Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Renewable Energy Expansion: The Role of Global Supply-Chain Vulnerability in the Energy Transition

Jamil Ahmed Channa

jamilchanna@gmail.com

Shah Abdul Latif University

Professor, Dr. Naveed Ahmed Shaikh

Shah Abdul Latif University

Dr. Rehana Naheed

Quaid-i-Azam University

Dr. Abid Iqbal

Iqra University

Dr. Majid Phul

Shah Abdul Latif University

Research Article

Keywords: monetary transmission; wind energy; supplier concentration; Herfindahl–Hirschman index; supply-chain resilience; panel data; local projections

Posted Date: November 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11674460/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Environmental Economics and Policy Studies

Monetary Tightening, Wind-Equipment Import Concentration, and Renewable Energy Expansion: The Role of Global Supply-Chain Vulnerability in the Energy Transition

--Manuscript Draft--

Manuscript Number:	
Full Title:	Monetary Tightening, Wind-Equipment Import Concentration, and Renewable Energy Expansion: The Role of Global Supply-Chain Vulnerability in the Energy Transition
Article Type:	Research Article
Corresponding Author:	Jamil Ahmed Channa, Ph.D Shah Abdul Latif University Khairpur, Sindh PAKISTAN
Corresponding Author Secondary Information:	
Corresponding Author's Institution:	Shah Abdul Latif University
Corresponding Author's Secondary Institution:	
First Author:	Jamil Ahmed Channa
First Author Secondary Information:	
Order of Authors:	Jamil Ahmed Channa
	Naveed Ahmed Shaikh, Professor
	Rehana Naheed, Assistant Professor
	Abid Iqbal, Assistant Professor
	Majid Hussain Phul, Assistant Professor
Order of Authors Secondary Information:	
Funding Information:	
Abstract:	Highlights • A monetary–supply-chain framework for wind-capacity expansion is proposed. • Wind-turbine supplier HHI is built from reconciled bilateral Comtrade data. • Monetary transmission is conditioned on equipment-sourcing concentration. • Wild-cluster bootstrap inference is used for a few-country panel. • Pilot estimates are imprecise; a shock-based full-sample design is specified.
Suggested Reviewers:	

Powered by Editorial Manager® and ProduXion Manager® from Aries Systems Corporation

Title Page (will not go to reviewers)

Monetary Tightening, Wind-Equipment Import Concentration, and Renewable Energy Expansion: The Role of Global Supply-Chain Vulnerability in the Energy Transition

Naveed Ahmed Sheikh[a], Jamil Ahmed Channa[b, *], Rehana Naheed[c],

Abid Iqbal[d], Majid Hussain Phul [e],

[a] *Vice Chancellor, The Shaikh Ayaz University, Shikarpur, Pakistan*

[b] *Department of Economics, Shah Abdul Latif University, Khairpur, Pakistan*

[c] *School of Management Sciences, Quaid e Azam University, Islamabad, Pakistan*

[d] *Department of Business Administration, Iqra University, Karachi, Pakistan*

[e] *Department of Economics, Shah Abdul Latif University, Khairpur, Pakistan*

** Corresponding author. E-mail address: jamilchanna@gmil.com*

Highlights

- A monetary–supply-chain framework for wind-capacity expansion is proposed.
- Wind-turbine supplier HHI is built from reconciled bilateral Comtrade data.
- Monetary transmission is conditioned on equipment-sourcing concentration.
- Wild-cluster bootstrap inference is used for a few-country panel.
- Pilot estimates are imprecise; a shock-based full-sample design is specified.

