{FE}_{cs} + \text{FE}_{ct} + \text{FE}_{st} + X_{ct,s}\gamma + \varepsilon_{cst}. \quad (\text{D.2})$$

The event study uses 2016 as the reference year and includes separate 2017 anticipation and 2018 transition interactions. The SBS first stage follows the wholesale price mapping: its post-2018 coefficient is 0.075 log points per unit of b_s in the full 2015–2024 sample (0.083 in the 2015–2023 IV sample).

The reduced form shows a positive 2017 coefficient of about 3.1 (SE about 1.0–1.2) per unit of b_s , followed by negative estimates from 2019 through 2023. In the preferred 2015–2023 pooled specification, the post-2018 coefficient is -5.227 (SE = 3.019, $p = 0.091$). A p90–p10 difference in electricity shares, 0.0252, corresponds to an implied enterprise-count change of -12.3% (95% CI $[-24.8\%, 2.2\%]$). In the full event-study sample, the 2015 coefficient is statistically indistinguishable from zero ($p = 0.676$), but the 2017 estimate is statistically distinguishable from zero. It is therefore more informative to interpret 2017 as anticipation or pre-adjustment than as evidence supporting a clean parallel-trends design. The 2024 coefficient is positive and imprecise; because the SBS data contain later statistical and framework changes and many other shocks by that date, 2024 is excluded from the preferred pooled estimate.


Two limited diagnostics clarify the reduced-form pattern. The 2017 gradient is driven by Germany: between 2016 and 2017, the electricity-share slope is -4.016 (SE = 0.879) in Germany, compared with -0.106 (SE = 0.728) in Austria. Dropping each broad NACE section leaves the post-2018 coefficient negative. The 2017 lead changes sign when manufacturing (section C) is removed, from 3.107 in the full sample to -5.156 ; however, section C contains 19 of the 43 sectors and much of the identifying electricity-share variation, so this result is best read as an influence diagnostic, not as evidence from a narrow manufacturing shock.

Finally, a simple just-identified IV uses $b_s D_{ct}$ as an instrument for $b_s \log P_{ct}$ in the 2015–2023 sample. The second stage gives

$$\hat{\theta}_{IV} = -62.796 \quad (\text{SE} = 36.268), \quad p = 0.091, \quad (\text{D.3})$$

The Anderson–Rubin p -value is 0.091. At the estimated first-stage price wedge, the corresponding p90–p10 enterprise-count effect is -12.3% (95% CI $[-24.8\%, 2.2\%]$). The conventional sector-expanded first-stage F statistic is not a useful measure of instrument strength here: the same country-level price shock is mechanically repeated across sectors, inflating the apparent number of independent observations. Moreover, with one instrument and one endogenous regressor, the

Figure D.15: German–Austrian split: SBS reduced-form event study




Note: Right panel of the SBS price-mapping and reduced-form event-study output. Estimates are scaled to the p90–p10 difference in the fixed model electricity share and report the enterprise-count effect relative to 2016; bars are 95% confidence intervals. The dashed red line marks the 2018 split and the grey line the 2021 SBS boundary.

exactly identified IV estimate is the reduced form divided by the first stage; it rescales the reduced form rather than providing independent identification. The exercise therefore illustrates the price-to-outcome mapping, but the single-treated-country setting and the difficult exclusion restriction preclude treating it as a standalone causal estimate.

D.5 Cross-country analysis: interconnection event studies

This appendix provides a descriptive country-level analogue of the bidding-zone exercise. It asks whether selected cross-country interconnector creations were accompanied by changes in wholesale price gaps and in the relative distribution of enterprise counts across sectors with different electricity exposure. The exercise is not a causal design: the sample contains only three shocks, and the country-level changes may coincide with other economic and institutional developments. I therefore report the first stage and reduced form separately and do not construct an instrumental-variables estimate. The inclusion rule is fixed before inspecting outcomes: retain exact cross-country creations for which both endpoint countries appear in the EU24 SBS panel and for which the reference year, commissioning year, and first full post year remain inside a single SBS statistical framework. This rule selects NordBalt between Lithuania and Swedish bidding zone SE4 (commissioned February 2016), COBRACable between Danish zone DK1 and the Netherlands (September 2019), and Cirkovce–Pince between Hungary and Slovenia (June 2022). The commissioning year is a partial-exposure year; the following calendar year is the first full post year.

Figure D.16: Interconnector creations: national wholesale-price first stages



Notes: Each panel plots the signed monthly log wholesale-price gap, defined as the price in the endpoint that was more expensive before commissioning minus the price in the endpoint that was cheaper, normalized by its mean over the preceding twelve months. The navy series is the national country-price gap used for the SBS analysis; the thin line is the monthly series and the thicker line is a three-month moving average. The maroon dashed series is the corresponding endpoint-bidding-zone gap and validates the country mapping. The vertical line marks commissioning.


The national first stage is substantial and persistent only for NordBalt. Relative to the reference year, the national log gap falls by 0.430 log points in the commissioning year and by 0.550 in the first full post year. For COBRACable, the gap falls by 0.109 in the commissioning year but is 0.043 above the reference level in the first full post year. For Cirkovce–Pince, the corresponding changes are +0.005 and −0.032. These figures motivate showing the three events separately rather than pooling them: the direction and persistence of the price response differ, and the national price is the relevant exposure for country-sector outcomes.

For each event, the reduced form is estimated on exactly balanced endpoint-country-harmonized-sector cells. Let β_s denote the fixed electricity cost share of sector s , let R_{ec} equal +1/2 for the endpoint that was initially expensive and −1/2 for the initially cheap endpoint, and let T_e denote the commissioning year. The event-specific specification is $\log M_{cst} = \alpha_{cs} + \lambda_{ct} + \mu_{st} + \theta_{e,0}\beta_s R_{ec}1\{t = T_e\} + \theta_{e,1}\beta_s R_{ec}1\{t = T_e + 1\} + \rho\alpha_s \log w_{ct} + \varepsilon_{cst}$, where M_{cst} is the SBS enterprise count, α_{cs} , λ_{ct} , and μ_{st} are country–sector, country–year, and sector–year fixed effects, and the final term controls for the labour-share-weighted national wage. The reference year is omitted. The plotted coefficients scale each estimate by the common p90–p10 electricity-share difference, 0.0252, to express the implied enterprise-count change for a high–relative to a low-electricity-share sector.

The reduced-form estimates are imprecise and do not show a common response consistent with the first stage. For NordBalt, the scaled effect is +0.48% in the commissioning year and −4.05% in the first full post year, with very wide intervals. For COBRACable, the corresponding estimates are −4.74% and −8.14%. For Cirkovce–Pince, they are +2.48% and −5.16%. Thus, even the event with the clearest and most persistent price convergence does not deliver a precisely estimated, uniformly signed extensive-margin response, while the events with weak or nonpersistent price changes produce different reduced-form paths. The sector-clustered intervals should be read as exploratory sampling intervals: with one shock per event and only three events, they do not restore the shock-level inference required for a causal interpretation. I consequently treat these figures as an accessible country-level illustration of the price-to-enterprise-count mapping, not as evidence that interconnector construction caused the reported changes.

I exclude the remaining events for design reasons. ALEGrO crosses the 2020–2021 SBS framework break and has little local price incidence. NordLink and North Sea Link require Norway, and the latter also requires Great Britain, which are absent from the EU24 SBS panel. Viking Link requires Great Britain and its first full post year is the provisional 2024 release. The remaining observations are capacity expansions rather than clean creations; many dates are inferred from annual NTC jumps, and the pooled expansion price response is weak. Kriegers Flak is an expansion because the Kontek link already existed, and its outcome window also crosses the 2020–2021 SBS break.

Figure D.17: Interconnector creations: SBS extensive-margin reduced forms



Notes: Event-specific reduced-form coefficients from the balanced country-sector panels, scaled by the p90--p10 electricity-share difference. The reference year is normalized to zero; the commissioning year is a partial-exposure year and the full-post year is the first calendar year after commissioning. Bars are exploratory 95-percent intervals clustered by harmonized sector. They do not provide shock-level inference and the estimates are not causal.

D.6 Border versus bidding-zone price gaps: design and additional results

This appendix documents the border-versus-zone comparison used in Section 4 to establish that cross-border and within-country price gaps are the same economic object. The unit of observation is a bidding-zone pair by year, over 2015–2025, built from annual average day-ahead wholesale prices for every European bidding zone. The gap measure is the absolute difference in log annual prices between the two zones in a pair; the level gap in EUR/MWh is used as a robustness check. Pairs are classified into three types: within-country zone pairs, adjacent cross-border pairs, and non-adjacent cross-border pairs, where adjacency records whether the two zones share a physical interconnection. Interconnection intensity is measured as net transfer capacity relative to the smaller of the two zones' loads, so that a given border's capacity is expressed on the same scale regardless of zone size. Because zone fixed effects must not depend on the arbitrary alphabetical ordering within a pair, each zone enters through a symmetric indicator equal to the number of times it appears in the pair. Standard errors are clustered by zone pair. For the small Nordic sample, inference on the border coefficient is based on randomisation inference that permutes the border label across adjacent pairs over 1,000 draws, which does not rely on large-sample approximations with roughly twenty pairs.

Table D.13 reports mean gaps and the ratio κ of the mean adjacent-border gap to the mean within-country gap, for the Nordic and full European samples. In both samples cross-border gaps are no larger than within-country zonal gaps ($\kappa \approx 0.92$ in the Nordics, ≈ 1.08 in Europe), and the within-Norway crisis gaps that identify the main estimates exceed typical cross-border gaps, confirming that the price supports overlap.

Table D.13: Mean price gaps by pair type

(a) Nordic sample					
	Full period	Non-crisis	Crisis	EUR/MWh	Pairs
Within-country zones	0.286	0.183	0.559	15.6	17
Norway only	0.319	0.178	0.693	18.3	10
Border, adjacent	0.264	0.218	0.387	12.4	10
Border, non-adjacent	0.379	0.312	0.558	20.3	39
Ratio, adjacent/within	0.924	1.191	0.691	0.8	.
(b) European sample					
	Full period	Non-crisis	Crisis	EUR/MWh	Pairs
Within-country zones	0.156	0.113	0.259	9.0	38
Norway only	0.319	0.178	0.693	18.3	10
Border, adjacent	0.169	0.172	0.161	10.6	60
Border, non-adjacent	0.388	0.367	0.437	26.7	763
Ratio, adjacent/within	1.080	1.528	0.620	1.2	.

Notes: Cell means of the absolute log price gap by pair type, on the zone-pair panel over 2015–2025. Columns report the mean log gap over all years, in non-crisis years, and in the crisis years 2021–2023, the mean level gap in EUR/MWh, and the number of pairs. The *norway_only* row restricts within-country pairs to Norwegian bidding zones. The final row reports κ , the ratio of the adjacent-border mean to the within-country mean.

The full regression sequence behind the main-text panel — raw, capacity-conditional, headline with zone fixed effects, and the non-crisis and crisis subsamples — is reported in Table D.14 for the European sample and in Table D.15 for the Nordic sample, the single-market benchmark in which four multi-zone countries share one exchange and one market design. With roughly twenty adjacent Nordic pairs, the randomisation-inference p -value is the credible inference on the Nordic border coefficient, which is indistinguishable from zero.

Table D.14: Border premium in the European sample: full specification sequence

	(1)	(2)	(3)	(4)	(5)
	All pairs	Adjacent	Adjacent	Non-crisis	Crisis
National border	0.220*** (0.033)	0.048 (0.044)	0.148*** (0.049)	0.133** (0.054)	0.196** (0.085)
NTC / min. load		-0.022 (0.036)	0.027 (0.040)	0.010 (0.035)	0.081 (0.067)
Observations	8473	659	659	447	212
R ²	0.133	0.092	0.288	0.371	0.468
Zone FE	No	No	Yes	Yes	Yes
RI p-value (border)			0.002		

Notes: Regressions of the absolute log price gap on a national-border indicator, on the European sample of the zone-pair panel over 2015–2025. Column 1 reports the raw premium with year fixed effects on all pairs; column 2 adds interconnection capacity (net transfer capacity over minimum zone load) on adjacent pairs; column 3 adds symmetric zone fixed effects (the headline specification); columns 4 and 5 estimate the headline specification on non-crisis and crisis years. Standard errors clustered by zone pair in parentheses. The randomisation-inference p -value permutes the border label across adjacent pairs over 1,000 draws. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.15: Border premium in the Nordic sample: full specification sequence

	(1)	(2)	(3)	(4)	(5)
	All pairs	Adjacent	Adjacent	Non-crisis	Crisis
National border	0.070 (0.054)	0.102* (0.054)	0.130*** (0.034)	0.094*** (0.023)	0.208** (0.091)
NTC / min. load		0.011 (0.030)	0.058* (0.033)	0.031* (0.018)	0.117 (0.080)
Observations	726	180	180	126	54
R ²	0.428	0.359	0.432	0.479	0.450
Zone FE	No	No	Yes	Yes	Yes
RI p-value (border)			0.377		


Notes: Regressions of the absolute log price gap on a national-border indicator, on the Nordic sample of the zone-pair panel over 2015–2025 (NO, SE, DK, FI bidding zones). Column 1 reports the raw premium with year fixed effects on all pairs; column 2 adds interconnection capacity (net transfer capacity over minimum zone load) on adjacent pairs; column 3 adds symmetric zone fixed effects (the headline specification); columns 4 and 5 estimate the headline specification on non-crisis and crisis years. Standard errors clustered by zone pair in parentheses. The randomisation-inference p -value permutes the border label across adjacent pairs over 1,000 draws and is the credible inference for this sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The dispersion in these gaps reflects physical integration rather than the border line itself. Figure D.18 shows event studies around changes in grid configuration: when new zone borders


are switched on (the Swedish 2011 and German–Austrian 2018 splits) only small gaps open, while interconnector openings visibly close the gap between the connected zones. Table D.16 quantifies the latter with stacked signed-gap difference-in-differences: interconnector creations close pair gaps by about 0.26 log points, or 0.30 per gigawatt of added capacity, whereas validated capacity expansions on existing links have no significant additional effect.

