REM WORKING PAPER SERIES

From Fuel Dependence to Network Dependence: Cross-Border Electricity Interdependence and Strategic Vulnerability in Europe's Energy Transition

José Alves, João Estevão

REM Working Paper 0425-2026

July 2026

REM – Research in Economics and Mathematics
Rua Miguel Lúpi 20,
1249-078 Lisboa,
Portugal

ISSN 2184-108X

Any opinions expressed are those of the authors and not those of REM. Short, up to two paragraphs can be cited provided that full credit is given to the authors.

REM – Research in Economics and Mathematics

Rua Miguel Lupi, 20
1249-078 LISBOA
Portugal

Telephone: +351 - 213 925 912
E-mail: rem@iseg.ulisboa.pt

https://rem.rc.iseg.ulisboa.pt/

https://twitter.com/ResearchRem

https://www.linkedin.com/company/researchrem/

https://www.facebook.com/researchrem/

From Fuel Dependence to Network Dependence: Cross-Border Electricity Interdependence and Strategic Vulnerability in Europe's Energy Transition [*]

José Alves[†] João Estevão[‡]

July 2026

Abstract

Europe's energy geopolitics is usually told as a story of changing suppliers – Russian gas yesterday, Chinese clean technology tomorrow. Electrification adds a second, more local geopolitics, defined by who sits at the centre of the grid, who runs into bottlenecks, and who can call on flexibility when stress hits. Using public data alone, we build a bidding-zone-month panel covering 41 European zones over 2019-2025 and test six pre-stated hypotheses about how this internal layer redistributes price volatility, negative-price exposure, net imports, and cross-border price gaps. Three findings survive our identification checks. More cross-zonal capacity lowers net imports in average months, confirmed quasi-experimentally around the NordLink and Viking Link HVDC commissionings. Higher renewable shares raise within-month price volatility once network position is held fixed – about 1.8 EUR/MWh per ten percentage points of renewable share – concentrated in the network-central half of the panel. And the 2022 gas crisis widened the gap between EU-27 and non-EU European zones: integration transmitted the shock into the most-connected jurisdictions instead of dampening it. The flexibility-moderation prediction fails. Energy sovereignty in an electrified Europe is best understood as advantageous positioning within regional infrastructure, not separation from it; integration is double-edged.

Keywords: electricity interdependence; energy transition; strategic vulnerability; congestion; Europe; energy security

---

[*] This work was supported by national funds through FCT - Fundação para a Ciência e a Tecnologia, I.P., in the framework of the unit UID/0097/2025 and UID/06522/2025 – https://doi.org/10.54499/UID/06522/2025. The opinions expressed herein are those of the authors and do not necessarily reflect those of the authors' employers. Any remaining errors are the authors' sole responsibility.

[†] ISEG Research, ISEG - Lisbon School of Economics and Management, Universidade de Lisboa, Lisbon, Portugal; UECE - Research Unit on Complexity and Economics; CESifo (Center for Economic Studies and Ifo Institute). Corresponding author: jalves@iseg.ulisboa.pt.

[‡] Cegist, Instituto Superior Tecnico Universidade de Lisboa, Lisbon, Portugal; Advance/ISEG Research joaomjestevao@tecnico.ulisboa.pt

1. Introduction

In the months after Russia's invasion of Ukraine, one question defined European energy policy: who could replace its gas. Two years on, that question is still asked, and still loudly. But it no longer sits alone. A quieter one is taking its place: as Europe electrifies its way out of fossil fuels, who is becoming indispensable to the resulting system, and who is becoming exposed? This article argues that the answer is being decided less in the contracts Europe signs with external suppliers than in the wires, interconnectors, and bottlenecks that move electricity across the continent. This transition has two geopolitics: who sells the fuel, and where you sit on the grid. We are interested in the second.

The first is well charted. Critical-mineral rivalries, clean-technology manufacturing, hydrogen value chains, and the strategic exposure laid bare in 2022 have generated a substantial body of work, much of it sharp and historically informed (Scholten and Bosman, 2016; Scholten et al., 2020; Vakulchuk et al., 2020; Blondeel et al., 2021; Van de Graaf et al., 2020; Kuzemko et al., 2022; Colgan et al., 2023). But Bridge et al. (2013) made a point that has aged especially well: energy transitions are always spatial and infrastructural. They reorganise the geographies through which power, value, and vulnerability are distributed. In an electrified system, grids, balancing assets, and congestion points draw those geographies as much as extraction or manufacturing do – they decide who can move electrons when stress hits.

Nowhere is this second geography more politically charged than in Europe. The European Union pursues decarbonisation through market integration, regulatory harmonisation, and cross-border coordination, while political authority over energy remains territorially fragmented and contested. Goldthau and Sitter (2014, 2015) call this an unusual mix of liberal market-making and geopolitical exposure: governance through rules and competition inside a strategic environment shaped by coercion and asymmetric capacity. Recent shocks made the tension impossible to ignore. Mišík and Nosko (2023) read Europe's predicament as an energy-security paradox – resilience requires deeper cooperation, yet crises keep pushing member states toward unilateral defence. LaBelle (2024) argues this has reopened older questions of sovereignty and solidarity in European energy governance. Those questions motivate this article. But the answers, we think, depend on something the literature has not yet pinned down: the spatial structure of the electricity system and how it allocates the benefits and burdens of connection.

The classical case for that system is integration. Widening the geographic scope of trade should reduce balancing costs, smooth renewable variability, and lower prices. The evidence broadly agrees. Newbery et al. (2016) estimates sizeable gains from European market integration. Newbery et al. (2018) argues that high-renewables systems need market design that better values flexibility and transmission. Child et al. (2019) shows that a renewable European system depends on combining flexible generation, storage, and cross-border exchange, not on self-contained national systems. National autarky, on this reading, is not merely inefficient but incompatible with deep decarbonisation.

Yet the integration that produces efficiency also redistributes exposure, and not arbitrarily. It depends on where interconnector capacity sits, how much of it is actually offered to the market, how congestion is managed, and which jurisdictions hold the dispatchable low-carbon plant, hydropower, or storage that lets a system absorb shocks internally. Fabra (2023) writes that the post-2022 crisis turned once-technocratic market-design questions into politically consequential ones, because the rules shape who absorbs the blow and who passes it on. Studies of flow-based market coupling sharpen the point. Finck (2021), Schönheit et al. (2021), and Lang et al. (2020) show that the mapping from physical network limits to tradable capacity is itself contestable. Poplavskaya et al. (2020) demonstrates that integrating day-ahead trading with redispatch only expands cross-border exchange by making congestion management more explicit. Together, this literature describes cross-border electricity less as a flat public good than as a structured field of asymmetric reliance.

That asymmetry is our starting point. Vulnerability in an electrified Europe, we argue, is shifting from a predominantly fuel-centred to a predominantly network-centred form. The question is no longer only who imports from whom, but who the system can no longer balance without and who pays when constraints bind. Bidding zones that combine central network positions with flexibility – hydropower, dispatchable low-carbon generation, storage – should enjoy advantages that peripheral or congestion-prone zones do not. Peripheral zones should pay for that geography in higher price divergence, greater stress-period import reliance, and a higher risk that renewable expansion ends in curtailment rather than usable supply. We treat these as testable claims, not received truths.

To test them, we draw on three conversations that have rarely talked to each other. The first, on the geopolitics of the transition, typically operates at high abstraction or focuses on external supply chains (Scholten and Bosman, 2016; Scholten et al., 2020; Vakulchuk et al., 2020; Blondeel et al., 2021). The second, on European electricity-market integration, is empirically careful but typically treats congestion and redispatch as questions of efficiency or regulatory design, not sources of strategic asymmetry. The third, on energy security, has moved beyond simple import-dependence indicators (Sovacool and Mukherjee, 2011; Cherp and Jewell, 2014; Ang et al., 2015), yet still tends to operationalise vulnerability at the aggregate national level, missing what happens inside the grid. We bring them into contact and offer three contributions. Conceptually, we reorient European energy geopolitics from external suppliers toward internal infrastructures. Theoretically, we link energy security, European integration, and electricity-market design through three mechanisms – network position, congestion, and flexibility. Empirically, we build a public-data panel from ENTSO-E, ACER, JRC-IDEES, PVGIS, and Open Power System Data, observed at the bidding-zone-month level for 2019-2025, and use it to ask who gains, who pays, and how the answers shifted during the gas-crisis window of 2022.

The article proceeds in six steps. Section 2 reviews the literature the argument builds on. Section 3 develops the theoretical framework. Section 4 describes the data and the stylised facts the framework must organise. Section 5 lays out the empirical strategy. Section 6 reports the results. Section 7 concludes.

2. Literature Review

The literature on the geopolitics of the energy transition has outgrown the early hope that renewables would dissolve the political stakes of energy supply. Scholten and Bosman (2016) were among the first to insist they would not: renewables change the material bases of power without abolishing asymmetry or contestation. The board is new; the game continues. Overland (2019) listed the myths that needed retiring – that renewables are uniformly distributed, that they decentralise by themselves, that they are intrinsically democratic, that they are geopolitically benign. Scholten et al. (2020) and Vakulchuk et al. (2020) traced where the stakes actually moved: into technology supply chains, infrastructures, standards, and new patterns of state-business coordination. Blondeel et al. (2021) synthesised the shift: what changes in an energy-system transformation is not the existence of geopolitics but its objects, scales, and mechanisms. The picture is consistent. Decarbonisation reorganises the spatial architecture of dependence; it does not abolish it.

But the vocabulary of that scholarship has stayed largely outward-facing. It is at its best on control over critical minerals, manufacturing leadership in strategic technologies, the politics of hydrogen value chains, and competition over industrial policy. Van de Graaf et al. (2020), for instance, show how emerging hydrogen markets could open a new arena for international power projection. These are not minor concerns, but they crowd out a second transformation inside regional systems – one that does not reduce to who supplies what to whom. Albert (2022) pushed in the right direction: specific technologies and infrastructures in a green transition have no ready substitutes. Taken seriously, that point makes cross-border grids and balancing assets strategically salient: not generically scarce, but indispensable to the particular system they serve. Bridge et al. (2013) made the more general observation earlier still: transitions are spatial projects that reorganise the geographies through which power, value, and vulnerability travel. Meckling (2019) adds the institutional layer – those geographies are deeply uneven politically. For Europe, both observations converge on the cross-border electricity system, at once unusually integrated and unusually territorial.

Here the energy-security literature has the most to offer: it has spent two decades sharpening the vocabulary the geopolitics-of-transition debate now needs. Sovacool and Mukherjee (2011) and Ang et al. (2015) argued, in different ways, that energy security is irreducibly multidimensional. A credible measure has to track exposure, resilience, affordability, and governance – an import share alone will not do. Cherp and Jewell (2014) went further: the relevant analytic object is vulnerability within a specific energy system, not an abstract inventory of more or less risky sources. For electricity, that distinction is decisive. Security depends not just on whether resources exist but on whether the system can be synchronised, balanced, kept adequate, and held together through real-time stress. Counting fuel-import shares captures none of that.

Post-2022, this strand matured rapidly and became far more concrete about Europe. Colgan et al. (2023) put a striking number on the economic cost of Europe's fossil-fuel exposure during the Russia-Ukraine crisis. Kuzemko et al. (2022) show the same shock accelerated decarbonisation and revived energy-security politics – pressures that did not always pull the same way. The crisis literature is also clear-eyed about the limits of treating Europe as one actor with one security interest. Mišík and Nosko (2023) call this the energy-security paradox: integration is the only durable route to resilience, yet each acute crisis pushes member states back toward unilateral defence. LaBelle (2024) extends the diagnosis to constitutional terrain: sovereignty and solidarity are being reframed around the governance of interdependence rather than import substitution.

This builds on a longer governance tradition. Goldthau and Sitter (2014, 2015) traced how market liberalisation and rule-based governance became central to Europe's energy strategy; Szulecki (2016) and Knodt et al. (2020) showed that even the Energy Union's soft coordination modes can bite. That tradition has yet to go inside the grid – into the mechanisms by which internal electricity infrastructure redistributes security across member states as the system electrifies and renewable shares rise.

A more direct answer begins to emerge in work on electricity-market integration. The headline finding is well established: larger, more integrated markets generate welfare, lower balancing costs, and ease decarbonisation. Newbery et al. (2016) estimated sizeable gains from European market integration. Child et al. (2019) showed that the high-renewable system Europe says it wants is far easier to balance when storage, flexible generation, and cross-border exchange are optimised jointly rather than nationally. Studies of market convergence – Sáez et al. (2019), Cassetta et al. (2022), and Karahan et al. (2024) – broadly support the same conclusion while documenting how unevenly convergence has progressed. The collective picture is hard to dispute. National self-sufficiency in green electricity is no longer just inefficient. For any seriously decarbonising European economy, it is a fantasy.

The same literature also shows that integration is not symmetry. Congestion, redispatch, bidding-zone configuration, and the translation of physical network constraints into tradable commercial capacity determine who benefits from cross-border integration and who is exposed when the system tightens. Finck (2021) analyses flow-based market coupling and shows its market outcomes depend critically on how network constraints are represented. Schönheit et al. (2020, 2021) demonstrate that seemingly technical choices – generation-shift keys, the construction of flow-based domains – materially redistribute cross-border trading opportunities. Lang et al. (2020) document how coupling rule changes shift redispatch requirements unevenly across space, and Poplavskaya et al. (2020) show that integrating day-ahead trading with redispatch only widens the cross-border channel by making congestion-management politics more visible. Fabra (2023), writing after the post-2022 electricity crisis, draws the broadest conclusion: market design is itself a distributional mechanism, shaping the incidence of volatility, rents, and risk across jurisdictions.

A smaller set of studies comes closest to our question, treating cross-border electricity itself as a security relationship. Moore (2017) argues that interdependence can enhance or undermine security depending on its institutional structure. Hawker et al. (2017) show that national capacity mechanisms can collide with the European single market, exposing the tension between domestic security provision and regional integration. Beyza et al. (2020) evaluate cross-border interconnections from a security-of-supply perspective; Escribano et al. (2023) remind us that interconnectors are politically contested infrastructures whose legitimacy is never automatic. Even this literature tends to frame security as reliability engineering or as a dispute between national and European regulatory logics. It rarely asks our central question: how does cross-border electricity dependence sort member states into those the system cannot do without and those it leaves exposed?

This is where matters stand. Three research streams have brought us close, each from a different angle, and each stops just short of the question this article asks. The geopolitics-of-transition literature shows that decarbonisation redraws the geopolitical map but rarely follows the redrawing inside regional electricity systems. The energy-security literature supplies the right conceptual vocabulary but has yet to deploy it where Europe's transition is actually being fought – at the bidding-zone level. The electricity-market integration literature has the empirical apparatus and has identified the asymmetries, yet tends to read congestion and redispatch as efficiency or regulatory questions, not strategic-asymmetry questions. The rest of this article puts the three conversations to work on the same problem.