Abstract

Renewable infrastructure is highly sensitive to the cost of capital, and its equipment comes from a concentrated set of international suppliers. This paper asks whether these exposures interact: does monetary tightening slow wind-capacity expansion more where wind-equipment imports come from few suppliers? Prior work shows that interest rates raise the levelised cost of renewables and that clean-technology supply chains are concentrated, but no study conditions monetary transmission to renewable investment on a country's sourcing structure. We propose a Monetary–Supply-Chain Transmission Mechanism (MSTM) in which supplier concentration governs bargaining power, lead times and exchange-rate pass-through, and so amplifies or buffers the financing effect of higher policy rates. We construct a Herfindahl–Hirschman index of wind-turbine (HS 850231) import suppliers from reconciled bilateral UN Comtrade data, merge it with IRENA capacity and BIS policy rates for 20 non-euro economies over 2012–2023, and estimate two-way fixed-effects interaction models, marginal effects, local projections and a leakage-safe machine-learning block with wild-cluster bootstrap inference. In the pilot sample currently available (7 countries, 60 country-years), the interaction is positive but imprecise (4.9 pp; bootstrap $p = 0.23$), changes sign under country exclusion, and no learner beats a mean forecast; these results describe the sample's limits, not the mechanism. The contribution is the framework, the measurement design and a reproducible pipeline that specifies the data required for credible inference. The policy relevance lies in treating supplier diversification as a hedge against financing-cycle risk.

Keywords: monetary transmission; wind energy; supplier concentration; Herfindahl–Hirschman index; supply-chain resilience; panel data; local projections

JEL classification: E52, F14, Q42, Q48, C23

1. Introduction

The 2022–2023 global tightening cycle arrived at the moment when wind and solar deployment needed to accelerate most, and it exposed a structural feature of renewable infrastructure that fossil generation does not share: almost all of its lifetime cost is paid up front. For an onshore wind project, capital expenditure and financing typically account for well over three-quarters of levelised cost, so a given rise in the cost of capital raises the levelised cost of wind electricity by far more than that of a gas plant (Hirth & Steckel, 2016; Schmidt et al., 2019). Egli, Steffen and Schmidt (2018) show that falling financing costs explained a substantial share of the decline in renewable costs in Germany over 2000–2017, and Serebriakova, Polzin and Sanders (2026) find, across 28 European countries over 2001–2024, that a 25-basis-point rise in the policy rate is associated with a 3.2 per cent fall in onshore-wind installed capacity while fossil, hydro and nuclear capacity are unaffected. Monetary policy, in short, has a technology-specific footprint.

A second structural feature has received separate attention. Wind-turbine manufacturing and trade are geographically concentrated. The International Energy Agency's Energy Technology Perspectives (IEA, 2023) documents that China dominates manufacturing in most clean-energy supply-chain segments, and Eurostat reports that in 2025 China supplied roughly 60 per cent of EU wind-turbine imports. For an importing country, the relevant exposure is not merely how much equipment it imports but how many suppliers it imports from. A country sourcing turbines from a single exporter faces supplier bargaining power over price and delivery, longer and less substitutable lead times, and exchange-rate pass-through concentrated in one currency pair. These are the conditions under which the trade literature expects shocks to propagate rather than dissipate (Baldwin & Freeman, 2022; Grossman, Helpman, & Lhuillier, 2023).

The two features have not been studied together. The monetary–renewables literature treats financing conditions as the whole story and holds equipment supply implicit; the supply-chain literature treats concentration as a risk to volumes and prices but not as a modifier of how financial shocks reach investment decisions. Yet the economics of irreversible investment implies they should interact. When a developer faces both a higher hurdle rate and an equipment market in which one supplier sets price and delivery, the option value of waiting rises on two margins at once (Dixit & Pindyck, 1994; Bloom, 2009). Conversely, a developer with several substitutable suppliers can renegotiate price, switch currency exposure or re-phase delivery to offset part of the financing shock. The conditional response of capacity expansion to monetary tightening should therefore depend on supplier concentration. That conditional response is the object of this paper.

We make three contributions. First, we propose and formalise a Monetary–Supply-Chain Transmission Mechanism (MSTM): an analytical framework, developed here rather than drawn from an established theory, in which the financing channel of monetary policy and the supply-chain channel of equipment sourcing enter the developer's investment decision jointly, and in which supplier concentration governs the sign and size of the interaction. The framework distinguishes explicitly between supplier concentration (how many sources, measured by a Herfindahl–Hirschman index of import shares) and import dependence (how much is imported relative to domestic production), because the two are often conflated and only the first is measured here.