Figure D.18: Grid-configuration events and price gaps

(a) Zone-border activations (Swedish and German–Austrian splits)



(b) Interconnector openings (seven events, stacked)



Note: Panel (a) plots the mean absolute log gap for treated and control pairs around the Swedish zonal split (2011m11) and the German–Austrian split (2018m10). Panel (b) plots the mean absolute log gap for treated borders and control pairs around seven interconnector commissionings, stacked in event time. The vertical line marks the event month.

Table D.16: Signed price convergence at interconnector events

	(1) Creations DiD	(2) Creations /GW	(3) Val. exp. DiD	(4) Val. exp. /GW
Post x treated	-0.262** (0.124)		-0.028 (0.065)	
Post x treated x GW		-0.295** (0.121)		-0.007 (0.070)
Observations	13354	13354	24249	24249
R^2	0.421	0.423	0.374	0.374

Notes: Stacked difference-in-differences on the signed pair gap (log price of the pre-event expensive zone minus that of the cheap zone), around interconnection events with pair and event-month fixed effects. Columns 1–2 use interconnector creations, columns 3–4 validated capacity expansions on existing links; the odd columns report the average post-event effect and the even columns the effect per gigawatt of added capacity. Standard errors clustered by event-pair in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.17 reports pooled AR(1) regressions of the pair gap with a border interaction, together with the share of pairs whose signed gap never changes sign over the sample. Cross-border gaps mean-revert faster — the AR(1) coefficient is lower by about 0.22 — and are somewhat less sign-stable than within-country zonal gaps, the one dimension on which the price variation is not identical across the two settings.

Table D.17: Persistence of price gaps by pair type

	(1) Europe	(2) Nordics
L.Abs. log price gap	0.769*** (0.100)	0.715*** (0.098)
National border=1 $\times$ L.Abs. log price gap	-0.219* (0.114)	-0.222* (0.113)
National border=0	0.000 (.)	0.000 (.)
National border=1	0.042** (0.021)	0.091*** (0.030)
Constant	0.038** (0.017)	0.059** (0.025)
Observations	757	209
R^2	0.362	0.524
Sign-stable share, within	0.385	0.333
Sign-stable share, border	0.317	0.200

Notes: Pooled AR(1) regressions of the absolute log price gap on its lag, a national-border indicator, and their interaction, with year fixed effects, on adjacent pairs. Column 1 uses the European sample, column 2 the Nordic sample. The bottom rows report the share of pairs whose signed gap never changes sign over 2015–2025, separately for within-country and cross-border pairs. Standard errors clustered by zone pair in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Two further tables support the headline. Table D.18 adds log distance and adjacency to the raw specification in the spirit of Engel and Rogers (1996); geodesic distance is a poor proxy for grid topology and is reported only for comparability, and the conditional border coefficient remains a modest positive premium in Europe and near zero in the Nordics. Table D.19 confirms that the European premium is not an artefact of the log transformation, mixed-source pair-years, the capacity denominator, or the fixed-effect implementation: level gaps, a mixed-source exclusion, a minimum-direction capacity measure, and a zone-by-year specification all reproduce a modest, significant border premium.

Table D.18: Gravity specification for price gaps

	(1)	(2)	(3)	(4)
	Nordics	Nordics	Europe	Europe
National border	-0.011 (0.037)	0.016 (0.037)	-0.041 (0.031)	0.215*** (0.046)
Log distance	0.172*** (0.042)	0.098** (0.046)	0.294*** (0.018)	0.239*** (0.017)
Adjacent	-0.056 (0.058)	-0.123** (0.052)	0.132*** (0.025)	0.084*** (0.028)
Observations	726	726	8473	8473
R ²	0.498	0.534	0.277	0.518

Notes: Regressions of the absolute log price gap on a national-border indicator, log geodesic distance between zone centroids, and an adjacency indicator, with year fixed effects. Columns 1–2 use the Nordic sample and columns 3–4 the European sample; the even columns add symmetric zone fixed effects. Standard errors clustered by zone pair in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.19: Robustness of the border premium

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	EUR gap N	No mixed N	Min NTC N	EUR gap EU	No mixed EU	Min NTC EU	ZxY FE EU
National border	7.403*** (2.362)	0.130*** (0.034)	0.125*** (0.031)	7.721*** (2.764)	0.138*** (0.048)	0.148*** (0.043)	0.136* (0.076)
NTC / min. load	3.269 (2.104)	0.058* (0.033)		1.374 (2.266)	0.031 (0.037)		-0.445** (0.174)
NTC (min dir.) / min. load			0.067** (0.032)			0.032 (0.039)	
Observations	180	180	180	659	544	660	411
R ²	0.496	0.432	0.434	0.428	0.299	0.288	0.826

Notes: Robustness of the capacity-conditional border coefficient. Columns 1 and 4 use the level gap in EUR/MWh; columns 2 and 5 exclude pair-years combining prices from different sources; columns 3 and 6 use a minimum-direction net-transfer-capacity measure; column 7 replaces symmetric zone fixed effects with zone-by-year fixed effects. Columns 1–3 use the Nordic sample, columns 4–7 the European sample. All specifications restrict to adjacent pairs and include year fixed effects. Standard errors clustered by zone pair in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E Firm selection at entry (Melitz extension)

A natural extension allows for firm selection at entry, following Melitz (2003). In this version, M_s is endogenous: potential entrants pay a fixed entry cost f_s^e and draw a portable productivity η_{fs} from a common Pareto distribution with shape $k > \sigma - 1$. This portable draw augments the baseline firm–location match a_{fjs} , which still smooths the location choice, with a component common across locations, so that more productive firms earn higher profits everywhere and selection operates at the entry margin. Firms enter if and only if their expected profits across locations exceed the entry cost. The free-entry condition is

$$\int_1^\infty \max_j \left\{ \frac{\pi_{js}(\eta)}{1 - \delta} \right\} k \eta^{-k-1} d\eta = f_s^e$$

This extension adds a selection margin to the baseline location mechanism. High-price countries attract fewer firms, and the firms that do enter are selected from a different part of the productivity distribution. Selection can partly offset the loss from lower firm mass because surviving firms are more productive on average, while agglomeration still makes the location of activity matter for aggregate productivity. The net TFP effect therefore depends on the relative strength of location misallocation, agglomeration, and selection. Empirically, this margin would be disciplined by comparing the productivity distribution of entrants and continuing firms across high- and low-price zones.

The main model abstracts from this extension for three reasons. First, location reallocation (M_{js}) is the first-order mechanism for spatial misallocation when $\gamma > 0$: agglomeration means that *where* production is located determines productivity, so the location margin is the primary object of interest. Second, the Norwegian evidence locates the adjustment in the reallocation of firms across locations rather than in the total number of firms: the cross-sector firm-count response is strong (Section 4), but aggregate firm counts evolve almost identically across price zones (figure D.13), suggesting that the aggregate entry margin free entry would endogenise — the total mass M_s and the selection of entrants — is not the dominant adjustment over the horizon studied. Third, adding free entry and selection would require solving an additional cutoff condition jointly with the spatial equilibrium, reducing the transparency of the comparative statics. The extension is therefore left as a robustness exercise for future quantitative work rather than built into the baseline model. More precisely, the baseline represents the extensive margin as the reallocation of a fixed total mass M_s across locations, while the Norwegian evidence identifies it through changes in active firm counts; because identification is cross-sectoral, the difference-in-differences recovers the same location semi-elasticity $1/(\psi - \gamma)$ under either reading, with the entry/exit margin the empirical counterpart of the free-entry margin activated here.

F Quantification details

This appendix documents the solved baseline: where each country sits on its merit-order supply curve, how the calibrated equilibrium performs against moments the calibration does not target, the numerical verification of the solution, and the algorithm used to compute it. The calibration inputs themselves are documented elsewhere: parameter values and their sources in Table 4 of Section 5.2, the construction of the supply-side primitives in Appendix C, the sensitivity of the results to the parameter ranges in Figure F.3 below, and the specification choices the ranges do not cover in Appendix G.

Baseline supply position. Table F.1 records where each country sits on its merit-order supply curve in the solved baseline, together with the wholesale price, fossil utilisation, and the supply elasticity ε_j that follows from equation (7). All twenty-three countries are on the fossil margin, so the log-linearised comparative statics of Section 3.5 apply directly throughout. Lithuania is the marginal case, importing the bulk of its consumption and running its own fossil fleet at 3% utilisation, so its price sits only just above the flat portion of its domestic curve. The ordering of ε_j is governed by spare fossil capacity relative to demand rather than by utilisation alone: France combines a moderate utilisation rate with the smallest fossil fleet relative to its own generation in the sample, and has the steepest supply curve as a result, while Lithuania holds fossil capacity seven times its own generation and has the flattest.

Untargeted-moment validation. Table F.2 reports the first-pass model fit against moments the calibration does not target, computed on the raw gross-output-share baseline. The headline row compares the inverted country–sector productivity shifter with measured labour productivity from structural business statistics: the two correlate at 0.78 across the 299 country–sector cells (Figure F.1), a genuine over-identifying check because the measured series enters nowhere in the calibration. The model also reproduces cross-country output shares (proxied by included-sector value added) and the level of wholesale electricity prices, the latter disciplined in part by the supply-side calibration and so reported as context. The remaining rows are shown for transparency rather than as fit tests: measured productivity alone is a weak fixed-price predictor of firm location (a correlation of 0.17), which motivates the productivity-measure upgrade noted in Section 5.4, while the baseline firm distribution and cross-country wages reproduce the data almost exactly because they are targeted by the inversion. The measured-productivity proxy is provisional and will be replaced by CompNet and EU KLEMS productivity.


Numerical verification. Separately from the economic fit, I verify that the solved equilibrium satisfies its own equations. Market-clearing and accounting residuals fall below 10^{-10} ; the

Table F.1: Baseline supply position by country

Country	Price (EUR/MWh)	Fossil util. (%)	Spare fossil cap. (%)	ε_j	ε_j at $f_j = 1$	Supply segment
France	42.4	38.3	61.7	2.452	2.745	Fossil margin
Bulgaria	44.2	70.3	29.7	2.156	2.324	Fossil margin
Sweden	35.6	36.2	63.8	1.868	2.148	Fossil margin
Austria	37.9	47.2	52.8	1.215	1.381	Fossil margin
Slovakia	39.1	42.8	57.2	1.213	1.355	Fossil margin
Belgium	42.4	51.2	48.8	1.184	1.330	Fossil margin
Slovenia	44.1	43.8	56.2	1.108	1.194	Fossil margin
Portugal	48.6	55.7	44.3	1.036	1.128	Fossil margin
Czechia	37.9	53.6	46.4	0.847	0.929	Fossil margin
Finland	36.6	41.1	58.9	0.818	0.929	Fossil margin
Germany	37.9	53.0	47.0	0.812	0.899	Fossil margin
Netherlands	41.4	58.5	41.5	0.738	0.819	Fossil margin
Romania	44.1	39.1	60.9	0.668	0.728	Fossil margin
Poland	41.8	52.6	47.4	0.571	0.618	Fossil margin
Spain	48.6	31.0	69.0	0.481	0.527	Fossil margin
Greece	53.5	38.8	61.2	0.420	0.453	Fossil margin
Estonia	37.1	48.0	52.0	0.418	0.495	Fossil margin
Italy	51.8	37.6	62.4	0.408	0.446	Fossil margin
Denmark	35.6	34.2	65.8	0.391	0.439	Fossil margin
Croatia	44.1	17.7	82.3	0.358	0.393	Fossil margin
Hungary	44.1	23.8	76.2	0.345	0.380	Fossil margin
Latvia	40.0	28.1	71.9	0.332	0.374	Fossil margin
Lithuania	40.0	3.0	97.0	0.065	0.073	Fossil margin


Note: Solved baseline equilibrium, EU-23, 2015–2019 calibration window. Price is the model-solved wholesale price. Fossil utilisation is fossil generation over fossil capacity, and spare fossil capacity is its complement. The supply elasticity $\varepsilon_j \equiv \partial \ln p_j^E / \partial \ln E_j^D$ is computed from equation (7) as $(\chi_j / \zeta) E_j^D / [K_j^F (1 - u_j)]$, where $\chi_j = (p_j^E - c_0^E) / p_j^E$ is the fuel share of the price; the final column reports the simpler $\chi_j = 1$ benchmark that ignores the non-fuel cost component. Rows are sorted by ε_j , descending. All twenty-three countries are on the fossil margin. Lithuania has the flattest curve in the sample: it imports 10.3 of its 12.5 TWh of demand, so its own fossil fleet runs at 3% utilisation and its solved price of 40.0 EUR/MWh sits only just above the flat portion of its domestic curve.