3. Theoretical framework

This paper asks who becomes strategically exposed in an electrified Europe – a question about how an integrated but heterogeneous regional electricity system distributes exposure across the jurisdictions it connects. Throughout, we work with a panel indexed by bidding zone $i \in \{1, \dots, I\}$ and month $t \in \{1, \dots, T\}$. The framework centres on what we call *network dependence*: the share of a jurisdiction's economic security and policy room for manoeuvre determined by its network position within, and its ability to use, the regional grid. In a fuel-centred system, the relevant accounting tracked suppliers, contracts, and physical disruption risks. They still matter in an electrified system – Europe's gas-import map is not irrelevant to its electricity politics – but they no longer suffice. A second layer of vulnerability joins them, rooted in transmission topology, balancing capability, congestion management, and the institutional rules that translate physical network limits into market outcomes. The framework that follows attempts to put that second layer on the same analytical footing as the first.

It draws on an economics-of-networks literature showing how interconnections amplify, redistribute, and concentrate shocks across structurally heterogeneous units (Allen and Gale, 2000; Acemoglu et al., 2012; Elliott et al., 2014; Atalay, 2017). Our contribution is to translate that abstract setting into the institutional and physical specifics of an integrated regional electricity system.

Three mechanisms formalise the second layer. The first is *network position, Nit,* a zone's centrality in, and ability to draw on, the regional grid. Interconnector degree – the count of adjacent zones with which i shares a physical border – is the natural baseline, but the empirically interesting object is the capacity-weighted version, which counts how much commercial cross-zonal capacity those borders actually offer:

$$N_{it} \;=\; \textstyle\sum_{j\in\mathcal{A}(i)} K_{ij,t}\,, K_{ij,t} \;=\; \text{available cross-zonal capacity from } i \text{ to } j \text{ in month } t, \quad (1)$$

where $\mathcal{A}(i)$ is the set of zones adjacent to i. The static (degree-only) version of N_{it} is what topology gives a zone; the time-varying (capacity-weighted) version is what regulators, transmission system operators, and market rules let it use. The distinction is political as much as technical: capacity-allocation choices make the same physical interconnection more or less useful, with distributional consequences.

The second mechanism is *congestion*, C_{it}, which measures how far nominally available interdependence is actually deliverable. Two zones can be physically connected, with cross-zonal capacity on paper, yet remain poorly integrated whenever flows are constrained or prices diverge persistently. Our preferred construction combines two complementary indicators – a price-spread component and a capacity-utilisation component:

$$C_{it} \;=\; \underbrace{\frac{1}{|\mathcal{A}(i)|}\textstyle\sum_{j\in\mathcal{A}(i)} \mathbb{E}_{\tau\in t}\left[\,\left|P_{i,\tau}-P_{j,\tau}\right|\,\right]}_{\text{price spread to neighbours}} \;+\; \lambda\,\underbrace{\frac{1}{|\mathcal{A}(i)|}\textstyle\sum_{j\in\mathcal{A}(i)} \mathbb{E}_{\tau\in t}\left[\frac{|F_{ij,\tau}|}{K_{ij,\tau}}\right]}_{\text{utilisation pressure on borders}}, \quad (2)$$

where $P_{i,\tau}$ is the hourly day-ahead price, $F_{ij,\tau}$ is the physical flow from i to j in hour $\tau, \mathbb{E}_{\tau\in t}[\cdot]$ denotes the within-month average over hours, and $\lambda \geq 0$ weights the components. The price-spread term captures the result of binding constraints; the utilisation term, proximity to them. A high reading signals a system either bumping against the rails or already off them.

The third mechanism is *flexibility*, F_{it}, the resources a zone can mobilise to absorb intra-day and intra-month variation without leaning on its neighbours. Hydropower, storage, dispatchable low-carbon generation, demand response, and well-positioned thermal capacity all count. A zone-month composite summarises them,

$$F_{it} \; = \; \omega_h \, h_{it} \; + \; \omega_s \, s_{it} \; + \; \omega_d \, d_{it} \; + \; \omega_n \, n_{it}, \qquad (3)$$

where h_{it}, s_{it}, d_{it}, and n_{it} are the shares of hydropower, storage, dispatchable low-carbon, and nuclear in the zone's installed-and-available capacity, and the weights $\omega_{\cdot} \geq 0$ reflect relative balancing usefulness. The composite is deliberately coarse. The point is not that this index is optimal but that flexibility must be measured by what it can do – cover residual-load swings, ride out stress hours, displace imports during scarcity – not by an abstract inventory of installed capacity.

Three quantities define what these mechanisms must explain. The first is the *net import share*,

$$\mathrm{NIS}_{it} \; = \; \frac{\sum_{\tau \in t}(\mathrm{Imports}_{i,\tau} - \mathrm{Exports}_{i,\tau})}{\sum_{\tau \in t} \mathrm{Load}_{i,\tau}}, \qquad (4)$$

which gives the fraction of consumption zone i meets from outside in month t. The second is *residual load*, $\mathrm{RL}_{i,\tau} = \mathrm{Load}_{i,\tau} - \mathrm{Wind}_{i,\tau} - \mathrm{Solar}_{i,\tau}$, and its within-month dispersion, $\sigma(\mathrm{RL}_{it})$, capturing how much variability the rest of the system must absorb. The third is *stress-period import reliance*, the conditional expectation of net imports during the month's most demanding hours:

$$\mathrm{SPIR}_{it} \; = \; \mathbb{E}_{\tau \in t}\big[\, \mathrm{NIS}_{i,\tau} \mid \mathrm{RL}_{i,\tau} \geq q_{0.9}\big(\mathrm{RL}_{i,t}\big) \;\text{ or }\; P_{i,\tau} \geq q_{0.9}\big(P_{i,t}\big) \,\big], \qquad (5)$$

where the threshold $q_{0.9}(\cdot)$ is the within-month 90th percentile of residual load or price. Average integration means little if reliance reverses in the hours that actually matter. Constructing SPIR_{it} requires hour-level net positions sliced to each zone's own stress windows. The monthly panel built here defines the object but does not yet carry it. The empirical tests below therefore run on the remaining outcomes, and the Conclusion flags SPIR as the natural extension. Alongside these three outcomes we track cross-border price divergence, $\mathrm{PD}_{it} = |\mathcal{A}(i)|^{-1} \sum_{j \in \mathcal{A}(i)} \mathbb{E}_{\tau \in t} \left| P_{i,\tau} - P_{j,\tau} \right|$, and the share of hours with negative prices – a curtailment-risk proxy for surplus renewables unable to clear.

These primitives make the main estimating equation compact. For an outcome Y_{it} drawn from {NIS_{it}, SPIR_{it}, PD_{it}, $\sigma(\mathrm{RL}_{it})$, negative-price share}, the baseline two-way fixed-effects specification is

$$Y_{it} = \alpha_i + \delta_t + \beta_1 N_{it} + \beta_2 C_{it} + \beta_3 F_{it} + \beta_4 (C_{it} \times F_{it}) + \mathrm{x}_{it}'\gamma + \varepsilon_{it}, \quad (6)$$

where α_i absorbs time-invariant zone heterogeneity (topology, climate, country-level market design), δ_t absorbs Europe-wide common shocks (gas-price moves, weather years, regulatory changes), and $\mathbf{x}_{it}$ collects time-varying controls (load, weather, and renewable shares). The coefficient β_1 tests whether better-positioned zones experience lower average vulnerability; β_2 whether binding congestion translates into measurable exposure. The interaction β_4 carries the flexibility argument: if flexibility is genuinely a buffer, the marginal effect of congestion on vulnerability,

$$\frac{\partial Y_{it}}{\partial C_{it}} = \beta_2 + \beta_4 F_{it}, \quad (7)$$

should shrink (less positive, or more negative) as F_{it} rises. That is testable.

These mechanisms are unlikely to operate uniformly across time. Crisis windows – periods of acute system stress – should sharpen, not blur, the distinction between buffer zones and peripheries. We model this by interacting the structural variables with a shock indicator $S_t \in \{0,1\}$,

$$Y_{it} = \alpha_i + \delta_t + \beta_1 N_{it} + \beta_2 C_{it} + \beta_3 F_{it} + \theta_1 (N_{it} \times S_t) + \theta_2 (F_{it} \times S_t) + \mathrm{x}_{it}'\gamma + \varepsilon_{it}, \quad (8)$$

where $S_t = 1$ for the gas-crisis window (2022 at the core, extended in sensitivity checks to 2021Q4 and 2023Q1) and zero otherwise. The shock-window hypothesis predicts $\theta_1 < 0$ on adverse outcomes (better network position helps *more* during stress) and $\theta_2 < 0$ likewise (flexibility helps *more* during stress).

The framework's substantive content reduces to falsifiable hypotheses with predicted signs and specific empirical counterparts. We state them now, before the data, so the empirical section reads as a verdict on each. Throughout, $L.\cdot$ denotes a one-month lag; $\sigma(\cdot)$ a within-month standard deviation; and r_{it} the renewable share of monthly generation.

H1 – Network position dampens average exposure. Capacity-weighted network position N_{it} is associated with lower cross-border price divergence and lower net-import share: $\beta_1 < 0$ in Equation (6) on PD_{it} and NIS_{it}. The prediction extends to stress-period import reliance ($SPIR_{it}$), which the panel defines but cannot yet measure.

H2 – Congestion raises exposure. Higher congestion C_{it} is associated with greater cross-border price divergence and a higher share of hours with negative day-ahead prices: $\beta_2 > 0$ in Equation (6) on PD_{it} and the negative-price-share outcome.

H3 – Renewable intermittency, conditional on network position. Holding N_{it} fixed, a higher lagged renewable share $r_{i,t-1}$ raises within-month price volatility: $\partial\sigma(P_{it})/\partial r_{i,t-1} > 0$ in specifications conditioning on $L.N_{it}$. The condition is essential: N_{it} and r_{it} correlate positively across zones, so an unconditional regression suffers the omitted-variable bias the framework's network channel makes precise.

H4 – Flexibility moderates the intermittency channel. Flexibility dampens the transmission from residual-load volatility to price volatility: F_{it}: $\beta_4 < 0$ in Equation 7, equivalently a negative cross-partial of $\sigma(P_{it})$ in $\sigma(\mathrm{RL}_{it})$ and F_{it}. This is the design's most demanding test – a second-derivative prediction tested with our coarsest measures – so we treat it as a falsifiable prediction worth testing, not the framework's headline claim.

H5 – Crisis dampening through the network. Better-positioned zones absorb aggregate price shocks with smaller increases in price divergence: $\theta_1 < 0$ on PD_{it} in Equation 8 when $S_t = \mathbf{1}\{t \in 2022\}$. The natural counter-hypothesis – integration transmits crisis prices into the most-connected zones rather than dampening them – predicts $\theta_1 > 0$ on the same outcome. The sign of θ_1 alone discriminates between the two.

H6 – Heterogeneity by network centrality. The renewable-intermittency channel of H3 is stronger, in absolute value, in network-central than in network-peripheral zones: splitting the panel at the within-sample median of $L.N_{it}$, the renewable-share coefficient on $\sigma(P_{it})$ is larger in the central than in the peripheral subpanel.

The six hypotheses are not equally easy to confirm. H1 and H3 directly predict conditional correlations between observable structural variables and outcomes. H2 asks the congestion composite to do work a coarser price-spread-only measure could not. H4 pins down a moderation pattern the data either delivers or refuses, with no room for partial confirmation. H5 stakes the framework against the naive integration-dampens-shocks reading; the sign of θ_1 alone determines which view the panel supports. H6 requires the H3 channel to differ systematically across the centrality split rather than hold uniformly.

The main empirical weight rests on the H1 - H3 - H6 triple, which together identify a network channel through which renewable intermittency translates differentially into price volatility across zones, and on H5's counter-hypothesis, the within-stress version of the same channel. Section 6 reports the verdict on each: H1, H3, and H6 hold in their predicted direction; H2 and H4 do not; H5 must be replaced by its counter-hypothesis.

What does the framework predict, and what does it not? It predicts asymmetry, not autarky. Integration remains welfare-improving on average, in the sense familiar from the electricity-market literature (Newbery et al., 2016; Child et al., 2019); the framework adds that the gains do not arrive symmetrically. Zones with high N_{it} and high F_{it} see lower NIS_{it} in average months and lower SPIR_{it} in the months that matter. Zones with low N_{it}, $highC_{it}$, and low F_{it} end up the other way. The framework does not say that interdependence is bad, that integration should be reversed, or that energy sovereignty requires self-sufficiency. It says the opposite: in an electrified Europe, sovereignty is best understood as a structural position within infrastructure, not as separation from it. The political question that follows is who pays for being on the wrong side of that geography. We turn to the data with that question in mind.

4. Data and Stylized Facts

The unit of analysis is the European bidding zone, observed monthly from January 2019 to December 2025 – 41 zones across 84 months, just over 3,400 zone-months after dropping a handful of missing-data corners. We chose the bidding zone, not the country, because that is where Europe's regulatory geography lives. Day-ahead prices clear at that level. Cross-zonal capacity is allocated at that level. The political economy of congestion – who is offered how much capacity on which border, on what terms, and for how long – runs in zone coordinates. Aggregating Italy's seven price areas into a single "Italy" or Sweden's four into a single "Sweden" would erase exactly the variation we want to study.

The panel therefore keeps multi-zone countries split (Italy into seven, Sweden into four, Denmark into two, Norway into five) and covers the EU-27 zones plus the neighbours ENTSO-E publishes alongside them: Switzerland, Great Britain, and the five Norwegian price areas. For border-level outcomes (price divergence, physical flows, capacity utilisation) we build a border-month subpanel from the same raw data. Where a zonal indicator deteriorates – typically for non-EU neighbours during regulatory reconfigurations – we fall back to a country-month aggregation that preserves the identifying logic. Appendix A lists every variable with its precise definition and source.

The 2019-2025 window serves substantive and empirical purposes that reinforce each other. Substantively, it opens in the late phase of pre-crisis European market integration, runs through the 2022 gas shock, and ends with the system having had two-plus years to settle into its current political equilibrium. Empirically, the window yields enough within-zone variation in prices, congestion, residual-load conditions, and weather stress for the comparative work the framework requires. These years are also politically meaningful in ways the analysis cannot ignore: they cover the period over which energy sovereignty in Europe stopped being an academic question.