Second, we construct the measurement design the question requires. Supplier concentration is computed from bilateral UN Comtrade import values for HS 850231 (wind-powered generating sets), with partner sums reconciled to reported world totals and no imputation of missing years. Wind capacity comes from IRENA, summing onshore and offshore, on-grid and off-grid components; monetary conditions from BIS monthly policy rates averaged within

calendar years, with partial years excluded. The sample is 20 advanced and emerging economies with independent national monetary policy over 2012–2023, euro-area members excluded so that no common policy rate is treated as 20 separate shocks.

Third, we build and release an estimation framework calibrated to what a panel of 20 countries and 12 years can support: two-way fixed effects with lagged, centred concentration and marginal effects reported across the observed distribution; restricted wild-cluster bootstrap inference rather than conventional clustered standard errors; a parsimonious dynamic specification with explicit Nickell-bias discussion and no System GMM; local projections that switch from an association design to an identified-shock design once a high-frequency monetary-shock series is merged; and a machine-learning block that uses leave-one-country-out and chronological validation so that predictive claims cannot be inflated by leakage.

We are explicit about where the paper stands. The trade data available at the time of writing cover supplier concentration for seven of the twenty countries, because the anonymous Comtrade download quota was exhausted before the panel was complete. We report the pilot estimates in full, but they are diagnostics of the sample, not evidence on the mechanism: with seven clusters and sixty observations, no coefficient in Sections 6–8 is interpretable as a population quantity, and we say so wherever a number appears. The paper's value at this stage lies in the framework, the measurement and the reproducible machinery that make the question answerable once the remaining data—which exist and are freely obtainable—are merged.

The remainder of the paper proceeds as follows. Section 2 reviews seven strands of literature and identifies the gap. Section 3 develops the MSTM and four hypotheses. Section 4 describes data construction and the audit. Section 5 sets out the econometric framework and its inference. Sections 6–8 present pilot results, the machine-learning block and sensitivity analysis. Section 9 discusses policy, and Section 10 concludes.

2. Literature review

2.1 Monetary transmission and renewable energy investment

The macro-monetary literature has long established that contractionary policy depresses investment, and that the effect is heterogeneous across firms and sectors according to balance-sheet strength and external-finance dependence (Ottonello & Winberry, 2020; Cloyne, Ferreira, Froemel, & Surico, 2023). Applications to renewable energy are more recent. Chen and Lin (2024) estimate a structural VAR on US quarterly data from 1990 to 2023 and find that monetary tightening reduces renewable production, with solar most responsive. Sun, Khan, Liu and Lei (2022) and Hashmi, Syed and Inglesi-Lotz (2022) report comparable associations for G7 and US data respectively. The most direct antecedent of the present paper is Serebriakova, Polzin and Sanders (2026), whose fixed-effects panel of ten utility-scale technologies across 28 European countries finds that onshore wind is among the technologies most sensitive to policy-rate increases, with interaction terms on investment intensity and external-finance dependence pointing to interest-rate and balance-sheet channels. Mbassi, Samba and Elomo Zogo (2025) widen the lens to energy consumption under inflation targeting. Lucchetta (2025) frames the trade-off between price stability and renewables formally. None of these studies conditions the monetary response on the structure of equipment sourcing; the supply side of the investment is implicit.

2.2 Financing constraints in capital-intensive energy infrastructure

The cost-of-capital literature explains why the monetary footprint is technology-specific. Hirth and Steckel (2016) show analytically and numerically that the levelised cost of capital-intensive low-carbon generation rises far more steeply with the discount rate than that of fuel-intensive generation, so higher capital costs tilt the merit order towards fossil plants. Egli, Steffen and Schmidt (2018) decompose the fall in German wind and solar costs and attribute a

large share to improved financing conditions and financial-sector learning. Steffen (2020) assembles cost-of-capital estimates for renewable projects and documents wide cross-country dispersion, which matters for any panel design because the same policy-rate change maps into different hurdle-rate changes across countries. Schmidt et al. (2019) then show that a return of interest rates to pre-2008 levels could more than offset expected technology-cost declines for European wind. These studies motivate H1, but they treat the equipment price as exogenous to financing conditions.