Figure F.1: Untargeted validation: model-inverted productivity shifter versus measured labour productivity



Note: Model-inverted country–sector base productivity Φ_{js} against measured apparent labour productivity (gross value added per person employed, Eurostat structural business statistics, 2015–2019 average), both in logs and demeaned by sector, across the 299 country–sector cells of the EU-23 sample. The measured series enters the calibration nowhere. The dashed line is the fitted regression rather than the 45-degree line: the two series are commensurate in units but not in dispersion, since the inverted shifter is an equilibrium object with a standard deviation of 0.16 log points against 0.61 for measured productivity, so the fitted slope is 0.21 and the check is one of co-movement. Labelled points are the cells the fit misses by the most. The two largest are Danish and Dutch mining and quarrying (section B), where the measured series is dominated by offshore North Sea oil and gas extraction: apparent labour productivity there reaches 963 and 624 thousand euros per person employed against a sample median of 45, reflecting resource rents and capital intensity in an extractive activity rather than a location advantage of the kind the location model represents. The inversion accordingly places both well below what the measured series alone would predict, by 0.50 and 0.37 log points. Excluding the two raises the correlation from 0.78 to 0.82.

Figure F.2: Out-of-sample validation: model crisis-shock prices versus observed 2021–2023 prices




Note: Solved country wholesale prices when the 2015–2019 calibrated economy is subjected to the 2021–2023 crisis environment (crisis-period fossil supply costs, interconnector capacities, and residual electricity load, with all parameters and fundamentals held at their baseline calibration), against observed 2021–2023 average wholesale prices, EU-23, with the 45-degree line. The correlation of log prices is 0.73; the dispersion of log prices is 0.18 in the model against 0.19 observed. As with the baseline wholesale-price moment, the crisis fossil supply anchoring draws on observed prices, so the discipline falls on the equilibrium demand, trade, and reallocation response rather than on the price level. The fit is in relative rather than absolute prices. Modelled prices lie below observed ones in all 23 countries, by 0.30 log points on average, because the country-specific fossil-segment wedge κ_j is calibrated to the 2015–2019 window and held fixed: the crisis solve therefore prices the marginal plant at crisis fuel and carbon costs but carries only baseline scarcity rents, and misses the premia over marginal cost that scarcity, hoarding, and hydro opportunity cost added during 2021–2023. The exercise is accordingly informative about the geography of crisis prices, which is what the reallocation response depends on, and not about their level. The largest outlier is France, whose observed crisis price reflects the exceptional nuclear outages of 2022, when several reactors underwent extended maintenance, reducing low-cost supply — a capacity shock the model does not capture, as it holds nuclear capacity at its baseline level.

Table F.2: Untargeted-moment validation (first pass)

Moment	Role	n	Correlation	R^2
Productivity shifter vs. measured productivity	headline validation	299	0.78	0.61
Country output/GDP shares	untargeted (output proxy)	23	0.67	0.45
Wholesale electricity prices	context (supply-side calibration)	23	0.97	0.94
Measured productivity alone $\rightarrow$ firm shares	fixed-price diagnostic	299	0.17	0.03
Cross-country wages	equilibrium context	23	1.00	1.00
Baseline firm distribution	calibration identity	299	1.00	1.00

Note: Model-versus-data fit for untargeted moments, on the raw gross-output-share baseline (EU-23). Correlations and R^2 are across the units in column n (country-sector cells or countries). The top three rows are validation or context moments the calibration does not target; the bottom three are shown for transparency: the fixed-price diagnostic feeds measured productivity into the location equation at fixed prices and wages, while the baseline firm distribution and cross-country wages are reproduced by construction. Measured productivity is a provisional Eurostat structural-business-statistics labour-productivity proxy.

Figure F.3: Sensitivity of the full-union TFP gain to calibration parameters



Note: Left panel: the full-union first-pass TFP gain under one-at-a-time parameter sweeps — γ over $[0, 0.08]$, σ over $[3, 6]$, and ζ over the $\{1.5, 2, 2.5, 3\}$ grid — and a 95% band propagating the standard error of the estimated location gradient δ_M through the calibration. Diamonds mark two earlier calibration vintages: the archived nine-sector calibration and the thirteen-sector run on the pre-update interconnector network. Right panel: the Norwegian extensive-margin gradient δ_M with 95% confidence intervals when the sector exposure is measured at two-digit, thirteen-sector, and nine-sector granularity. Moving from two-digit exposure to the thirteen-sector partition the model solves on shifts the gradient by 0.50, about a sixth of a standard error, so the calibration is insensitive to granularity at the model's own resolution. Coarsening further, to the nine NACE sections, makes the gradient 3.53 larger in magnitude, about nine tenths of a standard error; the three confidence intervals overlap throughout, so this is a movement of the point estimate within sampling error rather than evidence against the coarse aggregation. Its sign is the conservative one: a larger $|\delta_M|$ implies a smaller $\psi - \gamma$ and hence a larger union gain, so calibrating at the model's own resolution understates rather than flatters the counterfactual. The left panel's δ_M band already carries this: it perturbs δ_M by one and by 1.96 calibration standard errors — a range that contains the nine-sector point estimate — and moves the headline gain only between 5.84 and 6.32 basis points.

market-coupling dispatch satisfies its Karush–Kuhn–Tucker conditions with zero complementarity slack; input validation returns no failures (100 checks pass, 5 warnings, no failures); the four pre-counterfactual stress shocks solve with economically signed responses; and the outer fixed point is locally stable, with a spectral radius of about 0.79 under the baseline damping.

Solution algorithm. The calibrated equilibrium is solved by the following nested fixed-point iteration, initialised at the baseline or, for a counterfactual, at the solved baseline allocation.

1. *Location and labour.* Given electricity prices p_j^E , wages w_j , and demand shifters B_s , evaluate the location values and set $M_{js} = M_s \text{softmax}_j(v_{js})$; adjust each w_j to clear the modelled-sector labour market $\sum_s L_{js} = \bar{L}_j$, iterating location and wages to convergence.
2. *Electricity market.* Add the fixed residual load H_j to price-responsive firm electricity demand. Without trade, clear each country on its supply schedule (Lemma 2). With trade, solve the market-coupling dispatch as one convex welfare program: maximise demand benefit net of generation cost over country firm demand, signed interconnector flows, and generation by source, subject to nodal balance and the net-transfer-capacity and generation-capacity bounds; nodal prices, flows, and congestion rents are the solution's duals. The segmented and full-union scenarios are the same program with the transmission bounds set to zero or made slack.
3. *Goods market.* Given the updated locations, wages, and prices, compute the sectoral price indices P_s and aggregate income I , and update the demand shifter $B_s/B_s^0 = (I/I^0)(P_s/P_s^0)^{\sigma-1}$.
4. *Outer update.* Combine steps 2–3 into a new outer state $(\ln p_j^E, \ln B_s)$ and take a damped step with safeguarded Anderson acceleration, warm-starting the dispatch from the previous flows and falling back to the damped Picard step when the accelerated step is not an improvement.
5. *Convergence and stability.* Repeat until the electricity-price, goods-market, labour-market, and network-balance residuals fall below tolerance. At the solution, verify local stability from the spectral radius of the Jacobian of the damped update.

Counterfactuals repeat steps 1–5 with the perturbed primitives (generation capacity, transmission bounds or scenario, or a single union-wide price), warm-started from the baseline solution.

G Robustness and extensions

This appendix reports the specification choices and modelling restrictions flagged in Section 6.1. Parameter uncertainty is collected in Figure F.3; what follows concerns choices of functional form and scope, which the sensitivity ranges do not cover.

G.1 Substitution between electricity and labour

The baseline production function is Cobb–Douglas, so electricity cost shares are technological constants. Relaxing this to a CES form with elasticity η between electricity and the non-electricity composite leaves the model unchanged at $\eta = 1$ and makes the cost function convex below it, so that mean-preserving price convergence delivers a first-order gain by Jensen’s inequality.

Solving the union at $\eta \in \{1, 0.75, 0.5, 0.3, 0.15\}$ shows that this gain is real but sectoral. The aggregate falls from 5.95 to 5.36 basis points as η drops to 0.15, because inelastic electricity demand shrinks the volume of fossil generation the union displaces and hence the fuel bill saved; the location channels move by less than 0.1 bp. Manufacturing rises from 13.4 to 14.2 bp over the same range and the exposed divisions from 35.1 to 38.0 bp, since it is the high- β_s sectors whose cost function becomes materially convex. Substitution redistributes the gain toward exposed sectors rather than scaling it.

The long-run elasticity close to unity discussed in Section 3.7 is an aggregate-energy object, reached through directed technical change. For electricity in manufacturing specifically the long-run estimates are lower: translog estimates on capital, labour, energy and materials put it between 0.27 and 0.44, and electricity is the poorest substitute among energy carriers. This is why the union is solved across the range rather than at unit elasticity alone. Within that range, unit elasticity is the generous choice for the aggregate and the conservative one for the sectoral figures, and the whole range moves the economy-wide gain by about a tenth, entirely through the fuel bill.

G.2 The fossil supply curve

Curve shape. Country supply curves are drawn from a common Pareto heat-rate distribution scaled to each country’s fossil capacity. The alternative truncates that common distribution instead, so that countries with small fossil fleets retain only their most efficient plants and face flatter curves. The implied elasticities differ sharply for those countries, by a factor of roughly seven for France and by one to two orders of magnitude for Sweden and Austria. The union gain rises by nine percent under truncation, entirely through the fuel-saving channel, and the firm allocation is invariant to four decimal places. The constrained counterfactuals are more exposed, since near-flat country curves resist the price convergence a partial grid would otherwise deliver: the planned-grid exercise falls by about a quarter and the natural-potential comparison by considerably more.

Neither specification is identified by the available plant data, which resolve dispersion only for the largest fleets and are silent precisely for the small clean-fleet countries where the two differ. The common-Pareto form is retained on parsimony and on the implausibility of the near-price-taking elasticities truncation implies for those countries; the alternative is treated as a specification band on the constrained legs rather than as a competing calibration.

Slopes by fuel. The headline dispatch treats fossil generation as a single composite block with a common curvature ζ . A disaggregated version calibrates separate blocks by fuel, with steeper curvature for gas and for coal than the pooled fit, and is used for the emissions accounting of Section 7.4. Solved as an equilibrium alternative it leaves the aggregate union gain essentially unchanged. The composite is retained as the headline because per-fuel projections are not available for the forward-looking scenarios.¹²⁹

Calibration provenance. The heat-rate distribution is fitted to plant-level efficiencies, which are observed for 353 plants of the active fossil fleet, 315 of them German. Germany is the only country whose fleet supports a credible independent fit, at $\zeta = 1.90$ against the pooled value of 2; the remaining country fits rest on five to twelve plants each, and eighteen of the twenty-three model countries have no efficiency observations at all. Observed frontier heat rates are close to common across the countries where they can be measured, between 1.59 and 1.74, so the common frontier is not the binding assumption. The common curvature is: it is an assumption for most of the sample rather than a validated fact. The electricity union is not exposed to it, since a single clearing price places every country on the pooled curve; the baseline congestion term and the constrained-integration counterfactuals are.

Placement of the country cost wedge. Each country carries a multiplicative wedge on its fuel term, reconciling the modelled merit order with observed prices. Because both the multiplicative and an additive placement are re-anchored to the same observed price and utilisation, the baseline supply elasticities are identical under either, and the choice bears only on how fuel and carbon price changes pass through in counterfactuals. The wedges are dispersed between roughly 0.76 and 1.21 across countries, and that dispersion is not fleet vintage: country frontier efficiencies explain almost none of it. It is better read as rents, scarcity and market structure, which is what a multiplicative wedge is designed to absorb.

Curtailed renewable output. Baseline low-cost supply is calibrated to observed generation, which already embeds the curtailment that segmentation itself causes, so the union is denied cheap volume that integration would in part release. Allowing the counterfactual to draw on it raises the gain by about 0.4 basis points on a deliberately generous release and by roughly half that on measured curtailment, at a near-constant 0.014 bp per TWh released. The entire move is the fuel-saving channel. The channel is bounded by what cross-border integration can actually relieve: of some 15 TWh/year of recent European curtailment, the large majority is internal

¹²⁹Blocks whose observed generation meets or exceeds registered capacity, mostly bioenergy and other fossil, about 8% of fossil generation, are treated as near must-run: capacity is set so that baseline generation runs at 90% utilisation and the cost scale is anchored to the observed price.

congestion, most of it German north–south redispatch, which coupling across borders does not address.

G.3 Endogenous generation capacity

The counterfactuals hold generation capacity fixed, so integration changes how existing plant is dispatched and where firms locate, but not what Europe builds. Since integration changes the prices investors expect to receive, this restricts a margin that plausibly responds.