The data architecture rests entirely on public sources, by design: any reader with the same downloads can rebuild the panel. The electricity-market core comes from the ENTSO-E Transparency Platform, which publishes hourly day-ahead prices, actual total load, generation by production type, cross-border physical flows, scheduled commercial exchanges, net positions, day-ahead cross-zonal capacities, annual installed-generation capacity, and generation-outage records. ACER adds the cross-zonal capacity and 70-percent-target documentation, the market-monitoring reports, and the congestion evidence ENTSO-E does not publish itself. JRC-IDEES contributes country-level structural features (sectoral energy use, generation mix, balancing-relevant infrastructure) that move too slowly to register in zone-month within-variation but matter for cross-sectional interpretation. Weather inputs – hourly 2-metre air temperature, 10-metre wind speed, and global horizontal irradiance at each country capital – come from PVGIS, which runs from 2005 through 2023. We had planned to use the ERA5 reanalysis, but its delivery times proved too slow; PVGIS is a defensible substitute. Open Power System Data serves for time-series prototyping and validation. No empirical claim rests on proprietary balancing-market data, restricted transmission models, or vendor indices.

Variables on both sides of the regressions follow the framework. Four outcomes capture different dimensions of strategic vulnerability. *Cross-border price divergence* (PD_{it}) is the within-month mean of the absolute day-ahead price gap between zone i and each adjacent neighbour, averaged across borders. *Negative-price exposure* is the share of the month's hours with a day-ahead price below zero. *Net import share* (NIS_{it}, Equation 4) and its high-stress counterpart *stress-period import reliance* (SPIR_{it}, Equation 5) are the share of monthly load supplied from outside, on average and during top-decile residual-load or price hours respectively. The stress-period version requires hour-level net positions inside each zone's stress windows, which the current monthly build does not carry, so it is defined here and left unestimated. *Residual-load volatility* is the within-month standard deviation of load net of wind and solar generation; it summarises how much intra-month variability the rest of the system must absorb.

On the regressor side, network position N_{it} enters as interconnector degree (a static count of adjacent zones) and as its capacity-weighted version, the sum of monthly available cross-zonal capacity to those neighbours (Equation 1). Congestion C_{it} is the price-spread-plus-utilisation composite of Equation 2; flexibility F_{it} is the hydro-storage-dispatchable-nuclear composite of Equation 3. The controls vector $\mathbf{x}_{it}$ contains the zone's load level, monthly mean 2-metre temperature, 10-metre wind speed, surface solar radiation, and the renewable share of generation (wind, solar, and hydro combined). Appendix A gives the full catalogue with units and source codes; Table 1 reports the univariate distributions.

Table 1: Summary statistics for the main outcomes, network variables, and controls

	N	*Mean*	*SD*	*Min*	*P50*	*Max*
Day-ahead price (EUR/MWh)	3360	94.947	80.148	0.966	73.483	547.598
Price volatility (EUR/MWh)	3360	37.445	31.81	0.254	28.856	239.034
Negative-price share	3360	0.012	0.03	0	0	0.321
Cross-border price gap (EUR/MWh)	3300	18.28	27.254	0	9.603	249.717
Mean hourly load (MW)	3364	8111.577	12071.993	501.595	3850.066	71641.313
Residual load (MW)	2491	5412.645	7335.37	-1190.04	3051.473	55477.043
Residual-load volatility	2490	1377.925	2166.056	95.451	732.552	19335.76
Net position (MW)	3120	-66.66	1759.439	-8224.428	-167.587	8307.681
Net import share	3120	0.078	0.598	-0.91	-0.031	3.942
Wind share of generation	2994	0.206	0.242	0	0.131	1
Solar share of generation	2994	0.058	0.074	0	0.026	0.568
Hydro share of generation	2994	0.273	0.323	0	0.116	0.998
Nuclear share of generation	2994	0.112	0.199	0	0	0.777
Interconnector degree	3444	3.61	1.82	1	4	11
Temperature (deg C)	2040	11.967	7.658	-7.516	11.682	29.781
Wind speed at 10m (m/s)	2040	2.996	0.964	1.089	2.9	7.006
Solar GHI (W/m^2)	2040	150.854	90.173	3.552	150.992	337.639

Notes: Bidding-zone-month panel covering 41 European zones over 84 months (2019m1-2025m12); observation counts vary by variable because of differential data coverage. All variables are defined in Appendix A. Sources: authors' calculations from ENTSO-E day-ahead prices, scheduled commercial exchanges, generation by type, and day-ahead net transfer capacities; PVGIS hourly weather at country capitals (TMY 2005-2023); OPSD time series; and ACER Market Monitoring Report 2023.

Before the regressions, three features of the panel stand out. None proves anything alone; together they form the descriptive geography the regressions will be asked to organise.

First, day-ahead prices have changed shape, not just level. Figure 1 traces the cross-zone monthly mean of zone-level day-ahead prices from January 2019 through December 2025, with the p10 - p90 envelope across zones each month. The mean sits around 40 EUR/MWh in the late pre-crisis years, spikes above 300 EUR/MWh in the third and fourth quarters of 2022, and retraces toward 70-90 EUR/MWh by the end of the panel. The envelope tells the more interesting story. Cross-zone dispersion was modest before 2021, widens sharply during the gas crisis, and stays visibly larger afterwards – even once the mean has come down.

Figure 2 sharpens the point from two angles: a year-by-year boxplot of the cross-border price gap (price_gap_mean) on the left, and a panel-mean time series of the monthly negative-price share (price_negative_share) on the right. Negative prices, a decade ago a niche feature of one or two zones, now appear regularly. The panel-mean monthly negative-price share is about 1.2 percent, but several zones reach 30 percent or more in individual summer months. The price-level story is well known. The price-shape story – negative-price hours in summer, persistent cross-zone gaps in winter – carries the distributional consequences, and it is what the framework targets.

Figure 1. Panel-wide trajectory of day-ahead electricity prices, 2019-2025.

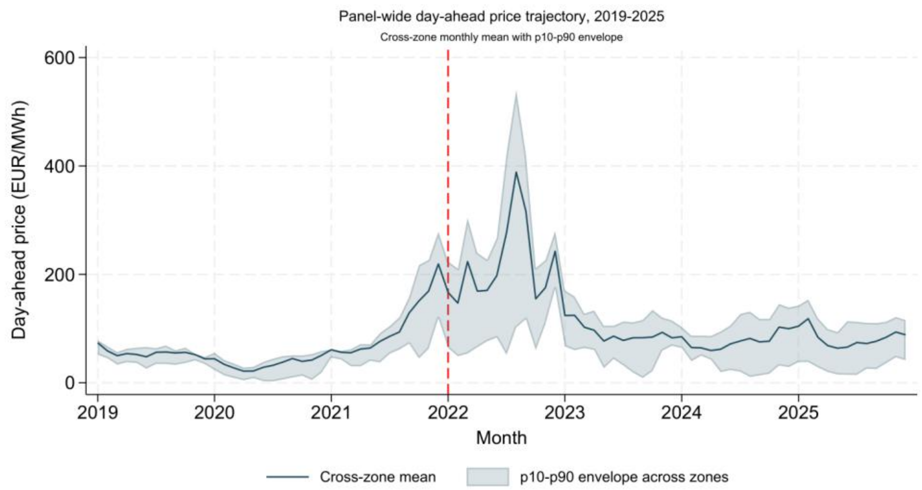

Notes: Cross-zone monthly mean (solid line) with the p10-p90 envelope across zones (shaded band). The vertical dashed line marks January 2022, the anchor of the gas-crisis window.

Figure 2. Cross-border price divergence and negative-price exposure across the panel.

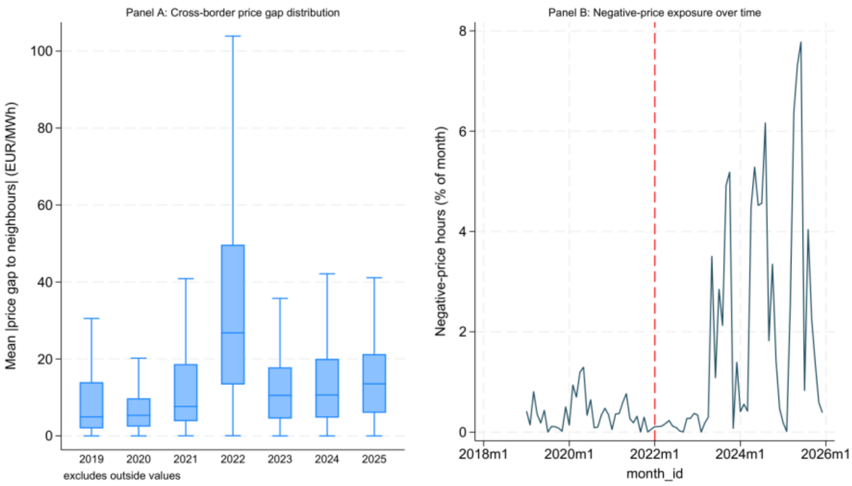

Notes: Left panel: distribution of within-month cross-border price gaps (PD_{it}) by year, across zones. Right panel: panel-mean monthly share of hours with a day-ahead price below zero, with the January 2022 gas-crisis anchor marked by the dashed vertical line.

Second, network position is profoundly unequal across the panel. Figure 3 ranks the 41 zones by interconnector degree – the count of adjacent zones with which each shares a physical border. The range runs from one (Ireland, with its single Great Britain cable; the peripheral Italian zones, each with one mainland tie) to eleven in the Germany-Luxembourg zone. The capacity-weighted version, built from the monthly available cross-zonal capacities of Equation (1), is more skewed still: a few zones command large fractions of the total cross-zonal capacity offered to European markets. Within zones, available day-ahead capacity moves visibly with capacity-allocation policy, not only with physical infrastructure. That is partly why congestion is a political object as much as a technical one: the same border can be offered as a 70-percent corridor or a 30-percent one, depending on what transmission system operators and regulators decide.

Figure 3. Network position across European bidding zones

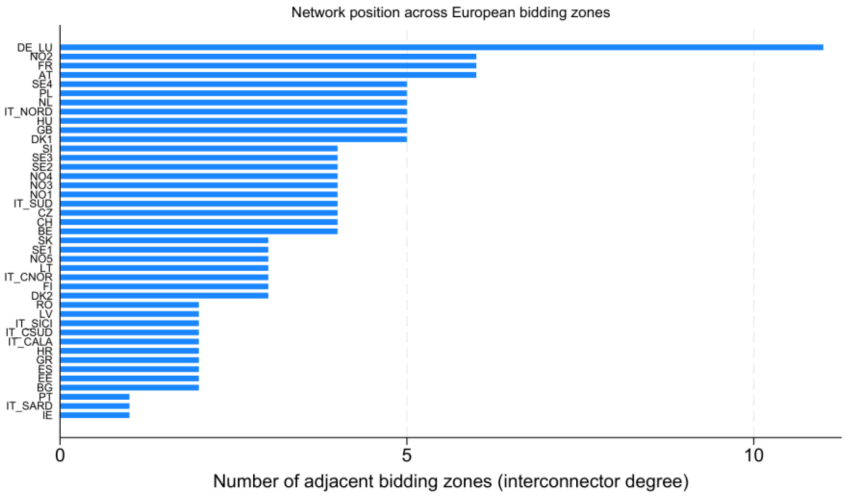

Notes: Bars show the interconnector degree (number of adjacent bidding zones with which each zone shares a physical border).

Third, the renewable expansion that has reshaped Europe's electricity supply has been uneven across zones, and residual-load volatility (residual_load_std) tracks that unevenness almost mechanically. Spain, Germany-Luxembourg, and Denmark have absorbed the largest renewable shares (renewable_share) over the panel period, and their monthly residual-load dispersion has grown most. Lower-share zones have moved less on both. None of these patterns is yet evidence of strategic exposure – they are correlations, not coefficients. They are the cross-sectional and temporal heterogeneity the regressions in the next section are designed to turn into structured tests.

5. Empirical Strategy

The estimation problem has a built-in tension, and the strategy below works with it rather than against it. The cross-section is small: 41 bidding zones, and the framework's structural variables – network position N_{it} and flexibility F_{it} – are slow-moving within zones, near-static for some. Pure within-zone identification of those coefficients is weak by construction. But the panel runs for 84 months and is rich in stress-period variation, so month-year fixed effects do real work, and the interactions at the framework's centre – congestion with flexibility, structural endowments with the gas-crisis window – are inherently time-varying. We exploit both kinds of variation. Each specification carries its own identifying assumptions; we do not ask all of them to estimate one unified treatment effect.

Our workhorse is a two-way fixed-effects panel regression. For outcomes $Y_{it} \in \{\mathrm{PD}_{it}, \text{neg-price share}_{it}, \mathrm{NIS}_{it}, \sigma(\mathrm{RL}_{it})\}$, we estimate

$$Y_{it} = \alpha_i + \delta_t + \beta_1 N_{i,t-1} + \beta_2 C_{i,t-1} + \beta_3 F_{i,t-1} + \beta_4 \left(C_{i,t-1} \times F_{i,t-1}\right) + \mathrm{x}_{it}'\gamma + \varepsilon_{it}, \quad (9)$$

where α_i are zone fixed effects, δ_t month-year fixed effects, and x_{it} collects the time-varying controls (load level, monthly mean temperature, 10-metre wind speed, surface solar irradiance, renewable share, and the Eurostat external-dependence indicators). The structural regressors enter with a one-period lag, for reasons below. We estimate by within-group OLS, with both fixed effects absorbed by two-way demeaning rather than estimated explicitly.

What does Equation (9) identify, and what does it not? The zone fixed effect α_i sweeps out everything time-invariant about a zone – physical topology, climate normals, country-level market design, the long-run technology mix. The time fixed effect δ_t absorbs Europe-wide shocks: gas-price moves, weather years, regulatory changes hitting every zone in the same month. That leaves the structural coefficients with within-zone variation in the time-varying components of N_{it}, C_{it}, and F_{it}. So $\hat{\beta}_1$ is identified from the capacity-weighted, allocation-dependent part of N_{it} rather than its degree-only baseline; $\hat{\beta}_3$ from the part of F_{it} that responds to capacity additions, retirements, and seasonal availability; $\hat{\beta}_2$ from month-to-month movement in the congestion composite. The interaction $\hat{\beta}_4$ carries the framework's central prediction: the marginal effect of congestion on vulnerability,

$$\frac{\partial \mathbb{E}[Y_{it}|\cdot]}{\partial C_{i,t-1}} = \hat{\beta}_2 + \hat{\beta}_4 F_{i,t-1}, \quad (10)$$

is the estimator-side analogue of Equation (7), identified from within-zone, within-month-year variation in the product $C_{i,t-1} \times F_{i,t-1}$.

We do not pretend this is causal in the textbook sense. No instrument exogenously varies network position, and no quasi-experimental break shifts flexibility allocation across otherwise-identical zones. What we offer is a defensible conditional-correlation design in the system-specific spirit of Cherp and Jewell (2014): how structural exposure moves with structural endowment, holding aggregate European conditions and time-invariant zone characteristics fixed. Where the data let us sharpen identification we do – through the shock-window comparison, the event studies below, and robustness samples that drop the most idiosyncratic zones. The estimates support the inference the framework asks for, not the recovery of a clean treatment effect.