2.3 Trade concentration and supply-chain resilience

The international-trade literature after 2020 reframed concentration as a source of systemic risk. Baldwin and Freeman (2022) survey what is known about shock propagation through global value chains and emphasise that exposure depends on the number and substitutability of sources, not on import volumes alone. Grossman, Helpman and Lhuillier (2023) formalise the policy question: whether governments should promote international diversification or reshoring, showing that private sourcing decisions under-diversify relative to the social optimum when disruptions are correlated. These models imply that the *structure* of sourcing changes how a given shock transmits to downstream investment, which is exactly the conditioning role the MSTM assigns to supplier concentration. The literature is, however, almost entirely about supply disruptions and price shocks; financial shocks interacting with sourcing structure are not studied.

2.4 Renewable-technology imports

Eschmann and Jochem (2025) estimate a structural gravity model of wind and solar PV trade and find that rapid renewable-capacity growth produced only a minimal increase in wind-technology import dependence, so expansion need not mechanically raise reliance on imports. The IEA's Energy Technology Perspectives (2023) provides the stylised facts: production of

most clean-energy technologies is concentrated in China, but wind components—towers, blades, nacelles—are heavy and nonetheless widely traded. This matters for measurement. HS 850231 captures assembled wind-powered generating sets; a country can be diversified in turbine sets and concentrated in a component, or the reverse. We therefore interpret our concentration index as turbine-set supplier concentration and say so.

2.5 Geopolitical and trade vulnerability

Ulug and Andrei (2025) study the energy transition under simultaneous geopolitical and macro-financial shocks and find that the two reinforce each other, slowing transition indicators more than either alone. Cai and Wu (2021) document time-varying interactions between geopolitical risk and renewable consumption. This "twin-shock" framing is the closest conceptual neighbour to our interaction hypothesis, but it is specified at the level of aggregate risk indices rather than the country's actual supplier structure.

2.6 Monetary conditions and green investment

A growing literature asks whether central banks have a carbon bias. Serebriakova et al. (2026) interpret their findings as evidence of an unintended bias in contractionary operations; Chen and Lin (2024) show counterfactually that monetary policy can offset or amplify other structural shocks to renewables. We take no position on the normative question but note that any carbon bias would be larger in countries whose renewable equipment is concentrated in few suppliers if H3 holds, which gives the interaction a direct bearing on that debate.

2.7 Supplier diversification and infrastructure development

Evidence on diversification as a resilience strategy for infrastructure is thin and mostly descriptive. The policy literature (IEA, 2023) recommends diversification without quantifying its protective value against financial shocks. H4 is designed to supply that quantification.

2.8 The gap and the positioning of this study

The gap is at the intersection: no study measures a country's wind-equipment supplier concentration from bilateral trade and asks whether it conditions the response of capacity expansion to monetary tightening. The nearest studies either hold supply implicit (2.1, 2.2), hold finance implicit (2.3, 2.4) or use aggregate risk indices in place of sourcing structure (2.5). Table 1 summarises.

Table 1 Literature-gap matrix

Study	Outcome	Monetary variable	Supply-side variable	Interaction	Sample	Identification	Small-G inference
Serebriakova, Polzin & Sanders (2026)	Installed capacity, 10 technologies	ECB and national policy rates	Investment intensity, external-finance dependence	Rate × technology trait	28 European countries, 2001–2024	FE panel; 2SLS in LCOE specification	Not addressed
Chen & Lin (2024)	Renewable production	Identified monetary shock (SVAR)	None	None	US, 1990Q1–2023Q2	Structural VAR	Not applicable
Schmidt et al. (2019)	LCOE scenarios	Interest-rate scenarios	None	None	Germany/Europe, scenario	Simulation	Not applicable
Eschmann & Jochem (2025)	Wind/PV technology imports	None	Import dependence (gravity)	None	Global bilateral trade	Structural gravity	Not addressed
Grossman, Helpman & Lhuillier (2023)	Sourcing choice (theory)	None	Supplier diversification	Disruption × sourcing	Theoretical	Model	Not applicable
Ulug & Andrei (2025)	Transition indicators	Macro-financial risk index	Geopolitical risk index	Risk × risk	Cross-country	Panel	Not addressed
This study	Wind-capacity expansion (Δln, net MW)	Policy-rate change; identified HF shocks (planned)	Supplier HHI from bilateral HS 850231 imports	Rate change × lagged HHI, marginal effects	20 non-euro economies, 2012–2023	TWFE; local projections; shock design once merged	Wild-cluster bootstrap, t(G−1), LOCO