The structure of the model confines how much is at stake. Under a union every country faces the same electricity price, so electricity drops out of relative location incentives and the firm allocation is the same whatever has been built (Proposition 6). Capacity can therefore reach the union gain only through the fuel bill. Two consequences follow. The within-manufacturing figure of 13.4 basis points is untouched, because it carries no fuel term. And of the 5.95 bp economy-wide gain, only the 1.89 bp fossil-resource component is exposed at all; the 4.06 bp of location channels are fixed by the invariance. This holds while baseline capacity stays at its observed level, which the last paragraph relaxes.

Gonzales, Ito and Reguant (2023) evaluate the Chilean interconnection twice, once holding solar capacity fixed and once letting free entry respond, which is the contrast at issue here. Their entry condition pins revenue per installed MW, so adjustment comes as quantity rather than price: without the interconnection, capacity in the two solar zones falls to 20% and 17% of what was actually built, while revenue per MW is unchanged. Their generation-cost saving from integration rises from 73 to 219 million USD per year once investment responds, a factor of three; their internal rate of return, which is net of the capital cost of the induced capacity, rises from 7.0% to 19.7%, a factor of 2.8. Applying those factors to the fuel-saving component alone raises the economy-wide gain from 5.95 to between 7.1 and 9.7 basis points. A local capacity supply curve calibrated to plausible investment elasticities, and closed with a cannibalisation fixed point on the union price, gives 6.5 to 7.8 bp on unrelated evidence.

A natural objection is that European capacity in the calibration window was built under feed-in tariffs, auctions and contracts for difference rather than at wholesale prices, so that changes in wholesale prices need not reach investment at all. The Chilean evidence cuts the other way. Solar entered the Atacama steadily between 2015 and 2017 while the local spot price sat near zero, which is rational only under the anticipation that the announced interconnection would raise it, and the auction contracts of those years cleared at 78 and 52 USD/MWh in expectation of long-run local prices. Long-term contracting transmitted the expectation of integration into investment rather than insulating investment from it.

That calculation prices only the capacity integration induces. Capacity responding to price under segmentation runs the other way: it is drawn to the countries where electricity is most expensive, compressing the dispersion integration exists to remove. The two cannot be netted here,

because the location channel depends on which country's price moves and not only on how far dispersion falls.¹³⁰ The net direction of the two channels is therefore not signed. What the exercise establishes is the scope rather than the sign: the manufacturing result is unchanged throughout, and the economy-wide figure is exposed only through its fuel-saving component.

G.4 Endogenous generation capacity under partial integration

The argument above rests on a single price, and a realistic grid does not deliver one. Where prices stay dispersed, new capacity changes relative prices, the location channels are exposed, and the invariance property no longer applies, so this case has to be solved rather than transported.

Free entry is a price ceiling, and a price ceiling is a flat segment in the merit order: investors supply any quantity at a country's break-even price and nothing below it. It therefore enters as one more generation block per country, priced at that threshold and bounded by remaining technical potential, and the market-clearing conditions deliver the entry condition without iteration and without assuming which countries share a price. The threshold is recovered from the observed capacity stock, as Gonzales, Ito and Reguant (2023) recover their entry cost. Because the block carries its levelised cost, the counterfactual pays for the capacity it induces instead of receiving it free.

Solved on the planned grid, this induces 235 TWh of new capacity in eight countries, two thirds of it in Denmark and Sweden, and displaces 157 TWh of fossil generation. The grid's contribution to price convergence does not survive it. The grid works by raising prices where electricity is cheap, and that is exactly where entry becomes profitable, so the new capacity competes the increase away and returns those countries to their own thresholds; the expensive countries are left untouched, because congestion stops the additional output reaching them. Cross-country dispersion widens rather than narrows, from a standard deviation of 2.4 to 3.8 EUR/MWh, and the location channels turn negative, at -0.7 against $+0.4$ basis points: the planned grid combined with an entry response allocates production slightly worse than no grid at all. What survives is a fuel bill that more than triples, from 1.1 to 3.9 bp, against a capital charge of 0.9 bp for the new capacity. The grid is worth more once generation responds, 2.3 bp against 1.5, but it has stopped being a lever on where production locates and become one on what generation costs.

That capital charge enters the productivity measure the way the fuel bill does, through the share of it that leaves Europe, since payments to European factors are internal transfers exactly as domestic operating costs already are in the fossil term. No published estimate of the local-content share of European renewable projects exists, so the share is constructed: component cost weights are combined with import shares anchored on trade and manufacturing-capacity evidence, giving

¹³⁰ Across the partial-integration scenarios of Section 7.1, two variants of the same benchmark sit at almost identical price dispersion and yet differ by 1.5 basis points in their location channel, and the scenario reaching the lowest dispersion of any returns among the smallest location gains. No dispersion-based scaling can therefore price the offsetting channel.

0.14 for onshore wind and 0.36 for utility-scale solar, and these are weighted by the technology composition of the induced capacity itself, which is 82% wind because entry concentrates in countries whose solar resource is poor. The resulting share is 0.18. This is a construction rather than a measurement, and its weakest input, the turbine’s share of onshore project cost, has cross-checks that all sit below the value used, so the charge is if anything overstated.

The conclusion is accordingly conditional on that parameter, and the relevant threshold is close enough to state. The entry response stops adding value at an import share of 0.33, nearly twice the estimate used but below the figure for utility solar alone, so a build dominated by solar rather than wind would overturn it. What does not depend on the parameter is the finding that motivates this exercise: the capital charge falls entirely on the fuel-saving channel, so the sign of the location channels, and with it the conclusion that a planned grid stops improving the allocation of production once generation is free to follow prices, holds across the whole range.

G.5 Transmission losses

The baseline counterfactual abstracts from electricity lost during cross-border transmission. I relax this assumption by treating transmission as an iceberg technology. For each directed flow $f_{ab} \geq 0$ sent from country a to country b , the receiving country gets

$$q_{ab} = (1 - \lambda)f_{ab}, \quad \ell_{ab} = \lambda f_{ab}, \quad \lambda = 0.03,$$

where q_{ab} is the electricity delivered and ℓ_{ab} is lost in transit. The lossy dispatch uses separate non-negative sending-end flows for the two directions of each interconnector; the original network topology and capacity convention are retained, with the full-union bounds made slack as in the lossless union. Because receiving-country nodal balances count q_{ab} , the additional generation required to cover ℓ_{ab} enters generation costs and the fuel bill automatically. For a positive uncongested flow from a to b , the corresponding price condition is $p_a^E = (1 - \lambda)p_b^E$, rather than equality of the two prices.

The 3% loss factor is a representative central value rather than an estimated model parameter. Nord Pool documents the sending- versus receiving-end accounting and gives a 2.4% worked example, while current Nordic market settings specify 2.9% for Skagerrak, 3.1% for NordLink, and 3.2% for NorNed (Nord Pool, n.d, Nordic Regional Coordination Centre, 2024). These values also align with the European Commission’s engineering estimates of 2.4% for line-commutated HVDC and 3% for voltage-source HVDC systems (Table 4) (European Commission, 2023). Applying the 3% loss to every cross-country transfer reduces the manufacturing TFP gain from 13.38 to 11.52 basis points, the gain within the four electricity-intensive divisions from 35.15 to 30.09 basis points, and the economy-wide gain from approximately 5.95 to 4.77 basis points. These correspond to attenuations of 14%, 14%, and 20%, respectively. The loss-adjusted union sends 675 TWh across

modelled borders and dissipates 20.3 TWh in transit.

This is an operating-loss robustness exercise, not a network-investment counterfactual. Additional cable capacity can relax a binding quantity constraint, but it does not change the loss incurred per MWh sent; in the full-union experiment the transmission bounds are already slack. A lower-loss route or technology would therefore require an explicit investment and network-design extension. The uniform factor is deliberately simple: actual losses vary with distance, technology, and loading. Accounting for them attenuates rather than overturns the productivity gain from integration.

G.6 Firm mass and entry

Total sector mass is exogenous, so the counterfactuals reallocate production within Europe without changing how much of it there is. Appendix E sets out the free-entry extension and the reasons the baseline abstracts from it; a reduced-form alternative, in which mass responds to sector profitability with an external entry elasticity, would capture the same world-market margin without a selection threshold.

For the electricity union that margin is not a gain, and its sign is not obvious. A union equalises prices without lowering them appreciably, and electricity-intensive production is already concentrated in the countries it makes more expensive, so the output-weighted electricity cost of European exposed production can rise even as aggregate productivity does. The competitiveness margin belongs instead to the build-out scenarios of Section 6.3, where the European price level genuinely falls.

G.7 Exposure through intermediates

Electricity cost shares measure direct purchases. Total exposure, tracing electricity embodied in intermediate inputs through the input–output structure, is about twice the direct share when only domestically generated electricity is counted and about 2.7 times including imported content, with the largest multiples in the least electricity-intensive sectors.

This is not a free multiplier on the results. The location multiplier is estimated per unit of directly measured exposure, so raising exposure and re-estimating the elasticity on the same basis leaves their product, and hence the reallocation any given price change induces, close to unchanged. Two further features work against a simple scaling: the embodied measure compresses the cross-sector exposure gradient that identifies the response, and the fuel-saving channel does not scale with exposure at all. The direct measure is therefore retained, and the restriction is read as understating exposure in the sectors furthest downstream rather than as a lever on the headline.

G.8 The labour market

Three assumptions govern how freely labour reallocates when firms move: how elastically labour is supplied to the modelled sectors, how far workers move across countries, and how freely they move across sectors within one. The baseline is at the restrictive extreme on the first two, since $\bar{L}_j$ is fixed and workers never cross a border, and at the permissive extreme on the third, since a single national wage w_j clears the sum over all thirteen sectors and labour therefore moves between chemicals and retail at a common wage. Each is relaxed below in turn and then all three together. Every variant nests the baseline exactly and reproduces the calibrated allocation at its own parameter values, so only the counterfactual response differs.

Elastic labour supply. The fixed endowment is replaced by an isoelastic supply curve anchored at the observed baseline, $L_j^S = \bar{L}_j(w_j/w_{j0})^\nu$, which recovers the baseline at $\nu = 0$ and leaves the rest of the model untouched. Because $\bar{L}_j$ counts employment in the modelled sectors only, ν is the elasticity of supply to that subset, and the relevant external estimate is an aggregate hours elasticity comparing stationary states rather than a Frisch elasticity: Chetty et al. (2011) report a steady-state Hicksian aggregate of 0.50 from macro evidence and 0.59 from micro, and Suárez Serrato and Zidar (2023) an effective elasticity of labour supply to a local labour market of 0.88, net of housing congestion.¹³¹ Solving the union at $\nu = 0.5$ raises the aggregate gain from 5.94 to 6.42 basis points and manufacturing from 13.4 to 13.9 bp; at $\nu = 0.88$ the figures are 6.73 and 14.3 bp.¹³² The mechanism is not the obvious one: gross cross-country reallocation *falls* slightly, while its covariance with the productivity fundamentals rises by about a fifth, because a high- $\bar{\Phi}_{js}$ country can now absorb arrivals without expelling as much incumbent production. Over half the aggregate improvement is the labour-intensive service sectors, which lose productivity under the union through the shared national wage, doing so less.

Cross-country mobility. Mobility replaces the fixed endowment with a logit allocation of a fixed European workforce, $L_j \propto \bar{L}_j \exp(\nu_m \ln w_j)$, so that $\sum_j L_j$ is conserved exactly and only its distribution responds. The relevant parameter is not the observed mobility stock, which the calibration matches by construction, but the responsiveness of flows to wage differences. Arpaia et al. (2016) decompose an asymmetric labour demand shock across EU countries and find that in the long run the increase in labour supply splits 40% activity rate and 60% mobility, which implies $\nu_m = 1.5 \nu$

¹³¹Both are real-wage objects while the supply curve above is written in nominal wages. The deflator is common across countries, so it cancels from the location comparison and leaves every reallocation channel unaffected; it bears only on the aggregate employment response, by at most ν times the change in the consumer price index, which the union puts below 7 basis points.

¹³²With $\nu > 0$ aggregate employment in the modelled sectors is no longer constant: the union lowers it by about 8 bp at $\nu = 0.5$ and 13 bp at the joint parameters below, because the employment-weighted average wage falls. The sectoral figures are unaffected, being output-weighted averages of productivity across countries within a sector, but the economy-wide productivity gain and the change in output now differ by the labour-share-weighted amount, 2 to 3 bp.

and hence $v_m = 0.75$ at the central v ; Beyer and Smets (2015) corroborate the order of magnitude and confirm that the familiar result that European and American mobility are comparable applies within countries, not across them. Adding this margin to the participation margin raises the aggregate gain to 7.01 bp and manufacturing to 14.6 bp. Mobility moves more labour across countries than participation does, but contributes almost nothing to the aggregate quantity, which is what conservation requires; it is also the cleaner of the two margins for the productivity accounting, since the workers it moves come from another modelled country whose contraction the model records, whereas the participation margin draws them from outside the model.