We cluster standard errors at the zone level. The default variance estimator is the Liang-Zeger sandwich with within-zone block summation,

$$\widehat{\mathrm{Var}}(\hat{\boldsymbol{\beta}}) = (\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1}\left[\textstyle\sum_{i=1}^{I}\tilde{\mathbf{X}}_i{}'\hat{\boldsymbol{\varepsilon}}_i\hat{\boldsymbol{\varepsilon}}_i{}'\tilde{\mathbf{X}}_i\right](\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1}, \qquad (11)$$

where $\tilde{\mathbf{X}}_i$ stacks the demeaned regressors for zone i and $\hat{\boldsymbol{\varepsilon}}_i$ the corresponding residuals. Clustering matters because of within-zone serial correlation: chronic congestion, persistent net-import dependence, and elevated residual-load dispersion all autocorrelate strongly. The price of clustering, with $I = 41$, is that Equation (11) is modestly downward-biased in small clusters. Two robustness reports complement it: a wild-cluster bootstrap-t on the main coefficients, which corrects the small-cluster bias[1], and two-way clustering at zone and month-year, conservative by construction but informative for interaction-term precision. When we use the country-month fallback in Equation (15), we cluster at the country level instead.

The structural regressors enter with a one-period lag, $\{N_{i,t-1}, C_{i,t-1}, F_{i,t-1}\}$, for two reasons. Econometrically, contemporaneous shocks can move price spreads and flows in the same direction within a month, biasing β_2 on C_{it}. Substantively, capacity-allocation decisions and the seasonal availability of flexible plant are taken *before* the month begins, not simultaneously with prices. Residual-load volatility $\sigma(\mathrm{RL}_{it})$ is by construction a within-month object and enters contemporaneously; robustness specifications also lag it, to verify the moderation result on $C_{i,t-1} \times F_{i,t-1}$ does not turn on simultaneity.

The shock-window analysis sharpens the structural test: do the relationships among N, C, F, and the outcomes themselves shift during the gas-crisis window? The estimator-side version of Equation (8) is

$$Y_{it} = \alpha_i + \delta_t + \beta_1 N_{i,t-1} + \beta_2 C_{i,t-1} + \beta_3 F_{i,t-1} + \theta_1(N_{i,t-1} \times S_t) + \theta_2(F_{i,t-1} \times S_t) + \mathrm{x}_{it}'\gamma + \varepsilon_{it}, \qquad (12)$$

[1] Rademacher weights $\omega_i \in \{-1, +1\}$ applied at the cluster level over B=999 resamples.

with $S_t = \mathbf{1}\{t \in \text{calendar 2022}\}$ in the core specification and $S_t = \mathbf{1}\{t \in \text{2021Q4-2023Q1}\}$ in sensitivity extensions. The coefficients $\hat{\theta}_1$ and $\hat{\theta}_2$ are identified off how zone-level deviations in N and F during 2022 differ from those in surrounding non-crisis months. The tests are deliberately conservative: they ask whether structural endowments helped *more* in the year when help mattered most, controlling for zone fixed effects and the average crisis effect through δ_t.

We also run two event studies, both with a $\pm$ 24-month window. The first uses the European gas-crisis anchor $t^\star$ = January 2022. For each outcome we estimate

$$Y_{it} \;=\; \alpha_i \;+\; \sum_{\substack{k=-24\\k\neq -1}}^{+24} \mu_k \; 1\{t - t^\star = k\} \;+\; \mathrm{x}_{it}'\gamma \;+\; \varepsilon_{it}, \qquad (13)$$

with the reference period $k = -1$ omitted, and trace the profile $\{\hat{\mu}_k\}_{k\neq -1}$ as the system enters and exits the crisis. We test parallel trends directly: jointly significant pre-event coefficients ($\hat{\mu}_k$ for $k < 0$) would undermine the event-study reading, so we report the joint F-test of $\{\hat{\mu}_k = 0,\ k = -24, \ldots, -2\}$ alongside the profile. Cluster-robust inference follows the structure of Equation (11). The second event study replaces the common calendar window with within-zone stress events. For each zone i, define the stress set

$$\mathcal{S}_i \;=\; \big\{t : \sigma(\mathrm{RL}_{it}) \;\geq\; q_{0.9}\big(\sigma(\mathrm{RL}_{i\cdot})\big)\big\}, \qquad (14)$$

that is, the months in the top decile of zone i's residual-load-volatility distribution. We re-estimate Equation (13) with the anchor $t^\star$ replaced by each element of $\mathcal{S}_i$, average the resulting $\hat{\mu}_k$ profiles across stress events within zone, and apply the same pre-trends F-test to each anchor. The designs are complementary by construction. The first captures a system-wide shock hitting every zone on the same date; the second captures within-zone stress arriving on different dates, revealing heterogeneity the first averages over. The four claims in Equations (1-8) should leave fingerprints in both.

The robustness battery asks five questions, each implemented as a specific deviation from Equation (9).

1. We re-estimate on the EU-27 subsample, $i \in \mathcal{I}^{\mathrm{EU27}}$, to verify the non-EU neighbours (CH, GB, NO1-NO5) do not drive the result.
2. We drop the partial 2025 calendar year, restricting to $t \in$ [Jan 2019, Dec 2024].

3. We test sensitivity to the weather controls by restricting the PVGIS coefficients in $\boldsymbol{\gamma}$ to zero.

4. We estimate a country-month fallback that load-weights zone-level data to the country level using

$$Y_{ct} = \frac{\sum_{i \in \mathcal{C}(c)} L_{it} Y_{it}}{\sum_{i \in \mathcal{C}(c)} L_{it}}, \qquad F_{c,t-1} = \frac{\sum_{i \in \mathcal{C}(c)} L_{i,t-1} F_{i,t-1}}{\sum_{i \in \mathcal{C}(c)} L_{i,t-1}}, \qquad N_{c,t-1} = \sum_{i \in \mathcal{C}(c)} N_{i,t-1}, \quad (15)$$

where $\mathcal{C}(c)$ is the set of zones in country c and L_{it} is monthly load.

5. We test the moderation result against functional-form alternatives, replacing the linear interaction $C_{i,t-1} \times F_{i,t-1}$ in Equation (9) with a discretised-quartile interaction,

$$\sum_{q=2}^{4} \tilde{\beta}_q \; C_{i,t-1} \, 1\{F_{i,t-1} \in Q_q\}, \qquad (16)$$

where Q_q is the q-th within-panel quartile of F (Q_1 is the omitted base category), and, separately, with a quadratic interaction $\tilde{\beta} \; C_{i,t-1} \, F_{i,t-1}^2$.

None of these is the definitive specification; the moderation result stands on the consistency of $\hat{\beta}_4$'s sign and magnitude across them.

The sample is incomplete, and the gaps are not random. Of the 41 zones × 84 months = 3,444 zone-months in the design, just over 3,400 carry at least the day-ahead price block. The baseline regressions use materially fewer (1,747 for the price-volatility column): the PVGIS weather controls cover only the EU-27 capitals and end in 2023, and the lag structure consumes the first month of each zone's series. Remaining attrition reflects months in which ENTSO-E published no relevant indicator for a zone – most often the non-EU neighbours during regulatory reconfigurations and the Italian sub-zones in the panel's early years. Generation by type is the most attrition-prone series; net positions and day-ahead prices are the least. We drop missing observations row-wise rather than impute them, so each headline regression uses the largest balanced-by-variable sample available. The country-month fallback in Equation (15) is used precisely to recover the zone-months that drop out under zonal specifications.

The strategy has three limitations that shape what we are willing to conclude.

First, headline identification of N_{it}'s effects is largely conditional-correlational: we exploit within-zone variation in monthly NTC availability, not instrumented capacity shocks. The exception is the staggered TYNDP DiD around the NordLink and Viking

Link commissionings (Section 6, Figure 6), which anchors the H1 prediction on net imports quasi-experimentally.

Second, the gas-crisis window is a single calendar event hitting all zones at once, so $\hat{\theta}_1$ and $\hat{\theta}_2$ in Equation (12) identify heterogeneity in how zones absorbed it, not the crisis effect itself. The EU-27 vs Norway/Switzerland/UK DiD in Figure 5 addresses this with a sensible comparison group but does not pass parallel trends in the conventional sense.

Third, with 23 zone clusters and roughly 1,200 zone-months in the headline H3 regression (Table 3, col. 1), the panel is well-powered for the effect size the framework predicts on $\sigma(P_{it})$ but only borderline-powered for the H1 and H5-counter coefficients, which weaken under the Holm correction for multiple testing. Future panels with more years or wider zone coverage should sharpen those secondary findings. Together, the caveats frame the results as triangulated evidence: quasi-experimental identification anchoring H1, conditional-correlation identification supporting H3 and H6.

6. Empirical analysis

We now read the framework against the data: the regressions below organise the descriptive geography of Section 4 (Figures 1, 2 and 3, and Table 1), with each table serving as a verdict on one or more of the six hypotheses of Section 3. H1 (network position lowers exposure), H3 (renewable intermittency conditional on network), and H6 (centrality heterogeneity) survive in their predicted direction. H2 (the congestion-utilisation channel) and H4 (flexibility moderation) do not. H5 (crisis dampening through the network) is rejected; the panel returns the counter-hypothesis instead – integration transmits the shock. The mix of verdicts is itself substantive: the rejections sharpen the framework's reading rather than refute it.

Table 2 reports the baseline regressions from Equation (9): four outcomes, one column each, with zone and month-year fixed effects and zone-clustered standard errors. Because the PVGIS weather controls cover the EU-27 capitals through 2023 and the structural regressors enter with a one-month lag, the estimating samples are smaller than the full panel. The models absorb 31 zone and 59 month-year fixed effects, leaving 1,747 zone-months for the price-volatility and negative-price-share regressions and slightly fewer for the two outcomes built from the cross-border data.

Table 2. Baseline fixed-effects models of strategic vulnerability

	Price vol.	*Neg-price share*	*Net import share*	*Price gap*
L. Renewable share	-0.733	-0.012*	-0.028	-14.137
	(7.868)	(0.007)	(0.081)	(15.448)
L. Load mean	0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.001)
Temperature (deg C)	-0.068	-0.000	0.011**	0.553
	(0.401)	(0.000)	(0.005)	(0.567)
Wind speed at 10m (m/s)	-1.582	0.006**	0.039	-3.756
	(1.067)	(0.002)	(0.024)	(2.429)
Solar GHI (W/m^2)	0.029	0.000	-0.000	-0.006
	(0.028)	(0.000)	(0.000)	(0.037)
Constant	40.262***	-0.008	-0.033	31.087*
	(7.260)	(0.012)	(0.123)	(15.725)
Observations	1,747	1,747	1,649	1,711
R^2 (within)	0.003	0.037	0.015	0.007
F (joint)	2.27	3.46	1.60	2.47

Notes: Two-way fixed-effects panel regressions of Eq. (9) on each outcome. Zone and month-year fixed effects are absorbed in all columns. Standard errors in parentheses are cluster-robust at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

Table 3. Network position, congestion, and cross-border dependence

	Price vol.	*Net import share*	*Price gap*
L. Network capacity (GW)	-0.160	-0.008*	-0.792
	(0.390)	(0.004)	(0.963)
L. Border utilisation	4.450	0.099	25.359
	(21.221)	(0.087)	(53.334)
L. Renewable share	18.258***	-0.083	5.465
	(5.835)	(0.151)	(11.592)
L. Load mean	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.001)
Temperature (deg C)	-0.100	0.012**	0.759
	(0.378)	(0.004)	(0.727)
Wind speed at 10m (m/s)	-0.734	0.053*	-2.180
	(0.967)	(0.031)	(2.633)
Constant	35.959***	-0.204	14.398
	(6.997)	(0.187)	(12.672)
Observations	1,167	1,075	1,133
R^2 (within)	0.015	0.045	0.013
F (joint)	2.06	3.80	1.41

Notes: Adds lagged capacity-weighted network position $N_{i,t-1}$ from Eq. (1) (in GW), built from ENTSO-E day-ahead net transfer capacities, and the border-utilisation component of the congestion composite from Eq. (2). Zone and month-year fixed effects are absorbed in all columns; remaining controls follow Table (2) except the solar-irradiance term, which is omitted from the network specifications. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

The two goodness-of-fit reads tell different stories. The within-zone R^2 values are small – 0.003 for price volatility, 0.037 for the negative-price share, 0.015 for net imports, 0.007 for the price gap – because zone and month-year fixed effects absorb most variation by design. The joint F-statistics on the structural-regressors-plus-weather block, reported in each table, show the block jointly significant at 5% for the negative-price share and the price gap, at 10% for price volatility, and not at conventional levels for net imports. Most within-zone, within-month-year movement is unexplained – exactly what a fixed-effects design with slowly-moving structural regressors should produce. What remains is the conditional pattern.

Coefficient signs are informative even where magnitudes are modest. Conditional on temperature and wind, the lagged renewable share enters weakly negative on the negative-price share ($\hat{\beta}$ = -0.012), cutting against the naive more renewables means more curtailment' reading. Wind speed at 10m enters positively for the negative-price share ($\hat{\beta}$ = 0.006): windier months bring more surplus supply, more negative-price hours, and more reliance on the network to clear it. Temperature enters positively for net imports ($\hat{\beta}$ = 0.011), consistent with summer drought pulling hydro from peak generation and cross-border flows filling the slack. Load levels carry no marginal explanatory power once the fixed effects are absorbed.

Table 3 delivers the framework's two central confirmations: more network capacity goes with lower import dependence (H1), and with network position held fixed, renewables raise price volatility (H3). It replaces the count-only interconnector degree with the capacity-weighted network position from Equation (1), $N_{it} = \sum_{j \in \mathcal{A}(i)} K_{ij,t}$, built from the ENTSO-E day-ahead NTC series and expressed in GW. Monthly available cross-zonal capacity does move within zones – a typical EU bidding zone's summed monthly NTC fluctuates by 10-20% around its time mean – so the network coefficient is identified rather than absorbed by the zone fixed effects. The table also adds the border-utilisation component of the congestion composite in Equation (2).

The within-zone R^2 values improve on the simpler baseline (price volatility from 0.003 to 0.015, net imports from 0.015 to 0.045), and the H1 prediction surfaces as the framework anticipated: $L.N_{it}$ enters negatively on net imports (-0.008). Given a sample mean of 7.6 GW and a within-sample standard deviation of 5.0 GW for $L.N_{it}$, a one-standard-deviation increase is associated with -4.0 percentage points of net-import share. The same coefficient is negative but imprecise for price volatility (-0.160) and the price gap (-0.792): H1 is confirmed for the NIS outcome and points in the predicted direction, though imprecisely, for the other two. Border utilisation does not enter significantly on its own; H2 in its naive form finds no support.