3. Theoretical framework and hypotheses

3.1 The Monetary–Supply-Chain Transmission Mechanism

We present the MSTM as an analytical framework developed for this study. Consider a developer in country *i* deciding whether to commission wind capacity in year *t*. The project is irreversible, so the developer invests when the expected value of proceeding exceeds the value of waiting (Dixit & Pindyck, 1994). Write the decision as

$$V_{it} = \frac{E[R_{it}]}{r_{it} + \delta} - P_{it}\left(H_{i,t-1}, e_{it}, \tau_{it}\right) - C_{it} \qquad (1)$$

where V is the net value of investing now, R the stream of revenue, r the project cost of capital, δ the depreciation rate, P the delivered equipment price, C installation and grid costs, H the supplier concentration of the country's wind-equipment imports, e the exchange rate against the supplier's currency and τ delivery delay. Monetary tightening raises r through the financing channel. Supplier concentration enters through P: with few suppliers, the importer has less bargaining power, pays a higher markup, bears exchange-rate pass-through concentrated in one currency, and faces longer and less substitutable lead times (Baldwin & Freeman, 2022; Grossman et al., 2023).

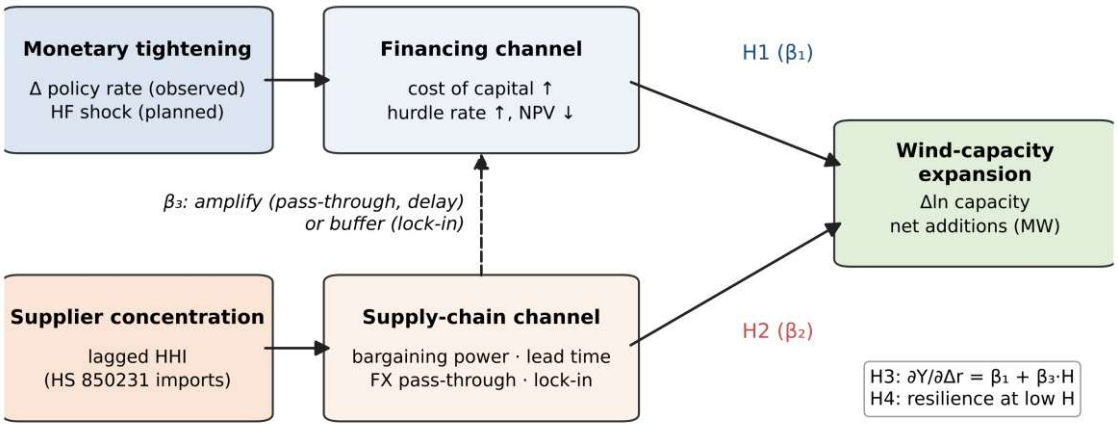


Fig. 1. Conceptual framework of the Monetary–Supply-Chain Transmission Mechanism (MSTM). Solid arrows are channels whose reduced form is estimable with the available variables; the dashed arrow is the conditioning effect β_3 tested under H3–H4.

The interaction arises because the two channels operate on the same irreversible decision. Differentiating the investment trigger with respect to the policy rate and to concentration gives the cross-partial that the paper estimates:

$$\frac{\partial^2 V_{it}}{\partial r_{it}\ \partial H_{i,t-1}} = \frac{\partial^2 V}{\partial r\ \partial P} \cdot \frac{\partial P}{\partial H} \qquad (2)$$

Three channels determine its sign. (i) *Cost-pass-through*: a concentrated supplier can raise prices when importers' financing costs rise, because importers have no outside option; the shock

is passed through twice and the cross-partial is negative (amplification). (ii) *Delay*: concentrated sourcing lengthens lead times, so a rate rise during the order window increases the carrying cost of committed capital and raises the option value of waiting; again amplification. (iii) *Lock-in*: a developer tied to a single supplier by framework contracts may be unable to postpone, so expansion continues through tightening and the cross-partial is positive (buffering). Which dominates is an empirical question, and Section 3.3 states both outcomes.