Sector-specific wages. The single national wage is replaced by a within-country Roy nest in which workers draw an idiosyncratic sector match with Fréchet shape κ , so that $L_{js}^S = \bar{L}_{js} \exp(\kappa d_{js}) / \Lambda_j$ with d_{js} the log deviation of the sector wage from its baseline. The country endowment is conserved for every κ , the baseline allocation is reproduced without recourse to sectoral wage data, and $\kappa \rightarrow \infty$ returns the common wage of the baseline. Galle, Rodríguez-Clare and Yi (2023) estimate this parameter for a model that is otherwise Caliendo and Parro (2015) with sector-specific wages, obtaining implied values between 1.42 and 2.79 across instrument constructions with a preferred $\kappa = 1.5$. At that value the aggregate gain falls from 5.94 to 4.46 bp and manufacturing from 13.4 to 7.5 bp, and at the knife edge $\kappa = 0$, where sectoral labour markets are entirely separate, to 4.01 and 5.7 bp. The reason is that the baseline lets an expanding electricity-intensive division draw on the entire national labour pool, so that its own expansion barely moves the wage it pays; under a finite κ it must bid within its own segment instead. Unlike the first two margins, here the reallocation and the productivity gain fall together and roughly proportionally, since sectoral segmentation damps a sector's response evenly across countries rather than changing which countries can absorb production.

All three together. Because the three act through the same market-clearing condition they interact, and multiplying the separate effects is not legitimate. Solved jointly, with the sector nest inside a country choice that responds to its inclusive value, at $v = 0.5$, $v_m = 0.75$ and $\kappa = 1.5$, the aggregate gain is 5.76 bp and manufacturing 9.2 bp. Both exceed what multiplying the separate results would predict, by 9% and 13% respectively, because elastic aggregate supply relieves part of the constraint that sectoral segmentation imposes. Taking all three at external estimates therefore lowers the manufacturing gain by about a third relative to the baseline, with the sectoral margin responsible for the whole of the reduction and the other two offsetting part of it.

What this leaves open. Two caveats bear on how much weight the joint figure carries. First, κ is estimated on American worker groups and applied here to thirteen broad European sectors compared across stationary states, a horizon over which sectoral mobility is plausibly higher than the preferred value, which is itself near the bottom of the estimated range; the damping reported

above is therefore the stronger end of what the source supports. Second, sector-specific wages bear on identification as well as on the counterfactual. The reduced-form estimate of Section 4 recovers the location multiplier because location–time fixed effects absorb the wage response, and that is exact only when one wage prevails in a location. With sector-specific wages the omitted term is a sector’s own wage response, which rises precisely where the sector expands, so the estimate is attenuated and the true multiplier larger than 15.8. Re-estimating on that basis would return part of what the sectoral margin removes, which makes the joint figure a conservative reading of the three margins rather than a central one.

G.9 Non-firm electricity demand

Non-firm electricity demand is held fixed at roughly half of load. Making it price-responsive at the external household elasticity would dampen the price response to firm entry, weaken electricity congestion and pass more reallocation through. The effect on production is small, and most of what elastic residual demand adds under a union is consumer surplus, a different object from the productivity gain reported here.

G.10 Wholesale and retail prices

Firms do not buy electricity at the wholesale price. They pay it plus network and distribution tariffs, supply margins and, depending on the tax concept, levies, components that differ across countries for reasons a wholesale union does not touch. To price that gap I re-solve the model with an additive retail wedge: a firm in country j and sector s pays $p_j^E + t_{js}$, where t_{js} is the difference between the ex-tax industrial retail price for the consumption band the cell occupies and the country’s wholesale price, averaged over the baseline window. Each country–sector cell takes the band its average firm’s annual purchase falls in, so a firm that relocates pays the destination country’s tariffs at the band its sector occupies there. The wedge is a fixed cost shifter: firms move across it and it does not respond to them. Its revenue is rebated lump sum, and both the location shifter and the firm demand coefficient are re-inverted at $p_j^E + t_{js}$, so the baseline reproduces the observed allocation as closely as the wholesale calibration does.

The measured wedge falls monotonically in consumption, from 101.0 EUR/MWh in the cells whose average firm buys under 20 MWh a year, to 24.8 EUR/MWh in the 2–20 GWh band where the electricity-intensive divisions sit, to 8.7 EUR/MWh in the largest band reached. That gradient is the volume discount the band structure represents, recovered rather than imposed. Weighted by electricity use it implies a pass-through of wholesale into firm-facing prices of 0.62 for the four electricity-intensive divisions, 0.53 for manufacturing as a whole and 0.49 economy-wide: the sectors that carry the location channels are also the ones whose bills are mostly wholesale.

The same correction applies to the location multiplier. The extensive-margin moment is estimated on bidding-zone wholesale prices, so if Norwegian firms paid spot plus a level network

charge, what the regression measured is the response to the firm-facing price scaled by its pass-through. Over the 2021–2023 divergence the Norwegian network charge in the relevant band is 14.0 EUR/MWh against zone prices from 30 to 122 EUR/MWh, giving a pass-through of 0.83 and $\psi - \gamma = 0.0530$ against 0.0639 in the published calibration.¹³³ This runs opposite to the pass-through attenuation of the counterfactual, so it is applied throughout rather than omitted as conservative. The correction is small against the uncertainty in the moment it adjusts, moving $\hat{\delta}_M$ by about one standard error.

Table G.1 re-runs the union against this calibration and splits integration into a wholesale and a retail leg. Equalising wholesale prices while the wedges remain is worth 8.5 basis points of manufacturing productivity against the 13.4 bp of the headline, and 23.8 bp within the electricity-intensive divisions against 35.1 bp, a scaling of about 0.63 and 0.68 respectively. The fossil-resource term falls as well, from 1.9 to 1.6 bp economy-wide, because firm electricity demand now responds to the wholesale price only through the pass-through and there is correspondingly less re-dispatch. Harmonising the wedges instead — within consumption band, at their demand-weighted mean, so that wedge revenue is preserved exactly — is worth 12.2 bp in the exposed divisions but –1.9 bp for manufacturing as a whole.

The two legs are complements at every level of aggregation, and in the exposed divisions they are worth more together than the wholesale union alone is credited with: 38.6 bp against the headline’s 35.1, exceeding the sum of the separate legs by 2.7 bp. Removing wholesale dispersion raises the return to removing the wedges, because a firm that has stopped paying for one source of cost dispersion is more responsive to the other. What the headline therefore understates is the value of integrating the price firms actually face; what it overstates is how much of that a wholesale union delivers on its own, which is about two-thirds in the exposed divisions and less elsewhere. Harmonisation moves production from the low-wedge countries — Sweden, Finland, the Netherlands, Denmark, Estonia, Slovenia — toward the high-wedge ones — Czechia, Spain, Belgium, Slovakia, Latvia, Bulgaria — with a productivity consequence that differs in sign across sectors, positive in the exposed divisions and negative in the manufacturing residual.

Three limitations attach. The band is assigned from the average firm in a cell, which is the object the model’s symmetric cells define but sits below the band where most volume is bought, since consumption within a cell is skewed. About one cell in six lies within 20% in logs of a band boundary, so its assignment, and hence the mean it harmonises to, is not sharp. And the wedge excludes taxes and levies, which is the concept the cost shares are measured on: including them would raise the demand-weighted wedge from 49.1 to 88.0 EUR/MWh, and from 27.1 to 54.7 EUR/MWh in the exposed divisions, so a firm holding no exemptions faces roughly twice

¹³³The pass-through varies from 0.68 in NO4 to 0.88 in NO2; the figure used weights each zone-year by its squared deviation of the log price, which is how the regression itself averages them. The unweighted mean of 0.81 and the value at the mean price of 0.86 bracket it. Norway’s retail price is published only as a national average, so the wedge here is the network component rather than the difference between retail and spot, which is negative in the southern zones.

the wedge used here.

Table G.1: Integration split into a wholesale and a retail leg

	Wholesale	Retail-wedge calibration		
	union (published)	Wholesale union	Wedge harmonised	Both
<i>Panel A: Economy-wide</i>				
Composition	2.535	1.575	0.183	1.814
Agglomeration	1.496	1.049	0.353	1.491
Allocative efficiency	0.021	0.027	0.017	0.040
Fossil-resource saving	1.887	1.597	−0.057	1.571
Total	5.939	4.248	0.496	4.915
<i>Panel B: Manufacturing</i>				
Composition	8.637	5.206	−2.098	3.279
Agglomeration	4.691	3.214	0.141	3.558
Allocative efficiency	0.049	0.063	0.014	0.072
Total	13.378	8.483	−1.942	6.910
<i>Panel C: Electricity-intensive divisions</i>				
Composition	22.361	14.671	10.266	26.159
Agglomeration	12.571	8.841	1.898	12.192
Allocative efficiency	0.215	0.243	0.026	0.250
Total	35.147	23.754	12.190	38.600

Note: Changes in production-side TFP in basis points relative to the relevant baseline. The first column is the published full electricity union of Section 6, computed on the wholesale calibration. The remaining three re-solve the model with the additive retail wedge t_{js} : equalising wholesale prices with the wedges retained, harmonising the wedges within consumption band at their demand-weighted mean with wholesale prices at baseline, and both together. The retail columns apply the Norwegian pass-through correction to the location multiplier. Composition, agglomeration, and allocative efficiency are the three location channels of Lemma 4, with output weights renormalised within each group. The fossil-resource saving is an economy-wide level term with no sector index and is therefore reported only in Panel A; the group totals are location channels only. Electricity-intensive divisions are paper, chemicals, non-metallic minerals, and basic metals.

G.11 Transmission investment and the returns to integration

Supporting the electricity union requires additional transmission capacity. I estimate the least upfront investment that supports the union allocation on the model's existing country pairs, and compare its cost with the monetary equivalents of the productivity and fuel-saving gains. The central estimate is EUR 120.7 billion for the baseline economy and EUR 106.3 billion beyond the planned grid for the 2040 economy. Under the valuation and discounting assumptions below, aggregate productivity gains and fuel savings exceed transmission costs in both comparisons.

Capacity and routing. The unconstrained union determines countries' net electricity injections, but these injections need not determine a unique set of bilateral flows. The solver's small quadratic flow regulariser selects a routing without regard to construction costs. I therefore hold the union injections fixed and choose flows and additional capacity jointly to minimise investment. Let $\mathcal{E}$ be the set of 36 country pairs, f_e signed annual flow on pair e , $x_e \geq 0$ additional bidirectional capacity in MW, and c_e its cost per MW. The problem is

$$\min_{\{f_e, x_e\}} \sum_{e \in \mathcal{E}} c_e x_e \quad \text{subject to} \quad Nf = b^U, \quad -h(\bar{k}_e^- + x_e) \leq f_e \leq h(\bar{k}_e^+ + x_e), \quad x_e \geq 0.$$

Here N is the country–edge incidence matrix, with positive entries at the exporting end; b^U is the vector of union net injections; $\bar{k}_e^+$ and $\bar{k}_e^-$ are reference capacities in the two directions; and $h = 0.00876$ converts MW into annual TWh. One additional MW raises both directional bounds and is charged once. The objective minimises upfront capital expenditure; annual operation and maintenance costs enter the subsequent return calculation. The exercise retains the annual, lossless country-level transport model and does not infer hourly peak requirements or electrical power-flow feasibility.

Investment costs. Costs come from ENTSO-E's TYNDP 2024 *Investment Candidates and Cost Assumptions* workbook (European Network of Transmission System Operators for Electricity, 2024b). I map its conceptual net-transfer-capacity candidates to model borders and calculate each border's central EUR/MW factor as total candidate capital expenditure divided by total incremental capacity in one direction. This uses 109 conceptual entries covering 35 of the 36 model pairs. Annual maintenance factors use the corresponding reported operating costs. The smallest and largest candidate unit costs define alternative cost scenarios; they are not confidence intervals. The workbook's net-transfer-capacity and bidding-zone representations are cross-checked rather than pooled, and real projects are not added to the conceptual candidates. The accompanying methodology confirms that candidate packages include the internal reinforcements needed to support the cross-border increase (European Network of Transmission System Operators for Electricity, 2025, p. 83).¹³⁴

Lithuania–Poland uses the adopted 2024 budgets for the overland Harmony Link project, giving EUR 1.32 million per MW for 700 MW of commercial transfer capacity (Polskie Sieci Elektroenergetyczne, 2024, LITGRID AB, 2024).¹³⁵ The published packages include necessary domes-

¹³⁴One anomalous Hungary–Slovakia maintenance entry is excluded from the maintenance average, with its original value preserved; its capital cost remains included.

¹³⁵The proxy combines PSE's approximately PLN 3 billion budget and Litgrid's EUR 220.03 million budget, converting at the [National Bank of Poland rate](#) of PLN 4.2621 per euro on 23 December 2024. [Litgrid's September 2024 project description](#) supplies the 700 MW commercial capacity; symmetry is an assumption. Maintenance is set to 1.45% of capital cost annually, the midpoint of the 0.7–2.2% band for non-mature projects (European Network of Transmission System Operators for Electricity, 2024c, p. 88). This proxy is held fixed across cost scenarios.

tic reinforcements without identifying their costs separately. I retain those packages, with neither an arbitrary subtraction nor an additional domestic-grid surcharge. Border-specific factors capture the differences embedded in these source packages; I do not impose separate mountain or submarine multipliers.