With the network controls in place, though, the lagged renewable share enters strongly positively on price volatility (+18.258) – a result buried by collinearity in the simpler baseline of Table 2, and the direct confirmation of H3. A 10-percentage-point increase in lagged renewable share is associated with about +1.8 EUR/MWh of additional within-month price volatility, roughly 5% of the panel mean (a one-standard-deviation increase of +21 p.p. maps to +3.8 EUR/MWh, or 0.12 outcome standard deviations). This is exactly the renewable-intermittency channel the framework predicted – conditional on a zone's available cross-zonal capacity, more renewables go with measurably more volatility, a pattern the unconditional baseline could not have recovered.

The H3 magnitude sits within the electricity-economics literature on renewable-intermittency price effects: Ketterer (2014) finds that German wind feed-in lowers the day-ahead price level and raises its volatility, the merit-order literature estimates that renewables compressed German wholesale prices by roughly 6-10 EUR/MWh over 2010-2012 (Cludius et al., 2014), and Hirth (2013) documents the parallel decline in the market value of variable renewables at higher penetration. Our +1.83 EUR/MWh per 10 p.p. of lagged renewable share falls within this order of magnitude while extending the result to a much wider cross-section of European bidding zones and the unusually volatile 2019-2025 window.

The H3 coefficient also survives a selection-on-unobservables stress test. Under an Oster (2019) bound with $R^2_{\max} = 1$ (the most conservative ceiling; Oster's recommended 1.3 x R^2 exceeds one once the fixed-effects structure enters the controlled R^2), the implied δ^* for driving the renewable-share coefficient to zero is -3.9. Overturning the H3 finding would require unobserved selection almost four times as strong as observed selection, operating in the opposite direction – well above the conventional $|\delta^*| > 1$ robustness threshold.

Table 4 tests H4, the moderation prediction. It interacts residual-load volatility – the renewable-intermittency signal as a single zone-month scalar – with the flexibility composite F_{it} from Equation (3). Figure 4 traces the implied marginal effect $\hat{\beta}_2 + \hat{\beta}_4 \, F_{i,t-1}$ from Equation (10) across the empirical range of $F_{i,t-1}$.

Table 4. Flexibility, residual-load volatility, and vulnerability

	Price vol.	*Neg-price share*	*Net import share*
Residual-load volatility	0.001	0.000***	0.000***
	(0.001)	(0.000)	(0.000)
Flexibility (hydro + nuclear share)	-5.652	0.001	0.628***
	(9.777)	(0.011)	(0.152)
Residual-load vol. × Flexibility	0.007	-0.000	0.000
	(0.005)	(0.000)	(0.000)
L. Renewable share	10.813	-0.003	-0.256***
	(6.576)	(0.005)	(0.093)
L. Load mean	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)
Temperature (deg C)	-0.397	-0.000	0.009*
	(0.400)	(0.000)	(0.005)
Constant	40.950***	0.007	-0.116
	(6.811)	(0.006)	(0.104)
Observations	1,469	1,469	1,430
R^2 (within)	0.012	0.022	0.057
F (joint)	3.14	28.11	5.33

Notes: Tests the flexibility-moderation specification in Eq. (10) with C≡residual-load standard deviation and F≡ hydro-plus-nuclear share of monthly generation. The implied marginal effect $\hat{\beta}_2 + \hat{\beta}_4 F_{i,t-1}$ across the empirical range of F is plotted in Figure 4. Zone and month-year fixed effects absorbed. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

Figure 4. Marginal effects of residual-load volatility across flexibility endowments.

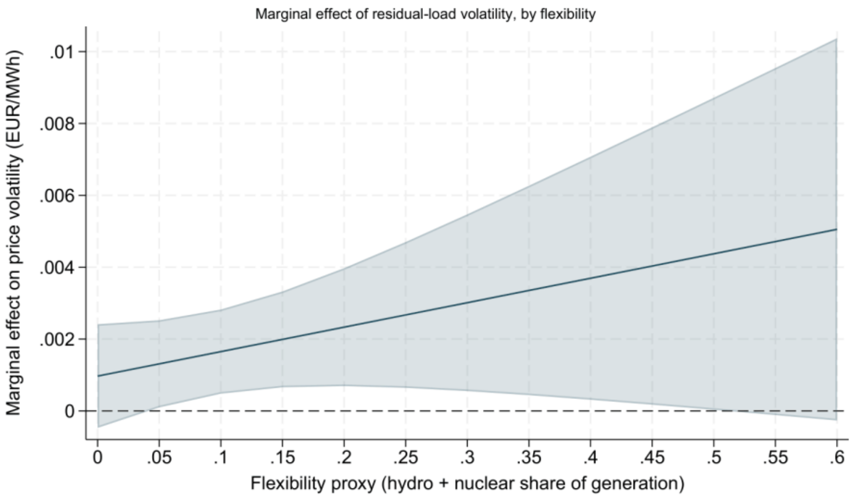

Notes: The solid line traces $\hat{\beta}_2 + \hat{\beta}_4 F_{i,t-1}$ from Equation (10); the shaded band is the pointwise 95\% confidence interval

H4 – the framework's signature prediction about flexibility moderation – is rejected outright in this baseline implementation. The interaction coefficient on price volatility is small (+0.007) and statistically indistinguishable from zero, so the predicted $\beta_4 < 0$ pattern fails; the marginal-effect line in Figure 4 stays flat across the empirical range of $F_{i,t-1}$. Appendix B's quartile-discretised version (Table 10) rises monotonically across flexibility quartiles (+0.001, +0.004, +0.010 from Q2 to Q4) – the opposite of the dampening H4 predicts. The explanation is unflattering but simple: within-zone variation in the hydro-plus-nuclear composite is too small and too noisily measured to identify moderation, most likely because our measurement confounds flexibility with network centrality – a point the Conclusion returns to. The negative-price-share regression returns a precisely positive but essentially zero correlation between residual-load volatility and negative-price exposure – right in sign, tiny in size.

The net-import-share regression is the most informative result in Table 4. Residual-load volatility raises net imports with a small but precise coefficient (p-value < 0.01); the flexibility composite enters sharply positive at $\hat{\beta}_3$ = 0.628, the lagged renewable share sharply negative at -0.256. Zones with more residual-load variability lean harder on net imports, as the framework's residual-load mechanism predicts. Zones with structurally higher hydro-plus-nuclear shares are also higher net importers on average – likely because the hydro-and-nuclear-rich zones (the Nordic core, parts of the Alpine system, France) are also the network-central ones, running higher trading throughput in both directions. The negative renewable-share coefficient fits renewables substituting for imports at the margin. This regression delivers the structured pattern the design was built to detect, even where the price-volatility specification does not.

Table 5, which interacts the structural regressors with a 2022 dummy via Equation (12), contains the section's central results. Two interactions are substantively non-trivial.

Table 5. Shock-window evidence

	Price vol.	Neg-price share	Net import share	Price gap
Renewable share	16.863***	0.025**	-0.015	4.125
	(5.766)	(0.009)	(0.233)	(13.859)
Network capacity (GW)	-0.433	0.000	-0.009*	-1.033
	(0.441)	(0.000)	(0.005)	(1.172)
Renewable share $\times$ 2022	-4.891	-0.015**	0.065	43.433**
	(14.618)	(0.005)	(0.160)	(18.909)
Network capacity (GW) × 2022	-0.186	0.000	0.002	2.092*
	(0.568)	(0.000)	(0.008)	(1.111)
L. Load mean	0.000	0.000**	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.001)
Temperature (deg C)	-0.118	0.000	0.012**	0.461
	(0.394)	(0.000)	(0.005)	(0.923)
Constant	37.477***	-0.010	-0.078	14.408
	(4.682)	(0.006)	(0.155)	(11.484)
Observations	1,164	1,164	1,070	1,128
R^2 (within)	0.017	0.051	0.027	0.041
F (joint)	3.09	4.10	4.80	2.48

Notes: Tests the shock-window specification in Eq. (12) with D_{2022}=1 for calendar year 2022 and zero otherwise. Each column interacts the structural regressors with D_{2022} on a different vulnerability outcome. Zone and month-year fixed effects absorbed; the level of D_{2022} is absorbed by the month-year fixed effects and is not reported separately. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

The network-capacity-by-2022 interaction directly rejects H5. The prediction was $\theta_1 < 0$ on the cross-border price gap – better-positioned zones absorbing the shock with smaller increases in price divergence. The estimate is $\hat{\theta}_1$ = +2.092: zones with more available cross-zonal capacity ended the gas year with *larger* price gaps to their neighbours, not smaller – at the shock-window mean gap of 20.4 EUR/MWh, +2.1 EUR/MWh per additional gigawatt of lagged network capacity. The data support the counter-hypothesis nested in H5, that integration transmits aggregate price shocks into the most-connected jurisdictions instead of dampening them; the sign of $\hat{\theta}_1$ alone discriminates between the two. Most plausibly, more-connected zones transmitted the European-wide shock faster and more completely, with differentials depending less on capacity headroom than on which side of the gas-driven marginal fuel each zone was trading. Network capacity's main effect on net imports stays negative and weakly significant (-0.009, p-value < 0.10): integration did not stop working in 2022 – H1 continues to hold for NIS – it delivered the crisis efficiently into the connected jurisdictions.

The renewable-share-by-2022 interaction is statistically zero for price volatility and net imports but strongly positive for the cross-border price gap (+43.4, p-value < 0.05): a 10-percentage-point higher pre-2022 renewable share was associated with an additional +4.3 EUR/MWh of price gap during 2022 – about +21% of the shock-window mean, on top of the network-capacity channel. This is evidence that renewable-heavy zones decoupled from the price level set by gas-driven marginal-fuel pricing elsewhere in the system. Whether that decoupling is benefit or exposure depends on the zone and the sign of the gap; politically, 2022 sharpened the boundary between renewable-heavy and gas-priced parts of the European system rather than homogenising them.

To put the H5-counter dynamic in policy units, vary only network position N_{it} and hold everything else fixed. Moving from the within-sample 25th percentile of $L.N_{it}$ ($\approx$ 3.2 GW) to the 75th percentile ($\approx$11.4 GW) raises the 2022 network-channel contribution to the predicted price gap from $\hat{\theta}_1 \times 3.2 \approx$ +6.7 to $\hat{\theta}_1 \times 11.4 \approx$ +23.8 EUR/MWh – a differential of roughly +17.1 EUR/MWh, about 93% of the full-panel mean price gap of 18.3 EUR/MWh (at the Q90-Q10 tails, 1.7 to 15.2 GW, the differential is +28.2 EUR/MWh). Integration is double-edged: it lowers net-import dependence in average months (H1) and amplifies crisis-month price-gap exposure in proportion to network position. The policy question is not whether to integrate but which zones to integrate at which level, and which congestion-management and price-cap instruments can dampen the within-crisis amplification.

Figure 5. Difference-in-differences around the 2022 gas-crisis window.

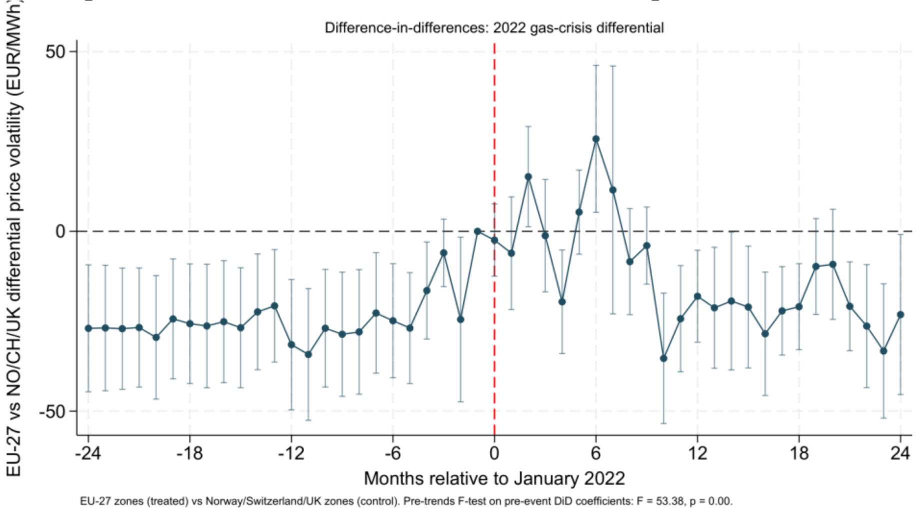

Notes: Coefficients $\hat{\delta}_k$ on (event-time k) × (EU-27 treatment indicator), with month $k = -1$ as the reference period and zone and month-year fixed effects absorbed. The treated group is the 34 EU-27 bidding zones; the control group is the seven non-EU European market zones (Norway 1-5, Switzerland, Great Britain). GB's day-ahead price series in the panel ends in 2020, so the effective post-period control group is the five Norwegian zones and Switzerland. Standard errors clustered at the zone level. The pre-trends F-test on $\{\hat{\delta}_k = 0,\, k = -24, \dots, -2\}$ is 53.58 (p-value=0.000) reported in the figure note.

Figure 5 reports a difference-in-differences estimate of the 2022 differential, built to address the pooled event study's parallel-trends problem. The 34 EU-27 bidding zones form the treated group; the controls are the seven non-EU European market zones (Norway 1-5, Switzerland, Great Britain), which share the geography and physical network but sit outside the EU's coordinated 2022 policy response (GB's day-ahead series ends in 2020, so the effective post-period controls are the five Norwegian zones and Switzerland). The specification is the standard dynamic DiD on price volatility, with zone and month-year fixed effects and zone-clustered standard errors.

The static post-2022 differential is $\hat{\delta}$ = +11.43 EUR/MWh and statistically significant: averaged over the 24 months from January 2022, EU-27 zones were about 30% more volatile than NO/CH/UK zones, conditional on the fixed-effects structure. The dynamic profile concentrates in the first nine months of 2022, peaks around +25 EUR/MWh, and fades back toward the pre-shock differential by mid-2023. The pre-trends test still rejects (F = 53.38, p-value < 0.01), but the rejection now reflects a converging level differential over 2020-2021 rather than the runaway drift that ruled out the pooled event study (F = 99.34; hence its relegation to Appendix B) – so the static coefficient is a defensible-but-not-textbook estimate of the 2022 EU-27-specific differential. Triangulated against the shock-window regression in Table 5, the two estimators agree in direction and substantively in magnitude: 2022 sharpened the boundary between the EU-27 and its non-EU neighbours, and the H5 verdict – integration transmits the shock instead of dampening it – is robust across specifications.