We stress the distinction that gives the framework its content. Supplier concentration H measures how many sources a country buys from; import dependence measures how much it buys relative to what it makes. China illustrates the difference: it is the world's largest wind manufacturer yet, in our data, its imports of HS 850231 come from few suppliers, so its HHI is high while its import dependence is negligible. The MSTM is about H, and the empirical design measures only H; claims about dependence are not made.

3.2 Channels that are and are not testable with available variables

Channel	Mechanism	Testable here	Variable
Financing	Higher policy rate raises project cost of capital	Yes (reduced form)	Annual policy-rate change; identified HF shock when merged
Supplier bargaining power	Concentration raises markup and reduces price negotiability	Partly: through the interaction, not directly	Lagged HHI; largest supplier share
Exchange-rate pass-through	Concentration narrows currency exposure to one pair	No: bilateral exchange rates not in panel	Flagged for extension
Lead time and delay	Concentration lengthens and rigidifies delivery	No: order-to-commissioning data absent	Flagged for extension
Investment irreversibility	Option value of waiting rises with both shocks	Indirectly: persistence and local-projection dynamics	Lagged dependent variable; horizons h = 0–3
Import dependence	Volume of imports relative to domestic production	No: not measured, by design	Not claimed

3.3 Hypotheses

H1 (monetary tightening). An increase in the policy rate is associated with slower wind-capacity expansion. *Mechanism*: the financing channel raises the hurdle rate on a capital-intensive, irreversible asset (Hirth & Steckel, 2016; Schmidt et al., 2019). *Variables*: policy-rate change (Δr) on capacity growth. *Expected sign*: $\beta_1 < 0$ at mean concentration. *Test*: TWFE coefficient on Δr with concentration centred at its sample mean, bootstrap inference. *Alternative*

outcome: $\beta_1 \approx 0$ if auction-based support schemes insulate project returns, or $\beta_1 > 0$ if tightening coincides with inflationary booms in which capacity pipelines were already committed.

H2 (supplier concentration). Higher wind-equipment supplier concentration is associated with slower or more volatile capacity expansion. *Mechanism*: reduced bargaining power and exposure to single-source disruption (Baldwin & Freeman, 2022). *Variables*: lagged HHI on capacity growth. *Expected sign*: $\beta_2 < 0$. *Test*: TWFE coefficient on lagged centred HHI. *Alternative outcome*: $\beta_2 > 0$ if concentration reflects a cost-minimising choice of the cheapest supplier, in which case concentration and expansion rise together (Eschmann & Jochem, 2025).

H3 (monetary–supply-chain interaction). The negative association between tightening and expansion is larger in magnitude where supplier concentration is higher. *Mechanism*: cost-pass-through and delay channels compound the financing shock. *Variables*: $\Delta r \times$ lagged centred HHI. *Expected sign*: $\beta_3 < 0$, so that $\partial Y/\partial \Delta r = \beta_1 + \beta_3 \cdot HHI$ becomes more negative as HHI rises. *Test*: interaction coefficient and the marginal effect evaluated at the 10th, 25th, 50th, 75th and 90th percentiles of lagged HHI with t(G−1) confidence bands; the hypothesis is supported only if the marginal effect is negative and distinguishable from zero in the upper part of the distribution. *Alternative outcome*: $\beta_3 > 0$ (buffering) under the lock-in channel.

H4 (supplier diversification). Countries with more diversified suppliers show greater resilience of capacity expansion during tightening. *Mechanism*: outside options allow price renegotiation and re-phasing. *Variables*: $\Delta r \times (1 - HHI)$. *Expected sign*: positive. *Test*: H4 is the mirror image of H3 under a linear interaction and is tested as a distinct proposition only through the threshold and quadratic specifications, which ask whether resilience is present above a diversification level rather than rising linearly. *Alternative outcome*: resilience may require a minimum number of suppliers rather than a lower HHI, in which case the supplier count, not the HHI, is the relevant variable.

The four hypotheses are testable propositions; none is an established finding, and the pilot results in Section 6 do not confirm any of them.