Reference economies and verification. Baseline additions are measured from the effective capacities used to solve the calibrated economy, including its feasibility adjustments to historical net transfer capacities. Using unadjusted historical capacities instead raises central capital expenditure to EUR 121.3 billion. The 2040 comparison holds projected electricity intensities, generation availability, residual demand and the carbon-price assumption common to the planned-grid and union equilibria. It costs only capacity beyond the planned TYNDP 2040 grid and measures benefits against that same grid; planned transmission and common generation investment are not charged again. I re-solve the equilibrium with the proposed finite capacities and check prices, firm allocations and productivity against the union. The central networks reported here pass these checks. In 2040, additions are rounded upward to whole MW before verification; the resulting EUR 5.5 million increment is included in costs.¹³⁶

Valuation and discounting. I value composition, agglomeration and allocative-efficiency gains separately from fuel savings. If g_s denotes the sector's location-channel gain in log points, μ_s its aggregate expenditure weight, and X_0 the model's reference gross output, annual monetary equivalents are $X_0 \sum_s \mu_s g_s$ for all model sectors and $X_0 \sum_{s \in \text{manufacturing}} \mu_s g_s$ for manufacturing. These are alternative scopes: the aggregate already includes manufacturing. Fuel savings are the difference in integrated fossil fuel bills between the reference and union equilibria. Electricity-bill changes, congestion rents and carbon-allowance transfers are not additional resource gains. The productivity equivalents use the model's EUR 16.35 trillion EU24 gross-output normalisation for its 23-country allocation; they are first-order output equivalents, not GDP or consumption-equivalent welfare.

The candidate workbook does not establish a uniform cost price year. I assume costs are 2024-equivalent and multiply the model's nominal 2015–2019 benefits by 1.24336, the Eurostat EU27 GDP deflator for 2024 relative to its 2015–2019 mean (Eurostat, 2026b).¹³⁷ The 2040 calculation retains this output valuation scale; it does not extrapolate GDP growth to 2040.

¹³⁶The finite-network checks retain the solver's convergence gates and require maximum absolute log differences below 10^{-4} for prices and firm mass and a manufacturing productivity difference below 0.001 basis points. A separate 2040 calculation starting from current effective capacities gives EUR 209.5 billion of capital expenditure, but its finite-network re-solve misses the convergence gate. That supplementary implementation remains provisional and is excluded from Table G.2. Low/high cost scenarios pass the linear-programme checks but are not separately certified by finite-equilibrium re-solves.

¹³⁷The calculation uses dataset `nama_10_gdp`, item `B1GQ`, unit `PD15_EUR`, geography `EU27_2020`: $127.255/102.3476 = 1.24336086$. This broad purchasing-power adjustment does not establish a common cost vintage or exactly deflate the model's nominal averages.

Following the TYNDP assessment convention, I use a 4% real discount rate, a 25-year horizon, investment at commissioning and zero terminal residual value (European Network of Transmission System Operators for Electricity, 2024a, pp. 27–28).¹³⁸ I hold each economy’s annual real benefits and maintenance costs constant over the horizon. For either productivity scope, the net present value is

$$\text{NPV} = -C + \sum_{t=1}^{25} \frac{B^{\text{location}} + B^{\text{fuel}} - O}{(1.04)^t},$$

where C is upfront investment, O annual maintenance, and B^{location} and B^{fuel} the annual benefits in the common assumed valuation. Capital expenditure is undiscounted because it occurs at $t = 0$. Present values are at commissioning; for 2040, they value flows in 2041–2065 in assumed 2024 purchasing-power units, rather than discounting them back to calendar 2024.

Results and interpretation. Table G.2 separates gross gains from costs. The baseline requires 84.3 GW on 13 borders. Cost-aware routing saves EUR 57.9 billion relative to retaining the original union flows. Aggregate location gains have a present value of EUR 128.7 billion, just below total transmission costs of EUR 130.6 billion. Fuel savings add EUR 59.9 billion, taking net present value to EUR 58.0 billion. Restricting productivity gains to manufacturing gives a larger EUR 99.2 billion net present value because it omits the negative nonmanufacturing contribution.

For 2040, 72.0 GW on eight borders supports the union beyond the planned grid. Aggregate location gains have a present value of EUR 443.3 billion against EUR 115.6 billion of total transmission costs, giving net present value of EUR 327.7 billion. Manufacturing gives EUR 332.3 billion. Fuel savings are negligible because fossil dispatch is already essentially zero in the reference economy. Productivity changes are computed directly between the 2040 planned-grid and union equilibria.

Extrapolation is the main cost uncertainty: borders exceeding source candidate volumes account for EUR 115.2 billion of baseline investment and EUR 101.8 billion in 2040. The low/high unit-cost scenarios give capital expenditure of EUR 105.2–135.4 billion and EUR 82.0–122.5 billion, respectively, but do not capture the uncertainty from scaling beyond those volumes. These are annual-model investment equivalents, conditional on linear costs and the monetary conversion, rather than engineering budgets or a proof that full integration maximises welfare.

H Counterfactual design

This appendix documents the additional primitives used in the counterfactuals of Section 6. All exercises use the calibrated equilibrium and network-clearing procedure described in Appendix F.

¹³⁸The *System Needs Methodology* also uses 4% and 25 years for interconnector annualisation (European Network of Transmission System Operators for Electricity, 2025, pp. 47–48). The guideline discounts to the study year; I report values at commissioning.

Table G.2: Transmission investment and gains from the electricity union

Reference network	Baseline economy Effective baseline	2040 economy Planned 2040
<i>Capacity and annual amounts</i>		
Additional capacity (GW)	84.3	72.0
Expanded country pairs	13	8
Annual maintenance (EUR bn)	0.64	0.60
Annual aggregate location gains (EUR bn)	8.24	28.38
Annual manufacturing location gains (EUR bn)	10.87	28.68
Annual fuel savings (EUR bn)	3.84	0.00
<i>Present values at commissioning (EUR bn)</i>		
Aggregate location gains	128.7	443.3
Manufacturing location gains	169.9	448.0
Fuel savings	59.9	0.0
Total gains: aggregate location plus fuel	188.6	443.3
Total gains: manufacturing location plus fuel	229.8	448.0
Upfront investment	120.7	106.3
Discounted maintenance	9.9	9.3
Total costs	130.6	115.6
NPV: aggregate location plus fuel	58.0	327.7
NPV: manufacturing location plus fuel	99.2	332.3

Notes: Central border cost factors; 4% real discounting over 25 years; constant annual real benefits and maintenance; upfront investment at commissioning and no terminal residual. All monetary entries use the conditional 2024 valuation described in the text. Productivity gains combine composition, agglomeration and allocative efficiency. Manufacturing is included in the aggregate and is reported as an alternative scope. Each NPV subtracts the same full transmission cost once. The 2040 column costs additions beyond the planned grid and includes whole-MW rounding. Generation investment is common to the reference and union within each comparison. Fuel savings in 2040 are numerically zero. Values are independently rounded.

The full electricity union makes every internal transmission constraint slack and clears a common price against the pooled baseline generation fleet, so it introduces no additional data input. The natural-potential comparison reallocates the existing EU-wide renewable stocks using JRC ENSPRESO technical capacity for solar PV and on- and offshore wind, and gross theoretical hydropower potential from Hoes et al. (2017); combining this allocation with the union defines the comprehensive-coordination benchmark without further primitives. The announced-trajectory exercises instead use the 2040 national capacity, generation, demand, and grid projections from ENTSO-E’s TYNDP 2024 scenario family. The subsections below document these constructions, the projected sectoral electricity cost shares used in the electrification exercises, and a companion exercise that closes the existing network rather than extending it.

H.1 Natural-potential inputs

The natural-potential comparison re-sites the four renewable stocks using two potential measures. For solar PV and for on- and offshore wind, the potential $\bar{\theta}_j^r$ is technical capacity from the JRC ENSPRESO database (Ruiz et al., 2019); installed capacity is reallocated in proportion to $\bar{\theta}_j^r$, preserving the EU-wide total within each technology, and generation is computed at destination-country yields. Offshore potential is defined for the coastal model countries; landlocked countries hold no offshore stock. For hydro, the potential is gross theoretical potential from Hoes et al. (2017) (Appendix C), and the reallocation is on a generation basis, since installed hydro capacity includes pumped storage and capacity-based yields are not comparable across countries. Nuclear and fossil capacity remain at their baseline locations, and fossil utilisation clears residual demand through the baseline merit-order system. The reallocated stocks are small relative to the ceilings — a uniform one to three percent of technical potential per technology — so the proportional rule never binds against a ceiling.

H.2 Announced renewable trajectories: projected inputs

The announced-trajectory scenarios combine four projected inputs at the 2040 horizon, harmonised to the EU-23 sample.

Announced capacity. The core input is the ENTSO-E TYNDP 2024 national-trajectory dataset at the 2040 horizon, using the maximum-capacity corridor (the minimum corridor is retained as a sensitivity). The published dataset covers solar PV (rooftop and utility) and onshore wind; across the EU-23 the announced build-out raises solar capacity from 103 to 1,433 GW and onshore wind from 128 to 619 GW. Nuclear capacity follows the TYNDP National Trends scenario, including the German exit and the Polish entry, with new-entrant capacity converted at the EU fleet-average yield (111 to 98 GW). Where announced onshore ambition exceeds the ENSPRESO technical ceiling — Austria, Finland, and Germany, by 167.6 TWh in total — capacity is clipped at the ceiling.

Offshore wind and hydro. Offshore wind is not part of the published trajectory rows and is staged separately from TYNDP National Trends capacity (257.7 GW), converted to generation with per-country capacity factors; it contributes 954 TWh, the largest single addition. Hydro is not an announced policy path: 2040 hydro generation is the mean inflow across 36 climate years (346.5 TWh EU-wide), with countries lacking projections held at baseline.

Demand, network, and fossil capacity. Electricity demand is scaled to the TYNDP 2024 projections, applying country growth factors between 1.02 and 1.75 — the midpoint of the Distributed Energy and Global Ambition scenarios — to the non-firm residual load, while firm electricity demand remains endogenous. The scenarios run on the TYNDP-2040 grid of the planned-integration exercise. Announced fossil phase-out paths are not imposed: fossil capacity remains at baseline and utilisation adjusts endogenously; in equilibrium fossil dispatch falls to essentially zero, so the omission does not bind. In aggregate, projected zero- and low-marginal-cost generation reaches 4,729 TWh against a scaled load of 3,277 TWh: announced capacity oversupplies even at 2040 demand.

Re-sited variant. The re-sited scenario reallocates the projected solar and onshore-wind totals across countries with the technical-potential rule of the natural-potential comparison, holding EU-wide totals fixed within each technology; nuclear follows the announced path.

H.3 Projected electricity cost shares

Construction. The electrification exercise of Section 6.3 replaces the baseline electricity cost shares with projected 2040 values, $\beta_s^{2040} = \beta_s \times r_s$, where β_s is the 2010–2011 Exiobase anchor of Section 5.2 and r_s is the EU-aggregate growth of sectoral electricity use relative to sectoral activity between 2019 and 2040 in the TYNDP 2024 Distributed Energy scenario — the same scenario family that supplies the generation, demand, and grid projections above. The ratio is EU-aggregate and country-invariant by construction: TYNDP’s country detail embeds the scenario’s own assumptions about where industry sits, and importing that pattern would answer the model’s location question by assumption. Relative prices are held constant, so r_s moves technology rather than prices, and the labour shares α_s remain at their calibration values. TYNDP’s base year 2019 closes the calibration window, so the 2019–2040 ratio chains onto the baseline shares without double-counting. A separately measured 2024–25 cost share is used only as a corroborating check rather than as the anchor, because its drift from 2010–11 is dominated by the rise in the relative electricity price while physical intensity fell, whereas the projection requires β_s as a technology object.

Sector mapping and activity measures. TYNDP reports industrial electricity demand and activity by subsector, mapped to the model sectors as follows: paper (C17) from Paper; chemicals (C20) from Chemicals and Fertilizers; basic metals (C24) from Steel, Aluminium, and Other metals; residual manufacturing from Food, Refineries, and Aggregated other; and information and communication (J) from Datacentres, which TYNDP files under industry but which belong to J in the model’s accounting (see the level check below). Non-metallic minerals (C23) has no TYNDP subsector of its own and takes the manufacturing-aggregate ratio, excluding datacentres so that the proxy remains a manufacturing object. The activity denominator is physical production for the three metals subsectors and useful energy demand elsewhere. Because useful energy demand embeds process efficiency, a falling activity index would overstate the implied cost-share rise, so activity ratios are floored at one; the floor binds for non-metallic minerals (raw ratio 0.97) and residual manufacturing (0.86). The datacentre assignment rests on a level check against the model’s baseline electricity use: the TYNDP residual-manufacturing basket matches the model’s residual-manufacturing electricity to about 3% once datacentres are excluded and overshoots by about 17% when they are included, while TYNDP’s datacentre demand alone (75 TWh in 2019) is the size of the model’s entire J sector (64 TWh).

Held sectors. Mining (B) and construction (F) do not appear in TYNDP’s industry accounts, transport (H) is excluded because its projected electricity growth is dominated by vehicle charging rather than the modelled sector’s production costs, and the market services (G, I, M, N) carry no TYNDP output denominator; all are held at their baseline shares, which understates economy-wide electrification. The resulting factors are collected in Table H.1, and the cost shares they imply are shown in levels in Figure H.1: electrification is concentrated in chemicals, non-metallic minerals, residual manufacturing, and paper, is nearly flat in basic metals, and information and communication declines as datacentre electricity grows more slowly than the underlying activity.