Figure 6 pushes the H1 verdict beyond the in-text battery to a quasi-experimental design. The capacity-weighted network position N_{it} used in Table 3 mixes physical-infrastructure variation with capacity-allocation policy, and both can move endogenously with the outcome. To isolate the physical channel, we estimate a staggered difference-in-differences around two major TYNDP HVDC additions within the panel window: NordLink (1.4 GW, NO2-DE_LU, commercial operation from spring 2021, coded 2021m5) and Viking Link (design capacity 1.4 GW, DK1-GB, coded 2024m1, initially limited to roughly 800 MW by Danish-side grid constraints). Both endpoints of each cable are treated from the commissioning month forward; the 37 remaining zones serve as never-treated controls. Three caveats apply. North Sea Link (1.4 GW, NO2-GB, October 2021) adds a second capacity step to the already-treated NO2, while IFA2 (1.0 GW, 2021) and ElecLink (1.0 GW, 2022) connect Great Britain with France, which sits in the control pool. GB's price and net-position series in our panel end in 2020, so the Viking Link arm contributes post-treatment variation through DK1 alone. Re-estimating without FR and GB leaves the estimates essentially unchanged (OLS-TWFE -0.087, SE 0.025; Callaway-Sant'Anna 12-month post-event -0.133), so neither the contaminated control nor the data gap drives the result. We report both a classical OLS two-way-fixed-effects static DiD and the Callaway-Sant'Anna doubly-robust estimator, which accommodates heterogeneous treatment effects across cohorts.

The net-import-share result is the clean H1 confirmation the article was looking for. OLS-TWFE returns $\hat{\beta}$ = -0.087: commissioning a new HVDC cable cuts directly-affected zones' net-import share by 8.7 percentage points relative to never-treated controls. The Callaway-Sant'Anna estimator within a $\pm$ 12-month event-time window agrees ($\hat{\beta}$ = -0.134), and its pre-event average is statistically zero ($\hat{\beta}_{\text{pre}}$ = -0.048, p-value = 0.191) – the parallel-trends test the EU-27 DiD failed, the TYNDP DiD passes. H1's effect on NIS_{it} is thus directly identifiable from exogenous variation in physical capacity, and substantially larger than the -0.008-per-GW correlational coefficient in Table 3: the latter is identified from sub-gigawatt monthly NTC fluctuations, the DiD from a gigawatt-scale step change. H1 survives both identifications, and the quasi-experimental one is the stronger.

Figure 6. Callaway-Sant'Anna staggered difference-in-differences around two TYNDP capacity additions.

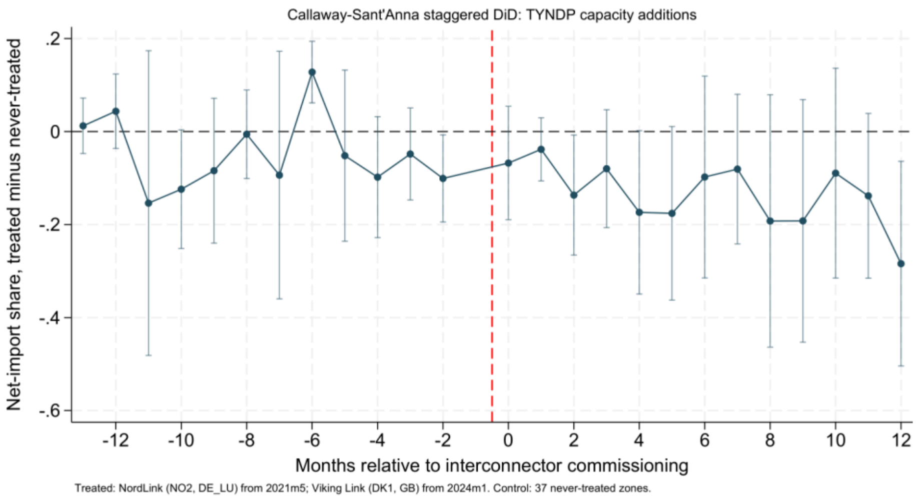

Notes: Treated zones are both endpoints of NordLink (NO2, DE_LU; commercial operation 2021m5) and Viking Link (DK1, GB; commercial operation 2024m1); the 37 remaining bidding zones serve as never-treated controls; GB contributes no post-treatment observations (its price and net-position series in the panel end in 2020), so the Viking Link arm is identified from DK1. Estimator: doubly-robust DR-IPW. The vertical dashed line marks the month before commissioning; the horizontal dashed line marks zero. Standard errors clustered at the zone level.

The cross-border-price-gap DiD is positive (+8.70 EUR/MWh, p-value = 0.068 in OLS-TWFE; +8.36, p-value = 0.032 in Callaway-Sant'Anna 12-month post-event) – another instance of the H5-counter dynamic. But the 12-month pre-event coefficient is also positive (+5.19, p-value = 0.063), so this DiD does not cleanly satisfy parallel trends and we do not lean on it causally. The NIS result is the substantive headline of the TYNDP quasi-experiment: more cross-zonal capacity, lower net-import dependence, exactly as the framework predicts and the correlational Table 3 suggests.

Table 6. Robustness checks and fallback country-month models.

	EU-27 only	2019-2024	No weather	+ Network
L. Renewable share	-0.895	-0.895	5.004	18.258***
	(7.876)	(7.876)	(5.701)	(5.835)
L. Load mean	0.000	0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Temperature (deg C)	0.024	0.024		-0.100
	(0.385)	(0.385)		(0.378)
Wind speed at 10m (m/s)	-1.858**	-1.858**		-0.734
	(0.906)	(0.906)		(0.967)
L. Network capacity (GW)				-0.160
				(0.390)
L. Border utilisation				4.450
				(21.221)
Constant	44.059***	44.059***	31.424***	35.959***
	(5.949)	(5.949)	(3.926)	(6.997)
Observations	1,747	1,747	2,953	1,167
R^2 (within)	0.002	0.002	0.002	0.015
F (joint)	1.53	1.53	1.21	2.06

Notes: Each column re-estimates the baseline price-volatility specification under an alternative sample or control set. Col. 1 restricts the panel to EU-27 zones (drops Norway, Switzerland, the UK); Col. 2 drops partial-year 2025; Col. 3 omits all PVGIS weather controls, expanding the sample to 2,953 zone-months; Col. 4 augments the baseline with lagged capacity-weighted network position $L.N_{it}$ and border utilisation, reducing the sample to 1,167 zone-months because NTC data is sparser than weather data. Zone and month-year fixed effects absorbed throughout. Cluster-robust SEs at the zone level.

Table 6 runs the baseline price-volatility regression through four alternatives: the EU-27 subsample, 2019-2024, a no-weather specification, and the baseline augmented with network position and border utilisation. Columns 1 and 2 are numerically identical, for an informative reason: the PVGIS weather controls exist only for the EU-27 capitals through 2023, so the weather-controlled sample is already EU-27-only and pre-2025. (The robustness columns, like Table 3's network specifications, omit the solar-irradiance control, so their coefficients are deliberately not digit-for-digit those of Table 2.) Column 3 expands the sample from 1,747 to 2,953 zone-months once those controls are removed: the weather variables, not the structural regressors, drive the sample restriction. Column 4 adds $L.N_{it}$ and border utilisation; sparser NTC coverage cuts the sample to 1,167 zone-months, but the within-zone R^2 rises from 0.002 to 0.015 and the lagged renewable share, indistinguishable from zero in the baseline, now enters strongly positive (+18.258, p-value < 0.01). The substantive findings of Table 3 therefore survive re-estimation as a baseline-augmented specification: conditional on network position, renewables raise volatility, not the other way round.

Beyond the in-text battery, Appendix B reports six further identification checks: a pre-trend correction for the event study, a placebo-year exercise, standardised baseline coefficients, a quartile-discretised flexibility moderation, alternative standard-error estimators, and a sample split by network centrality. Two matter for the main reading. The placebo-year test on the cross-border price gap (Table 8) returns null interactions for 2020, 2021, and 2023 and significant ones only for 2022: the shock-window finding is genuinely about the gas-crisis year. The network-centrality split (Table 12) tests H6 directly: the renewable-share coefficient on price volatility is roughly 50% larger in the more-connected half of the panel (+20.8***) than in the peripheral one (+14.0*), confirming H6 – H3's intermittency channel concentrates where the network can best redistribute the variability.

Two checks return less comfortable news. A linear pre-trend correction does not rescue the event study: the detrended pre-trend F rises to 646.8 instead of falling toward zero. The quartile-discretised flexibility interaction likewise rises monotonically rather than falling, suggesting flexibility and network centrality are confounded in our measurement. The standardised-coefficient and alternative-SE checks leave the baseline's qualitative reading unchanged.

Three further checks guard the inferential base; Appendix B reports the full estimates. First, the bounded [0,1] negative-price-share outcome does not drive the OLS signs: a fractional-response logit (Papke-Wooldridge GLM with zone and month-year dummies) returns average marginal effects of the same sign and similar magnitude (lagged renewable share -0.0086 against the OLS -0.0122; wind speed at 10 m +0.0048 against +0.0062). Second, pairwise residual correlations across the four baseline outcomes are below 0.10 in five of six pairs (the one moderate exception: $\hat{\rho}$ = 0.38 between price volatility and the price gap), so equation-by-outcome OLS is approximately as efficient as a joint SUR specification. Third, a Holm step-down correction across the four headline coefficients (H1 on NIS_{it}, H3 on price volatility, and the two 2022 interactions) controls the family-wise error rate: H3 survives at the 5% level (adjusted p-value = 0.020), the renewable-share interaction is borderline (adjusted p-value = 0.094), and the network-capacity interaction and H1 fall short (adjusted p-value = 0.146 each). We therefore read H3 – backed quasi-experimentally by the TYNDP DiD in Figure 6 – as the article's most robust claim, with H1 and the H5 counter-hypothesis directionally consistent but individually weaker.

Taken together, the estimates place the article's substantive contribution on the network side of the framework. H1, H3, and H6 identify a single coherent channel: capacity-weighted network position lowers net-import dependence; renewable intermittency raises price volatility once network position is conditioned on; and the intermittency effect concentrates in the network-central half of the panel, where the network is most able to redistribute the variability. H5 gives way to its counter-hypothesis – integration transmits aggregate price shocks into the most-connected zones, with 2022 the panel's clearest demonstration – and H4 is rejected in both linear and quartile-discretised specifications, for reasons the Conclusion diagnoses as measurement rather than theory. Within-zone R^2 values are modest by design, so the identifying power lives in the conditional patterns the surviving hypotheses describe rather than in any single high-precision coefficient. Those patterns are nonetheless economically and politically meaningful: the political geography of Europe's electricity transition is increasingly written in network coordinates, and the zones running highest on the network channel absorb the largest within-system price shocks during stress periods.

The findings reframe what solidarity means in practice. If vulnerability across Europe reflects not only fiscal capacity or fossil-fuel exposure but also grid topology and the way capacity and congestion are allocated, then solidarity cannot be reduced to emergency burden-sharing aftershocks arrive. It also concerns the *ex-ante* allocation of interconnection, congestion-management rules, balancing-resource deployment, and the politics of investment in system flexibility. The distributional politics of market design, in other words, are part of the politics of European cohesion – and the shock-window interactions in Table 5 make that lesson concrete even where the headline regressions in Table 2 are imprecise.

Two scope-of-claim points keep the empirical contribution from being over-read. First, the article does not deny the continuing importance of external suppliers, mineral supply chains, or industrial policy, and it does not argue that integration was a mistake: the negative network-capacity coefficient on net imports in Table 3 is consistent with integration delivering its expected first-order benefit, while the positive 2022 interaction in Table 5 reflects that same integration transmitting the gas-crisis shock. Second, the panel identifies conditional correlations within a two-way fixed-effects design, not textbook treatment effects, so the surviving hypotheses describe robust structural patterns rather than instrumented causal magnitudes. A member state may cut its fossil-fuel exposure and still remain strategically vulnerable if it occupies a peripheral place in the electricity network – a possibility our estimates make empirically visible rather than merely asserted.

The analysis implies a different way of thinking about policy trade-offs. The conventional question asks whether more interconnection is desirable; the more useful one – the one the framework is built to address – is which interconnections, under which rules, alongside which flexibility resources. The 2022 interactions in Table 5 make this concrete: zones with more capacity at their borders ended the gas year with larger price gaps, because the network delivered the shock efficiently, while Table 3 associates the same capacity with lower net imports in average months. Integration without attention to the distributional politics of price transmission, congestion allocation, and flexibility provision will keep deepening uneven exposure during stress periods even as it raises welfare on average.

7. Conclusions and policy implications

This article has argued that the political economy of Europe's energy transition is shifting from fuel dependence to network dependence, and has built the empirical machinery to make that shift visible in market data. As decarbonisation deepens, strategic exposure migrates from the fuel-import bill to the interior of the regional electricity system – to topology, capacity allocation, congestion management, and the politics of flexibility. The intuition is familiar from the engineering and market-design literatures. The translation is new: a framework the geopolitics of energy transition and the energy-security literature on vulnerability can read side by side, and a 41-zone, 84-month panel of public bidding-zone data on which to test it.

The contribution is to make that relocation measurable. Four primitives – network position, congestion, residual-load volatility, and flexibility – become estimable equations. A 41-zone public-data panel provides the testing ground, its identification logic and limits documented in Section 5 and Appendix B. The findings revise the framework's first-order priors as much as they confirm them.

The results centre on a single network channel, confirmed from several directions. Zones with more cross-zonal capacity rely less on net imports in ordinary months. With network position held fixed, higher renewable shares raise within-month price volatility – a result that withstands the article's strictest identification checks, and one the staggered DiD around the NordLink and Viking Link commissionings complements by confirming the import result quasi-experimentally. Then 2022 showed the channel's other face. Zones with more border capacity, and zones with higher pre-crisis renewable shares, ended the gas year with wider price gaps to their neighbours; EU-27 zones became about

+11.4 EUR/MWh more volatile than their non-EU European neighbours (Figure 5). Integration lowered import dependence in calm months and delivered the crisis efficiently in stressed ones, and the intermittency channel concentrates in the network-central half of the panel – centrality amplifies exposure rather than damping it.

Two predictions fail, and the failures are instructive. H2 (the congestion-utilisation channel) returns insignificant on its own once the price-spread component is absorbed: the naive congestion mechanism is not detectable in our panel. H4 (the flexibility moderation) is statistically zero in the linear specification, and the discretised version rises monotonically across quartiles, opposite to the prediction. The likely reason is measurement. Hydro-plus-nuclear flexibility is confounded with network centrality – the Nordic core, the Alpine system, and France are at once the high-flexibility and the network-central zones – so the linear $C \times F$ interaction does not isolate a flexibility-specific channel. Disentangling the two requires a higher-resolution flexibility variable – monthly reservoir levels, dispatchable-plant outage rates, balancing-market depth – that the present data architecture does not yet carry. The article's substantive contribution therefore lives in the H1-H3-H6 network triple and the H5 counter-hypothesis, not in the flexibility moderation the original argument placed alongside them.