4. Data and variable construction

4.1 Sources and sample

The panel covers 20 advanced and emerging economies over 2012–2023 (240 potential country-years): Australia, Brazil, Canada, Chile, China, Czechia, Hungary, India, Indonesia, Japan, Korea, Mexico, New Zealand, Norway, Poland, South Africa, Sweden, Thailand, the United Kingdom and the United States. Selection is purposive: each country sets its own policy rate and has material wind deployment or wind-equipment trade. Fig. 2 maps the coverage of every variable by country and year. Euro-area members are excluded because a single ECB rate cannot be treated as independent national tightening in a country-fixed-effects design. The sample is therefore not globally representative, and we do not generalise beyond it.

Table 2 Variable definitions and sources

Variable	Definition	Unit	Source
Wind capacity	Onshore + offshore, on-grid + off-grid installed capacity, all four components required	MW	IRENA, archived CSV (Zenodo 10952917, April 2024 vintage)
Net capacity additions	Capacity(t) − capacity(t−1); includes retirements and revisions	MW	Constructed
Capacity growth	100 × ln(capacity(t)/capacity(t−1))	%	Constructed
Supplier HHI	Σ (supplier share)², supplier share = partner import value / Σ partner values, HS 850231 imports, partner2 = 0, positive partners only	0–1	UN Comtrade bilateral annual
Largest supplier share	Max partner value / Σ partner values	0–1	Constructed
Policy rate	Mean of 12 monthly central-bank policy rates in the calendar year	%	BIS policy-rate statistics
Policy-rate change	Annual mean(t) − annual mean(t−1)	pp	Constructed
GDP growth	Real GDP growth	%	WDI NY.GDP.MKTP.KD.ZG
Real GDP	Constant 2015 USD, logged	ln USD	WDI NY.GDP.MKTP.KD
Inflation	CPI inflation	%	WDI FP.CPI.TOTL.ZG
Trade openness	(Exports + imports)/GDP	%	WDI NE.TRD.GNFS.ZS
Population	Total, logged	ln persons	WDI SP.POP.TOTL
CO_2 emissions	Excluding LULUCF, logged	ln $MtCO_2e$	WDI EN.GHG.CO2.MT.CE.AR5

Variable	Definition	Unit	Source
Renewable share	Renewable energy in final energy consumption (robustness; ends 2021)	%	WDI EG.FEC.RNEW.ZS

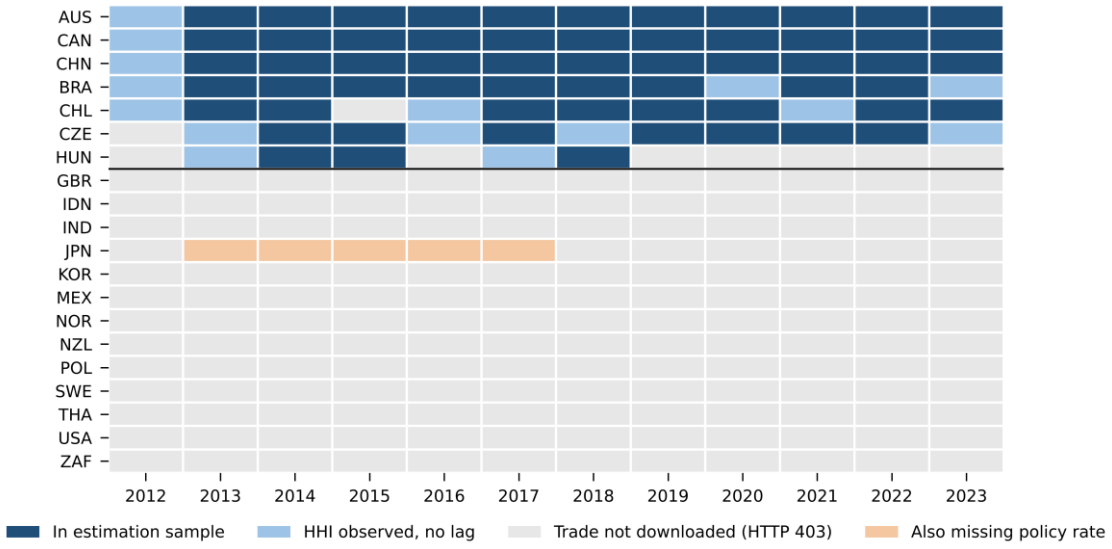


Fig. 2. Data coverage by country and year. Dark cells enter the lagged-interaction estimation sample (60 country-years, 7 countries); grey cells lack bilateral HS 850231 trade because the anonymous UN Comtrade quota was exhausted (HTTP 403); orange cells additionally lack a full-year BIS policy rate. Countries below the line contribute to descriptive statistics only.