H.4 A budget-neutral Italy-focused interconnection portfolio


This subsection asks whether the expenditure assigned to the TYNDP 2040 interconnection plan could generate larger productivity gains if allocated differently. I hold fixed the projected 2040 economy described above: sectoral electricity cost shares, generation availability, residual demand, the carbon price, and all non-network primitives remain unchanged. The exercise changes only cross-border transfer capacities. Its objective is within-manufacturing TFP, subject to preserving the investment, electricity-price, and aggregate-productivity outcomes of the TYNDP portfolio.

Policy problem. Let k collect the two directed transfer capacities on each of the 36 country pairs in the model, let k^0 denote the effective current grid, and let k^T denote the TYNDP 2040 grid. For

Table H.1: Projected 2040 electricity-intensity factors, TYNDP 2024 Distributed Energy

Sector	TYNDP subsectors	Electricity ratio	Activity ratio	r_s
C17 Paper	Paper	1.47	1.16	1.27
C20 Chemicals	Chemicals, Fertilizers	1.65	1.02	1.62
C23 Minerals	Manufacturing aggregate (proxy)	1.38	1.00 [†]	1.38
C24 Basic metals	Steel, Aluminium, Other metals	1.23	1.18	1.04
C-rest	Food, Refineries, Aggregated other	1.32	1.00 [†]	1.32
J Information & comm.	Datacentres	2.43	3.35	0.73
B, F, G, H, I, M, N	held at baseline	—	—	1.00

Notes: EU-aggregate ratios 2040/2019 from the ENTSO-E TYNDP 2024 Distributed Energy demand scenario. The electricity ratio is sectoral electricity demand growth; the activity ratio is physical production for the metals subsectors and useful energy demand otherwise; r_s is their quotient and scales the baseline cost share, $\beta_s^{2040} = \beta_s \times r_s$. [†]Activity ratio floored at one (raw values 0.97 for C23, 0.86 for C-rest), so that efficiency-driven declines in useful energy demand do not inflate the implied cost-share rise.

Figure H.1: Electricity cost shares by model sector: 2010–2011 baseline and projected 2040

Note: Electricity cost share of gross output for the thirteen model sectors. The lower segment is the 2010–2011 Exiobase anchor β_s of Section 5.2; the upper segment is the projected increment to $\beta_s^{2040} = \beta_s \times r_s$, with the factors r_s of Table H.1 from the ENTSO-E TYNDP 2024 Distributed Energy scenario. Electrification widens the cross-sector gradient rather than shifting it: the four exposed divisions start clustered between 3.0 and 3.9% and separate from the rest, while the seven sectors without a TYNDP activity denominator are held at baseline and do not move. The ordering within the exposed group also changes. Chemicals rises from 3.0 to 4.9% and paper from 3.9 to 4.9%, while basic metals — the most electricity-intensive division at baseline — is close to flat at 3.9 to 4.0%, since its projected electricity growth is nearly matched by its projected activity growth. Information and communication is the only sector to fall, from 0.7 to 0.5%, because datacentre electricity is projected to grow more slowly than the activity it serves.

border e and direction $d \in \{+, -\}$, define the positive addition relative to the current grid as

$$a_{ed}(\mathbf{k}) = \max\{0, k_{ed} - k_{ed}^0\}.$$

Following the investment-cost convention in Appendix G.11, one bidirectional expansion is charged once. If c_e is the central capital cost per MW on border e , total upfront expenditure is therefore

$$C(\mathbf{k}) = \sum_{e \in \mathcal{E}} c_e \max\{a_{e+}(\mathbf{k}), a_{e-}(\mathbf{k})\}.$$

The TYNDP portfolio defines the upfront CAPEX budget, $\bar{C} \equiv C(\mathbf{k}^T) = \text{EUR } 194.247$ billion. Operating and maintenance expenditure is excluded from this constraint. The exercise concerns the allocation of a fixed construction budget and therefore does not require assumptions about discounting or asset lifetimes.

For any candidate network $\mathbf{k}$, I solve the complete equilibrium in the projected 2040 economy. Let $G^M(\mathbf{k})$ denote the within-manufacturing TFP change relative to the current-grid 2040 equilibrium and let $G^A(\mathbf{k})$ denote the corresponding aggregate TFP change. Manufacturing comprises paper, chemicals, non-metallic minerals, basic metals, and remaining manufacturing, with expenditure weights renormalised within this group.

I also calculate a fixed-load weighted electricity price,

$$\bar{p}(\mathbf{k}) = \sum_j \lambda_j^0 p_j(\mathbf{k}), \quad \lambda_j^0 = \frac{L_j^0}{\sum_\ell L_\ell^0},$$

where L_j^0 is country j 's firm and residual electricity load in the current-grid 2040 reference equilibrium. Holding these weights fixed ensures that the price criterion does not mechanically change when the portfolio relocates electricity demand across countries.

The unrestricted policy problem would be

$$\begin{aligned} \max_{\mathbf{k}} \quad & G^M(\mathbf{k}) \\ \text{subject to} \quad & C(\mathbf{k}) \leq C(\mathbf{k}^T), \\ & \bar{p}(\mathbf{k}) \leq \bar{p}(\mathbf{k}^T), \\ & G^A(\mathbf{k}) \geq G^A(\mathbf{k}^T). \end{aligned} \tag{H.1}$$

The second constraint requires the candidate portfolio to weakly improve on the TYNDP weighted electricity price of EUR 28.504/MWh. The third prevents an increase in manufacturing TFP obtained at the expense of the rest of the economy: aggregate TFP must be at least the TYNDP gain of 4.058 basis points relative to the current-grid 2040 equilibrium.

Table H.2: Budget-neutral Italy-focused interconnection portfolios

Italy addition per border	Non-Italy TYNDP haircut	Δ manufacturing TFP (bp)	Δ aggregate TFP (bp)	$\Delta \bar{p}$ (EUR/MWh)
0.5 GW	2.39%	8.775	2.168	−0.114
1.0 GW	4.78%	17.309	4.264	−0.199
1.5 GW	7.17%	25.792	6.361	−0.238
2.0 GW	9.56%	30.198	7.519	−0.262
2.5 GW	11.96%	30.073	7.547	−0.288
3.0 GW	14.35%	29.722	7.551	−0.301
4.0 GW	19.13%	28.776	7.431	−0.293

Notes: Changes are relative to the TYNDP equilibrium in the projected 2040 economy. Every row has the same central upfront CAPEX, EUR 194.247 billion, and satisfies the price and aggregate-TFP constraints in equation (H.1). The Italy addition expands both directional bounds by the stated amount and is charged once as a bidirectional project. Manufacturing TFP uses weights renormalised across the five manufacturing sectors. Aggregate TFP includes manufacturing. The price is weighted by fixed current-grid 2040 loads.

A transparent restricted portfolio. Rather than claiming to solve equation (H.1) globally, I evaluate a one-dimensional family motivated by Italy’s position in the network. Let

$$\mathcal{E}_{IT} = \{AT - IT, FR - IT, GR - IT, IT - SI\}$$

be the four modeled Italian borders. For an Italy tranche t , measured in MW, the candidate network is

$$k_{ed}(t) = \begin{cases} k_{ed}^T + t, & e \in \mathcal{E}_{IT}, \\ k_{ed}^0 + \alpha(t)(k_{ed}^T - k_{ed}^0), & e \notin \mathcal{E}_{IT}, \end{cases} \quad d \in \{+, -\}. \quad (\text{H.2})$$

Thus, every Italian border receives the same bidirectional addition relative to TYNDP. All non-Italian TYNDP expansions receive one common scaling factor $\alpha(t) \in [0, 1]$. I choose $\alpha(t)$ by bisection until $C(k(t)) = \bar{C}$. This financing rule avoids selecting individual non-Italian projects after observing their model outcomes. It also leaves the portfolio governed by one policy parameter, t : a larger Italian tranche automatically implies a larger proportional reduction in the rest of the TYNDP expansion.

I evaluate tranches of 0.5, 1, 1.5, 2, 2.5, 3, and 4 GW on each Italian border. For every tranche, I reconstruct the complete 36-border capacity vector, match the TYNDP CAPEX budget to within 10^{-9} billion euros, and re-solve the general equilibrium. Every reported solution passes the model’s convergence, labour-market, network-balance, price-consistency, and transmission-optimality checks.

Results. The best manufacturing outcome among the evaluated tranches assigns an additional 2 GW of bidirectional capacity to each Italian border. Financing this expansion requires a 9.565% proportional reduction in every non-Italian TYNDP expansion, so $\alpha(2,000) = 0.90435$. The resulting network exhausts the same EUR 194.247 billion budget as TYNDP.

Relative to the current-grid 2040 equilibrium, the TYNDP network changes within-manufacturing TFP by -1.313 basis points. The Italy-focused portfolio instead raises it by 28.885 basis points. Its improvement relative to TYNDP is therefore 30.198 basis points. The gain does not come at the expense of the rest of the economy. Aggregate TFP rises by 11.577 basis points relative to the current grid, compared with 4.058 basis points under TYNDP, an improvement of 7.519 basis points. The fixed-load weighted electricity price also falls from EUR 28.504/MWh under TYNDP to EUR 28.242/MWh.

The dose response is nonlinear. Manufacturing gains rise steeply up to the 2 GW tranche, flatten between 2 and 2.5 GW, and then decline. Aggregate TFP continues to rise slightly until the 3 GW tranche. Under a fixed budget, additional Italian capacity eventually displaces increasingly valuable connections elsewhere. The preferred 2 GW case therefore maximises the stated manufacturing objective on the evaluated grid, while the aggregate objective would select a somewhat larger Italian tranche. The exercise brackets the manufacturing maximum between 1.5 and 2.5 GW but does not establish the continuous optimum.

For scale, the full electricity union raises within-manufacturing TFP by 35.275 basis points relative to TYNDP in the same projected 2040 economy. The 2 GW Italy-focused portfolio delivers 30.198 basis points, or 85.6% of this gain. This comparison does not imply that the Italy rule is the globally cost-minimising route to the union. It shows that a simple reallocation of the existing TYNDP budget toward Italy obtains most of the manufacturing productivity gain while improving both the aggregate-TFP and price outcomes delivered by the planned portfolio.

Scope. The exercise is a transparent restricted counterfactual rather than a global network optimum. It changes only the allocation of interconnection capacity and does not model construction sequencing, reliability benefits, or hourly power-flow constraints. The central cost schedule is linear beyond the published candidate-project volumes. This extrapolation affects both benchmarks: EUR 180.538 billion of the TYNDP cost and EUR 181.644 billion of the preferred Italy portfolio lie on borders whose priced additions exceed the observed project menu. The estimates should therefore be interpreted as comparable model investment costs rather than engineering project budgets.

H.5 The value of existing interconnection

The integration exercises all move toward a denser network; the opposite experiment prices the network that already exists. Closing every internal interconnector, holding generation capacity,

fuel and carbon prices, demand, and the fixed external positions at their calibrated values, leaves aggregate productivity 0.19 basis points *higher* than the baseline, so existing internal trade is worth about -0.2 bp on the production-side measure, against the 5.9 bp available from completing the network.¹³⁹ The near-zero aggregate is a cancellation rather than an absence of movement: the location channels gain 0.8 bp while the fossil-resource channel loses 0.7 bp, about €1.1 billion of additional fuel. Measured as within-group intensity, the location channels gain 4.1 bp in manufacturing and 7.8 bp in the four exposed divisions. In the 2040 economy the sign reverses and the magnitude becomes material: on the planned grid, aggregate productivity stands 2.8 bp above the same economy with no internal trade.

Two channels move against each other, and which dominates depends on the economy. Removing trade widens demand-weighted price dispersion by about a third while leaving the average price unchanged, so electricity-intensive production sorts harder on electricity cost. On today's generation mix that sorting is benign: the countries that cheapen when trade stops are the net exporters, and those are on average the more productive locations, so composition improves rather than deteriorating. What trade buys in the baseline economy is a cheaper dispatch of the existing fleet, not a better placement of production. By 2040 the build-out and the carbon price leave many countries separately cheap, and the binding loss from closing the network is clustering rather than sorting: electricity-intensive production spreads across the several cheap countries instead of concentrating in the most productive of them. Within manufacturing, trade is worth 6.6 basis points of agglomeration against -10.3 bp of composition; economy-wide, agglomeration accounts for essentially all of the 2.8 bp gain.

H.6 The German nuclear phase-out

A recurring question about European electricity policy is what Germany's nuclear phase-out cost. The model prices it by restoring Germany's 2010 pre-Fukushima fleet — 140.56 TWh/year of generation — to the German low-cost block and to nothing else: no other country, no other technology, and no change to demand, transmission, fuel or carbon prices.¹⁴⁰ In the calm calibration this adds 59.8 TWh/year to the fleet the 2015–2019 window already contains, and lowers the German wholesale price from 37.9 to 36.4 EUR/MWh. Only about half of the restored volume displaces German fossil generation; the rest is exported or absorbed by demand at the lower price.