The policy implication is therefore not a generic case for expanding or restricting interconnection. The more discriminating question concerns *which* interconnections, under *which* congestion-management rules, paired with *which* flexibility resources. The 2022 results make the trade-off concrete: the same integration that lowers net imports in normal months delivers price shocks with high fidelity into the most-connected zones in stress months. Moving a peripheral zone from the 25th to the 75th percentile of network position would have added 2022-window price-gap exposure roughly equal to the panel-mean price gap itself (Section 6). Congestion reform, market-coupling design, redispatch rules, storage deployment, and the political economy of transmission therefore sit at the core of Europe's contemporary energy order, not at its technical margins. Three extensions follow. The moderation test needs a sharper flexibility variable than the generation mix provides, and the analytical vocabulary should eventually reach redispatch volumes and explicit auction allocations. An hourly-resolution rebuild would let the stress-period import-reliance outcome the framework defines (Equation 5) be estimated rather than only specified. And the quasi-experimental design should generalise to the next round of HVDC commissionings – the Bay of Biscay interconnection, the

Celtic Interconnector, the Hansa PowerBridge – once enough post-commissioning months accumulate.

The broader lesson is that decarbonisation does not dissolve geopolitics; it relocates it. In Europe, one of the most consequential new locations is the cross-border electricity system – and the political stakes of how that system is built, allocated, and stress-tested are now as legible to public-data research as to public debate.

References

1. Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016.
2. Albert, M. J. (2022). The global politics of the renewable energy transition and the non-substitutability hypothesis: Towards a great transformation? Review of International Political Economy, 29(5):1766–1781.
3. Allen, F. and Gale, D. (2000). Financial contagion. Journal of Political Economy, 108(1):1–33.
4. Ang, B. W., Choong, W. L., and Ng, T. S. (2015). Energy security: Definitions, dimensions and indexes. Renewable and Sustainable Energy Reviews, 42:1077–1093.
5. Atalay, E. (2017). How important are sectoral shocks? American Economic Journal: Macroeconomics, 9(4):254–280.
6. Beyza, J., Gil, P., Masera, M., and Yusta, J. M. (2020). Security assessment of cross-border electricity interconnections. Reliability Engineering & System Safety, 201:106950.
7. Blondeel, M., Bradshaw, M. J., Bridge, G., and Kuzemko, C. (2021). The geopolitics of energy system transformation: A review. Geography Compass, 15(7):e12580.
8. Bridge, G., Bouzarovski, S., Bradshaw, M., and Eyre, N. (2013). Geographies of energy transition: Space, place and the low-carbon economy. Energy Policy, 53:331–340.
9. Cassetta, E., Nava, C. R., and Zoia, M. G. (2022). EU electricity market integration and cross-country convergence in residential and industrial end-user prices. Energy Policy, 165:112934.
10. Cherp, A. and Jewell, J. (2014). The concept of energy security: Beyond the four As. Energy Policy, 75:415–421.

11. Child, M., Kemfert, C., Bogdanov, D., and Breyer, C. (2019). Flexible electricity generation, grid exchange and storage for the transition to a 100% renewable energy system in Europe. Renewable Energy, 139:80–101.

12. Cludius, J., Hermann, H., Matthes, F. C., and Graichen, V. (2014). The merit order effect of wind and photovoltaic electricity generation in Germany 2008–2016: Estimation and distributional implications. Energy Economics, 44:302–313.

13. Colgan, J. D., Gard-Murray, A. S., and Hinthorn, T. (2023). Quantifying the value of energy security: How Russia's invasion of Ukraine exploded Europe's fossil fuel costs. Energy Research & Social Science, 103:103201.

14. Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. American Economic Review, 104(10):3115–3153.

15. Escribano, G., Gonzalez-Enriquez, C., Lazaro-Touza, L., and Paredes-Gazquez, J. (2023). An energy union without interconnections? Public acceptance of cross-border interconnectors in four European countries. Energy, 266:126385.

16. Fabra, N. (2023). Reforming European electricity markets: Lessons from the energy crisis. Energy Economics, 126:106963.

17. Finck, R. (2021). Impact of flow-based market coupling on the European electricity markets. Sustainability Management Forum, 29(2):173–186.

18. Goldthau, A. and Sitter, N. (2014). A liberal actor in a realist world? The Commission and the external dimension of the single market for energy. Journal of European Public Policy, 21(10):1452–1472.

19. Goldthau, A. and Sitter, N. (2015). A Liberal Actor in a Realist World: The European Union Regulatory State and the Global Political Economy of Energy. Oxford University Press.

20. Hawker, G., Bell, K., and Gill, S. (2017). Electricity security in the European Union – the conflict between national capacity mechanisms and the single market. Energy Research & Social Science, 24:51–58.

21. Hirth, L. (2013). The market value of variable renewables: The effect of solar wind power variability on their relative price. Energy Economics, 38:218–236.

22. Karahan, C. C., Odabasi, A., and Tiryaki, C. S. (2024). Wired together: Integration and efficiency in European electricity markets. Energy Economics, 133:107505.

23. Ketterer, J. C. (2014). The impact of wind power generation on the electricity price in Germany. Energy Economics, 44:270–280.

24. Knodt, M., Ringel, M., and Müller, R. (2020). Harder soft governance in the European Energy Union. Journal of Environmental Policy & Planning, 22(6):787–800.

25. Kuzemko, C., Blondeel, M., Dupont, C., and Brisbois, M.-C. (2022). Russia's war on Ukraine, European energy policy responses and implications for sustainable transformations. Energy Research & Social Science, 93:102842.

26. LaBelle, M. C. (2024). Breaking the era of energy interdependence in Europe: A multidimensional reframing of energy security, sovereignty, and solidarity. Energy Strategy Reviews, 52:101314.

27. Lang, L. M., Dallinger, B., and Lettner, G. (2020). The meaning of flow-based market coupling on redispatch measures in Austria. Energy Policy, 136:111061.

28. Meckling, J. (2019). Governing renewables: Policy feedback in a global energy transition. Environment and Planning C: Politics and Space, 37(2):317–338.

29. Mišík, M. and Nosko, A. (2023). Each one for themselves: Exploring the energy security paradox of the European Union. Energy Research & Social Science, 99:103074.

30. Moore, S. (2017). Evaluating the energy security of electricity interdependence: Perspectives from Morocco. Energy Research & Social Science, 24:21–29.

31. Newbery, D., Pollitt, M. G., Ritz, R. A., and Strielkowski, W. (2018). Market design for a high-renewables European electricity system. Renewable and Sustainable Energy Reviews, 91:695–707.

32. Newbery, D., Strbac, G., and Viehoff, I. (2016). The benefits of integrating European electricity markets. Energy Policy, 94:253–263.

33. Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. Journal of Business & Economic Statistics, 37(2):187–204.

34. Overland, I. (2019). The geopolitics of renewable energy: Debunking four emerging myths. Energy Research & Social Science, 49:36–40.

35. Poplavskaya, K., Totschnig, G., Leimgruber, F., Doorman, G., Etienne, G., and de Vries, L. (2020). Integration of day-ahead market and redispatch to increase cross-border exchanges in the European electricity market. Applied Energy, 278:115669.

36. Sáez, Y., Mochón, A., Corona, L., and Isasi, P. (2019). Integration in the European electricity market: A machine learning-based convergence analysis for the Central Western Europe region. Energy Policy, 132:549–566.

37. Schönheit, D., Kenis, M., Lorenz, L., Möst, D., Delarue, E., and Bruninx, K. (2021). Toward a fundamental understanding of flow-based market coupling for cross-border electricity trading. Advances in Applied Energy, 2:100027.

38. Schönheit, D., Weinhold, R., and Dierstein, C. (2020). The impact of different strategies for generation shift keys (GSKs) on the flow-based market coupling domain: A model-based analysis of Central Western Europe. Applied Energy, 258:114067.

39. Scholten, D., Bazilian, M., Overland, I., and Westphal, K. (2020). The geopolitics of renewables: New board, new game. Energy Policy, 138:111059.

40. Scholten, D. and Bosman, R. (2016). The geopolitics of renewables: Exploring the political implications of renewable energy systems. Technological Forecasting and Social Change, 103:273–283.

41. Sovacool, B. K. and Mukherjee, I. (2011). Conceptualizing and measuring energy security: A synthesized approach. Energy, 36(8):5343–5355.

42. Szulecki, K. (2016). European energy governance and decarbonization policy: Learning from the 2020 strategy. Climate Policy, 16(5):543–547.

43. Vakulchuk, R., Overland, I., and Scholten, D. (2020). Renewable energy and geopolitics: A review. Renewable and Sustainable Energy Reviews, 122:109547.

44. Van de Graaf, T., Overland, I., Scholten, D., and Westphal, K. (2020). The new oil? The geopolitics and international governance of hydrogen. Energy Research & Social Science, 70:101667.

Appendix

Table 7 lists every variable referenced in Section 4 and used in the Section 5 regressions, with its precise definition (including units) and primary source. Source codes are: ENTSO-E for the European Network of Transmission System Operators for Electricity Transparency Platform; ACER for the EU Agency for the Cooperation of Energy Regulators; PVGIS for the European Commission's Photovoltaic Geographical Information System (country-capital coordinates); JRC-IDEES for the JRC Integrated Database of the European Energy System; OPSD for Open Power System Data; and *Derived*, meaning the variable is built from one of these sources using the procedures described in Section 4 and documented in the replication archive.

Table 7. Variable definitions, units, and primary data sources.

Variable	*Definition (with unit)*	*Source*
	Panel identifiers	
Bidding zone	ENTSO-E bidding-zone code (e.g. FR, DE_LU, IT_NORD)	ENTSO-E
Country	ISO-2 country code (e.g. FR; DE for DE_LU; IT for the seven IT zones)	Derived
Month	First day of the calendar month; panel-time index for the regressions	Calendar
	Day-ahead prices (zone-month)	
Price, mean	Within-month mean of hourly day-ahead prices (EUR/MWh)	ENTSO-E (derived)
Price, std. dev.	Within-month standard deviation of hourly day-ahead prices (EUR/MWh)	ENTSO-E (derived)
Price, p_{10} and p_{90}	10th and 90th percentile of hourly day-ahead prices in the month (EUR/MWh)	ENTSO-E (derived)
Negative-price share	Share of hours in the month with day-ahead price below zero (0-1)	ENTSO-E (derived)
PD_{it}	Cross-border price divergence: mean absolute price gap to adjacent zones (EUR/MWh)	ENTSO-E (derived)
Max price gap	Maximum absolute price gap across a zone's adjacent neighbours (EUR/MWh)	ENTSO-E (derived)
	Load and residual load (zone-month)	
Load, mean	Within-month mean of hourly actual total load (MW)	ENTSO-E
Load, std. dev.	Within-month standard deviation of hourly load (MW)	ENTSO-E (derived)
Load, peak	Maximum hourly load in the month (MW)	ENTSO-E (derived)
Monthly consumption	Total monthly electricity consumption (MWh); integral of hourly load	ENTSO-E (derived)
Residual load, mean	Mean of $\mathrm{RL}_{i,\tau} = \mathrm{Load}_{i,\tau} - \mathrm{Wind}_{i,\tau} - \mathrm{Solar}_{i,\tau}$ within month (MW)	ENTSO-E (derived)
Residual-load volatility	Within-month standard deviation of $\mathrm{RL}_{i,\tau}$ (MW); used in Eq.~\eqref{eq:est-baseline}	ENTSO-E (derived)
Residual load, p_{90}	90th percentile of $\mathrm{RL}_{i,\tau}$ within the month, marking stress hours (MW)	ENTSO-E (derived)
	Generation by production type (zone-month)	
Wind generation	Total wind generation in the month (MWh)	ENTSO-E
Solar generation	Total solar-PV generation in the month (MWh)	ENTSO-E
Hydro generation	Total hydropower generation (run-of-river, reservoir, pumped) (MWh)	ENTSO-E
Nuclear generation	Total nuclear generation in the month (MWh)	ENTSO-E
Fossil generation	Total fossil-fuel generation (coal, gas, oil, lignite) (MWh)	ENTSO-E
Biomass generation	Total biomass generation in the month (MWh)	ENTSO-E
Total generation	Total generation across all production types in the month (MWh)	ENTSO-E (derived)
Generation share, s_{it}^{type}	Share of each production type: type-specific generation divided by total generation	Derived
Renewable share, r_{it}	Wind + solar + hydro shares of monthly generation	Derived
	Cross-border position (zone-month)	
Net position, mean	Within-month mean of zone net position (MW); positive denotes net export	ENTSO-E
Net position, std. dev.	Within-month standard deviation of zone net position (MW)	ENTSO-E (derived)
NIS_{it}	Net import share, Eq. (4): monthly mean of -(net position)/load	ENTSO-E (derived)
SPIR_{it}	Stress-period import reliance, Eq. (5): mean of NIS in top-decile RL or price hours	ENTSO-E (derived)
	Structural regressors (zone-month)	

Variable	*Definition (with unit)*	*Source*
Interconnector degree, d_i	Count of bidding zones adjacent to i (time-invariant per zone)	ENTSO-E topology (derived)
N_{it}	Capacity-weighted network position, Eq. (1): $\sum_{j\in\mathcal{A}(i)} K_{ij,t}$	ENTSO-E (derived)
C_{it}	Congestion composite, Eq. (2): price spread + utilisation pressure across borders	ENTSO-E (derived)
F_{it}	Flexibility composite, Eq. (3): hydro + nuclear shares of generation	Derived
	Weather controls (country-capital, monthly)	
2-metre temperature	Monthly mean 2-metre air temperature at the country capital ($°C$)	PVGIS
10-metre wind speed	Monthly mean 10-metre total wind speed at the country capital (m/s)	PVGIS
Global horizontal irradiance	Monthly mean global horizontal irradiance at the country capital (W/m^2)	PVGIS
	Long-run structural controls (country, annual)	
JRC-IDEES indicators	Sectoral energy use, generation mix, balancing-relevant infrastructure	JRC-IDEES 2021
	Validation series	
OPSD time series	Pre-computed hourly capacity factors and load/price series for European zones	Open Power System Data
	Indicator	
2022 shock, S_t	Indicator: 1 if calendar 2022, 0 otherwise; used in Eq. (12)	Calendar

B Identification Checks and Robustness Extensions

This appendix reports six identification and robustness checks addressing concerns from the audit of the main empirical results. Each check runs in the master do-file of the replication archive – whose verification block prints every statistic quoted here to the run log – and uses no data beyond that listed in Appendix A.