Production reallocates toward Germany in precisely the sectors the model makes exposed. German firm mass rises by 0.46% in manufacturing and by 1.36% in the four exposed divisions,

¹³⁹The exercise removes internal trade only: the fixed residual positions on omitted external borders are unchanged, so this is not autarky. The annual energy model carries no reliability or adequacy margin, so the figure is a lower bound on the cost of losing the network.

¹⁴⁰The addition is measured from observed generation rather than installed capacity, since the capacity series bundles nuclear with other fuels not elsewhere classified. Germany remains on the fossil segment of its supply schedule throughout, so the merit-order inversion is defined and the cost assigned to nuclear sets no price.

while total German firm mass *falls* by 0.05%. Cheaper power bids up the German wage and draws mass out of the labour-intensive service sectors, the same margin that produces the services drag in the union exercise.

The aggregate productivity effect is nonetheless small. Restoring the fleet raises production-side TFP by 0.75 basis points, and by 0.20 bp once the fleet's own operating cost is charged at 15 EUR/MWh.¹⁴¹ Of the gross figure, 0.69 bp is the saving on imported fossil fuel and only 0.06 bp is the location channel. That location term is small because it is very nearly a transfer: Germany's own contribution to the European location channels is +1.58 bp against -1.52 bp for the other twenty-two countries. The same cancellation holds an order of magnitude larger on the scopes reported in the text. Measured as within-group intensity, Germany gains 5.81 bp in manufacturing against -5.45 bp elsewhere, and 15.14 bp in the four exposed divisions against -13.76 bp, for European nets of 0.36 and 1.38 bp.


Read in reverse, the phase-out moved electricity-intensive production out of Germany and lowered Germany's contribution to European productivity, but it moved that production to other European locations rather than destroying it. The European aggregate therefore barely registers the policy, and what it does register is a fuel bill rather than a misallocation. The exercise is the sharpest illustration of the distinction this paper is built on: a national generation policy is a lever on *where* production sits, while integration is the lever on what Europe produces in aggregate.

H.7 Counterfactual exhibits

This subsection collects the counterfactual exhibits referenced from Section 6.


¹⁴¹This accounting is gross of the capital, retrofit and decommissioning costs of the fleet, and the annual energy model prices none of nuclear's capacity value, so the figure speaks to wholesale prices and to location rather than to the welfare of the phase-out.

Figure H.2: Firm-mass reallocation under the full electricity union




Note: Change in firm mass under the full electricity union relative to baseline, in percent; blue marks an increase, red a decrease, white no change; non-model countries in light grey. Panel (a): total firm mass (largest changes: Czech Republic +0.47%, Italy −0.33%). Panel (b): firm mass in the four electricity-intensive manufacturing divisions — paper, chemicals, non-metallic minerals, basic metals (largest changes: Greece +9.2%, Denmark −12.0%).

Figure H.3: Transmission congestion under the baseline, TYNDP-2040, and full-union grids




Note: Edges connect country centroids; edge width and colour encode the per-border congestion rent — the solved shadow price of the transmission constraint times the border flow — on one shared scale across panels. Total rents are €816 million per year under the baseline grid, €1,054 million under the TYNDP-2040 grid, and zero under the full union.

Figure H.4: Price compression across integration scenarios




Note: Baseline wholesale electricity prices by country, sorted (bars); country prices under US-equivalent integration (diamonds) and the TYNDP-2040 grid (circles), drawn at the same country positions; and the full-union clearing price (dashed line). EUR/MWh, 2015–2019 baseline calibration. The cross-country span falls from 17.8 (baseline) to 11.3 (US-equivalent integration) to 7.8 EUR/MWh (TYNDP-2040), and to zero under the union.

Figure H.5: Natural-potential re-siting: generation moved and firm-mass response




Note: Natural-potential comparison (renewable capacity re-sited in proportion to measured potential, segmented markets) relative to baseline. Blue marks an increase, red a decrease, white no change; non-model countries in light grey. Panel (a): change in renewable generation, in TWh (largest recipients: France +45, Spain +26, Romania +19 TWh; largest sources: Germany -111, Sweden -33, Austria -19 TWh). Panel (b): change in total firm mass, in percent (largest changes: Czechia +0.26, Bulgaria -0.49). The generation moved is large — 380 TWh gross — while the firm-mass response is small, and aggregate TFP falls: on a segmented grid the re-sited cheap power cannot reach demand.

Figure H.6: Renewable build-out under announced national plans, 2040




Note: Change in renewable generation under the announced-trajectories scenario relative to its baseline, in TWh; positive everywhere, largest in Germany (+761), Spain (+357), and France (+354 TWh). Non-model countries in light grey.

Figure H.7: CO₂ emission levels: baseline and coordinated scenarios




Note: Direct combustion CO₂ from modelled fossil generation by country, in Mt per year, on one shared colour scale across panels (pale = low, dark = high); non-model countries in light grey. Panel (a): baseline, total 774 Mt, largest in Germany (252), Poland (103), and Italy (83). Panel (b): the announced 2040 build-out combined with the union — fossil generation is exactly zero, so the map is empty on the shared scale. Panel (c): natural-potential re-siting combined with the union at current demand, total 781 Mt: coordinated re-siting at today's demand reconcentrates fossil dispatch into low-cost countries (Germany rises to 321 Mt) rather than lowering the EU total. Panels (a) and (c) are processed at the baseline carbon price, panel (b) at the 2040 carbon price.

Figure H.8: From the estimated multiplier to the macro elasticity




Note: Each row is a response of firm mass or of value added to a change in the electricity price firms face, divided by the relevant electricity cost share so that all four sit on the scale of the estimated multiplier. The first row is the within-location, cross-sector object the Norwegian design identifies, recovered here from the model's own simulated reallocation. The second is one sector compared across countries, one point per country-sector cell. The third and fourth are the same countries measured two ways, by total firm count and by value added. Solved from a symmetric $\pm 1\%$ firm-facing electricity price wedge applied to one country at a time with the revenue rebated lump sum, using central differences. Elasticities are per unit of the price change firms actually face rather than per unit of the imposed wedge: as firms cut demand the national market price falls back, so between 2% and 50% of the wedge is undone before it reaches them, by more in countries whose supply curve is steeper. Value added is production-sector revenue net of electricity purchases, not the fuel-net aggregate $\mathcal{V}$ of Lemma 4.

Figure H.9: Sector location responses and the wage channel



Note: Each point is one of the thirteen model sectors in one of the six wedged countries: the sector's labour-to-electricity cost ratio α_s/β_s against its location response, normalised as in Figure H.8. The dashed line is the closed form of Appendix B.8, the partial-equilibrium multiplier scaled by $1 + (\alpha_s/\beta_s)\lambda_j$, evaluated at the mean of the solved λ_j ; the shaded band spans the solved range of λ_j across countries. The vertical band marks where the response crosses zero, at $1/|\lambda_j|$. Sectors to its right expand where electricity becomes more expensive, drawn by the wage released as electricity-intensive firms leave.

Figure H.10: Dose-response: sector electricity intensity and reallocation under the union



Note: Each point is one of the thirteen model sectors: the electricity cost share β_s (horizontal axis) against the cross-country dispersion of firm-mass changes $|\Delta \ln M_{js}|$ under the full union (vertical axis; filled markers the maximum across countries, open markers the 90th percentile). The four electricity-intensive divisions (paper, chemicals, non-metallic minerals, basic metals) sit in the upper right; services cluster near the origin.

Table H.3: Counterfactual overview: aggregate TFP, decomposition, prices, and emissions

Scenario	TFP (bp)	Mfg (bp)	Comp. (bp)	Agglom. (bp)	Alloc. (bp)	Fossil (bp)	Price	Range	ΔCO_2 , Mt (%)	Δfossil (TWh)	max $ \Delta M $ (%)
<i>Today's capacity stock</i>											
Baseline (segmented markets)	0.00	0.00	0.00	0.00	0.00	0.00	42.7	17.8	0.0 (0.0)	0	0.00
Full electricity union	5.94	13.38	2.54	1.50	0.02	1.89	43.7 [†]	0.0	+9.4 (+1.2)	-56	0.47
Full union, today's carbon price	14.69	13.42	2.53	1.51	0.02	10.63	70.2 [†]	0.0	-320.1 (-48.5)	-576	0.47
US-equivalent integration	2.11	-6.71	0.65	0.68	0.00	0.78	43.2	11.3	-6.5 (-0.8)	-28	0.33
US-level capacity, uniform expansion [†]	0.19	-11.34	-0.23	0.06	0.00	0.36	42.9	14.0	-2.9 (-0.4)	-15	0.20
Planned grid expansion (TYNDP 2040) [†]	1.50	-0.17	0.04	0.36	0.01	1.09	43.2	7.8	-2.7 (-0.4)	-34	0.30
Natural-potential re-siting [†]	-1.20	-6.20	-1.09	-0.42	0.00	0.30	42.6	16.3	+1.3 (+0.2)	-10	0.47
Re-siting + union	5.96	13.42	2.53	1.51	0.02	1.90	43.7 [†]	0.0	+9.0 (+1.2)	-56	0.47
<i>2040 renewable build-out</i>											
Announced national plans, current grid	16.19	-6.72	-1.82	-4.23	0.13	22.11	22.8	38.5	-655.6 (-100.0)	-1346	1.79
Announced national plans, planned grid	18.21	-6.79	-1.32	-2.61	0.03	22.11	21.5	19.6	-655.6 (-100.0)	-1346	0.50
Announced plans + union	26.16	13.42	2.53	1.51	0.02	22.11	20.8 [†]	0.0	-655.6 (-100.0)	-1346	0.47
<i>Energy crisis</i>											
Crisis-calibrated union [†]	11.96	19.68	3.69	2.13	0.03	6.11	148.0 [†]	0.0	-38.6 (-4.0)	-34	0.55
Crisis shock (2021–2023 fuel prices)	3.21	-3.36	-0.63	-1.37	0.03	5.17	103.9	83.6	-523.7 (-67.7)	-937	0.35
Crisis shock + union	12.69	13.42	2.53	1.51	0.02	8.63	106.7 [†]	0.0	-566.5 (-73.2)	-974	0.47
Crisis shock + planned grid [†]	6.76	-8.34	-1.47	-0.76	0.01	8.98	110.6	46.0	-576.1 (-74.5)	-987	0.32
Crisis shock + US-level grid	8.37	0.38	0.14	-0.31	0.01	8.53	111.0	69.8	-554.2 (-71.6)	-982	0.48
Crisis shock + 2040 build-out	16.96	-3.94	-1.20	-4.06	0.11	22.11	19.1	30.2	-773.8 (-100.0)	-1346	1.36
Crisis shock + re-siting	0.61	-2.80	-0.80	-1.66	0.08	3.00	98.6	97.2	-483.8 (-62.5)	-891	0.51

Notes: EU-23 sample, 2015–2019 baseline calibration; one row per counterfactual scenario, grouped by capacity stock. TFP columns report the production-side decomposition of equation (14) in basis points; composition, agglomeration, allocative efficiency, and fossil-resource use sum to the total up to rounding. The manufacturing column reports the within-manufacturing TFP change of the location channels, with output weights renormalised within manufacturing as in table H.4. Price is the demand-weighted average wholesale price across countries, marked with a dagger (†) where the scenario clears at a single union-wide price; Range is the cross-country span (max–min); both in EUR/MWh. ΔCO_2 is direct combustion CO_2 from modelled fossil generation, allocating each country's fossil dispatch across its calibrated fuel blocks by cost-minimisation at the country shadow price, with the percentage change relative to the scenario's own baseline emissions in parentheses; Δfossil is the change in fossil dispatch; max $|\Delta M|$ is the largest absolute country-level change in total firm mass. Crisis rows are measured against their own crisis baseline, and the today's-carbon-price row against a baseline processed at the same carbon price; the 2040 rows state their grid — current, planned (TYNDP-2040), or union — in the scenario name. A double dagger (‡) on the scenario name marks solves accepted at the solver's relaxed convergence tolerance; entries marked – are not computed for that scenario.

Table H.4: Within-manufacturing TFP effects across counterfactuals

Scenario	Within-group TFP gain (bp)		Max $ \Delta M $ (%)	
	Manufacturing	EI divisions	Manufacturing	EI divisions
Full electricity union	13.38	35.15	4.08	11.93
Planned grid expansion (TYNDP 2040)	-0.17	0.84	3.23	9.50
Natural-potential re-siting	-6.20	-17.90	2.40	7.32
Announced national plans, planned grid	-6.79	-18.40	5.92	17.59
Announced plans + union	13.42	35.31	4.12	12.05
Crisis-calibrated union	19.68	49.78	8.78	23.91

Notes: Within-group TFP change in basis points for manufacturing (divisions C17 paper, C20 chemicals, C23 non-metallic minerals, C24 basic metals, and remaining manufacturing) and the four electricity-intensive divisions. Each column renormalises output weights within its own group and comprises the location channels (composition, agglomeration, allocative); the fossil-resource channel is aggregate-only and is not attributed to sectors. Hence scenarios involving a full union produce identical gains. For each group, max $|\Delta M|$ is the largest absolute percentage change across countries in firm mass summed within that group, matching the definition in table H.3. The crisis union is measured against its own crisis baseline.