Pooled event study and pre-trend correction. Figure 7 reproduces the pooled event study without a comparison group. It sits in this appendix because its parallel-trends F-test is decisively rejected (F = 99.34). The DiD in Figure 5 of the main text targets the same identification with a sensible comparison group – the EU-27/non-EU European split – and a smaller, if not eliminated, pre-trends violation (F = 53.38). We also re-estimate the pooled event study after subtracting a linear pre-event trend: fit $\hat{g} \cdot k$ to the pre-anchor observations using zone fixed effects, project that slope onto the whole window, and run the event-study regression on the detrended outcome. The estimated pre-event slope is $\hat{g}$ = 1.99 EUR/MWh per month; price volatility was rising at about two EUR/MWh per month over 2020-2021 before the gas crisis hit. Figure 8 shows the resulting profile. The post-event hump sits somewhat tighter against zero in the months before 2022, but the joint pre-trends F on the *detrended* pre-event coefficients fails too (it rises to F = 646.8). The drift in price volatility from 2020 onward is not a simple linear trend that a one-degree polynomial can absorb, so a mechanical correction cannot rehabilitate the pooled event study. The DiD specification in the main text and the shock-window regression in Table 5 remain the defensible identifications of the 2022 differential.

Figure 7. Pooled event-study dynamics around the 2022 gas-crisis window (without a comparison group).

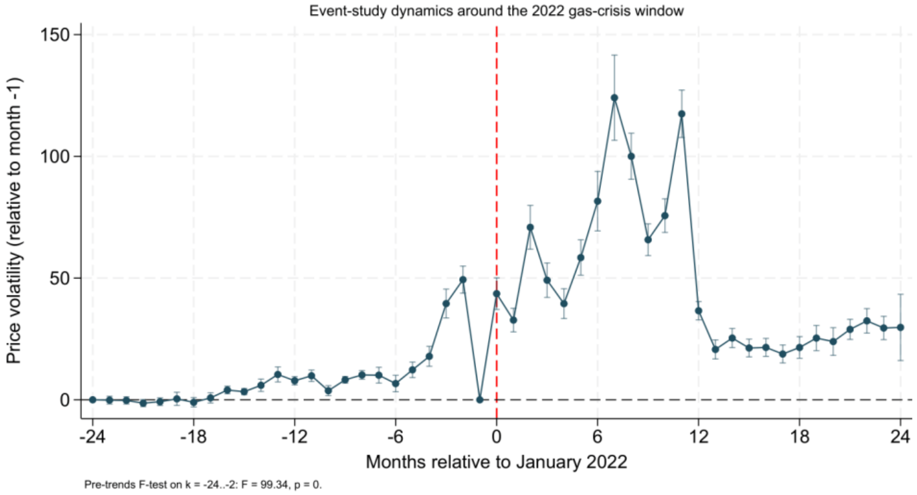

Notes: Coefficients $\hat{\mu}_k$ from Equation (13) on price volatility, with k=-1 (one month before the January 2022 anchor) as the reference period.

Figure 8. Event study after removing the linear pre-trend

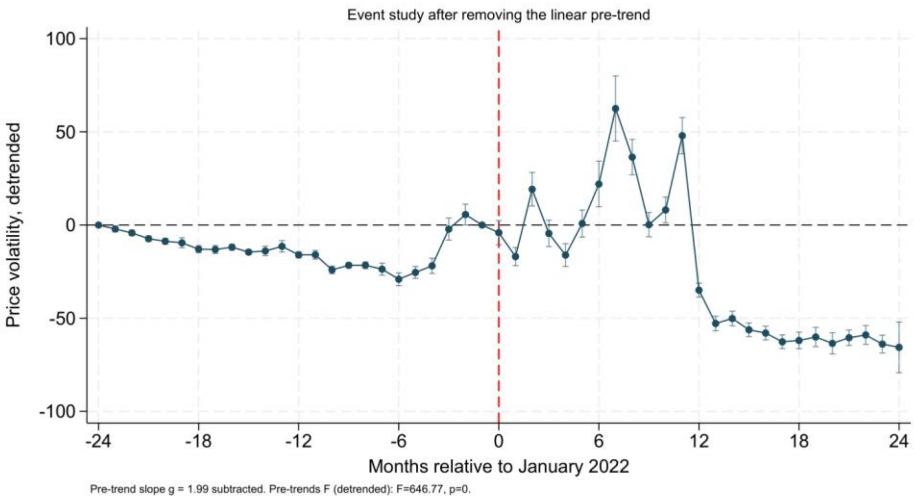

Notes: The pre-event slope $\hat{g}$ = 1.99 EUR/MWh per month is subtracted from the outcome over the full window before the event-study regression is re-estimated.

Placebo shock years. The shock-window result in Table 5 would be unconvincing if any arbitrary calendar year produced similar interaction coefficients. Table 8 re-estimates the shock-window specification for the cross-border price gap with the dummy year set to 2020, 2021, 2022 (the headline), and 2023. In the three placebo years the interaction coefficients are statistically zero; where their magnitudes are non-trivial, the standard errors are wide enough that the 95% confidence interval comfortably includes zero. Only in the 2022 column is either interaction significant (+43.4** for renewable share, +2.09* for network capacity). The 2022 result is not a generic artefact of the specification or sample but a feature of the gas-crisis year itself – which is what the framework's shock-window hypothesis predicts.

Table 8. Placebo shock years for the cross-border price-gap regression.

	2020 placebo	*2021 placebo*	*2022 (headline)*	*2023 placebo*
Renewable share	1.823	4.174	4.125	-3.231
	(14.980)	(15.833)	(13.859)	(12.503)
Network capacity (GW)	-0.900	-0.948	-1.033	-0.915
	(1.029)	(1.209)	(1.172)	(1.188)
Renewable share × shock year	14.364	-0.477	43.433**}	37.310
	(20.715)	(28.707)	(18.909)	(27.314)
Network capacity (GW) × shock year	-0.314	-1.055	2.092*}	1.938
	(0.870)	(1.338)	(1.111)	(1.386)
L. Load mean	0.000	0.000	-0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Temperature (deg C)	0.598	0.664	0.461	0.642
	(0.920)	(0.827)	(0.923)	(0.857)
Constant	17.799	16.556	14.408	10.841
	(10.712)	(10.387)	(11.484)	(8.313)
Observations	1,128	1,128	1,128	1,128
R^2 (within)	0.009	0.012	0.041	0.028

Notes: Each column re-estimates the shock-window specification of Eq. (12) on the cross-border price-gap outcome with the year-of-shock dummy set to the column header. Only the 2022 column returns statistically significant interaction terms; the placebo years (2020, 2021, 2023) return wider standard errors and non-significant interactions, confirming that the 2022 result is not a generic artefact of the specification. Zone and month-year fixed effects absorbed. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

Standardized coefficients. Table 9 re-estimates the baseline specification of Table 2 with all continuous outcomes and regressors standardised within sample, so each coefficient reads as the $\hat{\beta}$-SD change in Y associated with a 1-SD increase in X. On this scale, the strongest associations are wind speed on the negative-price share (+0.198**), temperature on net imports (+0.140**), and the lagged renewable share on the negative-price share (-0.125*). These magnitudes are economically modest – about a tenth to a fifth of one outcome standard deviation per one-SD regressor change – consistent with the framing in the main text: zone and month-year fixed effects absorb most of the variation in the vulnerability outcomes, and the structural regressors explain the residual conditional pattern only modestly.

Table 9. Standardized-coefficient version of the baseline regressions.

	Price vol. (z)	*Neg-price share (z)*	*Net import share (z)*	*Price gap (z)*
L. Renewable share (z)	-0.007	-0.125*	-0.014	-0.160
	(0.076)	(0.069)	(0.042)	(0.175)
L. Load mean (z)	0.029	0.077	-0.238	-0.010
	(0.146)	(0.350)	(0.258)	(0.368)
Temperature 2m (z)	-0.016	-0.094	0.140**	0.155
	(0.097)	(0.124)	(0.058)	(0.159)
Wind speed 10m (z)	-0.048	0.198**	0.063	-0.133
	(0.032)	(0.074)	(0.039)	(0.086)
Solar GHI (z)	0.081	0.073	-0.051	-0.021
	(0.079)	(0.081)	(0.050)	(0.121)
Constant	0.057*	-0.241***	-0.040	-0.020
	(0.032)	(0.051)	(0.026)	(0.073)
Observations	1,747	1,747	1,649	1,711
R^2 (within)	0.003	0.037	0.015	0.007

Notes: All continuous outcomes and regressors are z-scored within sample so each coefficient reads as the number of outcome standard deviations associated with a one-SD increase in the regressor. Zone and month-year fixed effects absorbed; fixed effects are not standardised. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

Quartile flexibility interaction. The H4 moderation in Table 4 ($\hat{\beta}_4$ = +0.007) is statistically zero, but a linear interaction can mask a non-linear pattern. Table 10 replaces the linear interaction $C \times F$ with quartile-discretised flexibility dummies interacted with residual-load volatility (Eq. 16). On the price-volatility outcome, the residual-load coefficient is essentially zero in the lowest flexibility quartile (Q1), and the differential moderation rises monotonically to +0.004* in Q3 and +0.010* in Q4. That reverses the simple flexibility-dampens-volatility prediction: in our panel, the zones with the highest hydro-plus-nuclear shares – also the network-central zones – appear to transmit residual-load volatility into price volatility more, not less. As the Conclusion argues, this measurement confounds flexibility with network centrality. We read the quartile result as an exploratory null and treat H4 as rejected in this design.

Table 10. Quartile-flexibility moderation of residual-load volatility.

	Price vol.	*Neg-price share*	*Net import share*
Residual-load volatility	-0.000	0.000***	0.000
	(0.002)	(0.000)	(0.000)
Flexibility quartile=2 × Residual-load volatility	0.001	0.000	-0.000
	(0.002)	(0.000)	(0.000)
Flexibility quartile=3 × Residual-load volatility	0.004*	-0.000	-0.000
	(0.002)	(0.000)	(0.000)
Flexibility quartile=4 × Residual-load volatility	0.010*	0.000	-0.000
	(0.005)	(0.000)	(0.000)
L. Renewable share	10.777	-0.003	-0.135
	(7.161)	(0.006)	(0.111)
L. Load mean	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)
Temperature (deg C)	-0.348	-0.000*	0.011**
	(0.430)	(0.000)	(0.004)
Constant	41.607***	0.009	-0.079
	(6.920)	(0.006)	(0.128)
Observations	1,469	1,469	1,430
R^2 (within)	0.017	0.025	0.054

Notes: Replaces the linear $C \times F$ interaction of Table 4 with quartile-discretised flexibility dummies interacted with residual-load volatility C (Eq. 16). Quartile 1 (lowest flexibility) is the omitted reference, so the reported coefficients on $C \times Q_k$ for $k \in \{2,3,4\}$ are differential moderation effects relative to Q1. Zone and month-year fixed effects absorbed. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

Standard-error robustness. Table 11 re-estimates the headline price-volatility baseline with three standard-error estimators: cluster-robust at the zone level (the manuscript's default), cluster-robust at the country level (the country-month fallback's choice), and Driscoll-Kraay, robust to both within-zone serial correlation and contemporaneous cross-zone correlation. The same OLS regression underlies every column, so point estimates are identical; only the standard errors differ. The conclusion that the baseline coefficients are mostly imprecise survives all three. Country-level clustering gives moderately larger SEs (e.g. on renewable share, 10.33 vs 7.87 at the zone level), as expected with fewer clusters; Driscoll-Kraay gives smaller SEs (3.58 on renewable share) because it deflates less aggressively for residual within-zone autocorrelation. None of them changes which coefficients are significant at conventional thresholds.

Table 11. Standard-error robustness for the headline price-volatility baseline

	Cluster (zone)	*Cluster (country)*	*Driscoll-Kraay*
L. Renewable share	-0.733	-0.733	-0.733
	(7.868)	(10.330)	(3.577)
L. Load mean	0.000	0.000	0.000
	(0.000)	(0.000)	(0.001)
Temperature (deg C)	-0.068	-0.068	-0.068
	(0.401)	(0.458)	(0.509)
Wind speed at 10m (m/s)	-1.582	-1.582	-1.582
	(1.067)	(1.075)	(1.237)
Solar GHI (W/m^2)	0.029	0.029	0.029
	(0.028)	(0.030)	(0.037)
Observations	1,747	1,747	1,747

Notes: Three SE estimators applied to the identical OLS regression of price volatility on the baseline regressors. Col. 1: cluster-robust at the zone level. Col. 2: cluster-robust at the country level (the country-month fallback's choice). Col. 3: Driscoll-Kraay standard errors, robust to both serial correlation within zones and contemporaneous cross-zone correlation; this estimator supports one-way absorption only, so the month-year structure is captured by explicit time dummies rather than by absorbed fixed effects. Point estimates are identical by construction; only the SEs differ. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.

Heterogeneity by network centrality. The framework predicts that the structural relationships differ between network-central and network-peripheral zones. We split the sample at the within-panel median of capacity-weighted network position N_{it} from Eq. (1) and re-estimate the network-mechanism regression on each half. Table 12 reports two outcomes: price volatility (cols. 1-2) and net import share (cols. 3-4). Two patterns stand out. First, the lagged renewable share enters strongly positively on price volatility in both halves but is roughly 50% larger in the central subsample (+20.8***) than in the peripheral one (+14.0*), consistent with the reading that network centrality *amplifies* the renewable-intermittency channel rather than damping it. Second, the within-zone R^2 for net imports is much higher in the peripheral subsample (0.084 vs 0.035): the structural regressors explain a larger share of within-zone variation there, a pattern compatible with peripheral zones being more constrained by the framework's structural variables and central zones showing more idiosyncratic month-to-month variation.

Table 12. Network-centrality heterogeneity in the network-mechanism regression

	Pvol Central	Pvol Peripheral	NIS Central	NIS Peripheral
L. Network capacity (GW)	0.246	-0.166	-0.010	0.001
	(0.268)	(0.915)	(0.010)	(0.011)
L. Border utilisation	12.381	4.855	0.048	0.120
	(9.305)	(26.622)	(0.198)	(0.080)
L. Renewable share	20.800***	13.970*	0.135	-0.223
	(6.707)	(7.111)	(0.267)	(0.135)
L. Load mean	-0.000	0.001	0.000	-0.000**
	(0.000)	(0.003)	(0.000)	(0.000)
Temperature (deg C)	-0.716	0.444	0.005	0.017**
	(0.437)	(0.474)	(0.004)	(0.006)
Wind speed at 10m (m/s)	-0.813	-0.843	0.028	0.063
	(0.966)	(1.248)	(0.022)	(0.044)
Constant	34.198***	30.958	-0.128	0.009
	(10.306)	(18.200)	(0.255)	(0.249)
Observations	515	649	507	565
R^2 (within)	0.039	0.014	0.035	0.084

Notes: Sample split at the within-panel median of lagged capacity-weighted network position $L.N_{it}$ from Eq. (1); the central subsample contains zone-months above the median, the peripheral subsample those at or below. Cols. 1-2 report the network-mechanism regression of Table 3 for price volatility in the two subsamples; Cols. 3-4 repeat for net-import share. Zone and month-year fixed effects absorbed within each subsample. Cluster-robust SEs at the zone level. *, **, and *** represent statistical significance at levels of 10%, 5% and 1%, respectively.