4.2 Dependent variables

Three outcomes are used, each with a distinct limitation (Fig. 3). Net capacity additions in MW are the most direct measure of expansion but are dominated by China, whose annual additions reach 75,931 MW against a sample median of 706 MW; any MW regression is in effect a China regression. They are also net: nine country-years in the panel show small negative values (retirements or statistical revisions, for example South Africa 2021, −21 MW), so additions must not be read as gross investment expenditure. Capacity growth, $100 \cdot \Delta \ln$ capacity, is scale-free and is the headline outcome; its limitation is that growth from a small base is mechanically volatile (Hungary and Czechia). The log capacity level, with its own lag as a regressor, is a third specification that treats the stock rather than the flow; it is the most persistent and least informative about year-to-year financing conditions.

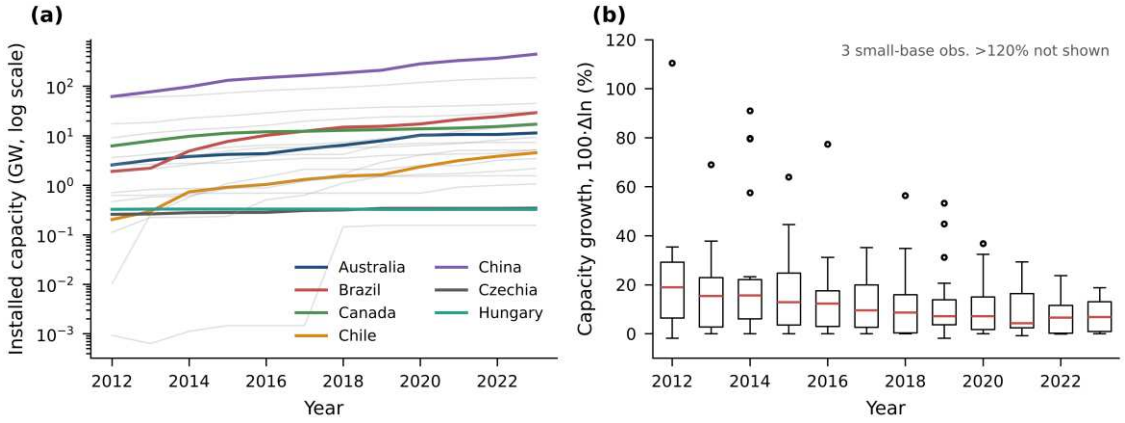


Fig. 3. Wind capacity, 2012–2023. (a) Installed capacity, log scale; coloured lines are the seven countries with verified supplier data, grey lines the remaining thirteen. (b) Annual capacity growth across all 20 countries.

4.3 Supplier concentration

For each reporter-year we retrieve bilateral annual imports of HS 850231 (wind-powered generating sets) with no second partner, retain partners with positive values, exclude the world aggregate from the share calculation and require the partner sum to reconcile to the reported world total within 2 per cent. The HHI is the sum of squared partner shares; its evolution and distribution are shown in Figs. 4 and 5. Two illustrative values anchor interpretation: two equal suppliers give 0.5, one supplier gives 1.0. We recomputed every HHI in the panel from the bilateral rows and obtained a maximum absolute discrepancy of 1.1×10^{-16}; all 65 usable observations reconcile at exactly 1.000. The HHI enters lagged one year so that concentration is predetermined relative to the capacity outcome, and it is centred on its sample mean (0.585) so that the coefficient on the rate change is interpretable at average concentration. Partner codes 899 and 490 (unspecified or other areas) appear in 28 raw rows and are retained as reported; a variant excluding them is a planned robustness check once the full sample exists. Concentration measured on HS 850231 is concentration of turbine-set suppliers and says nothing about towers, blades or gearboxes traded under other headings.

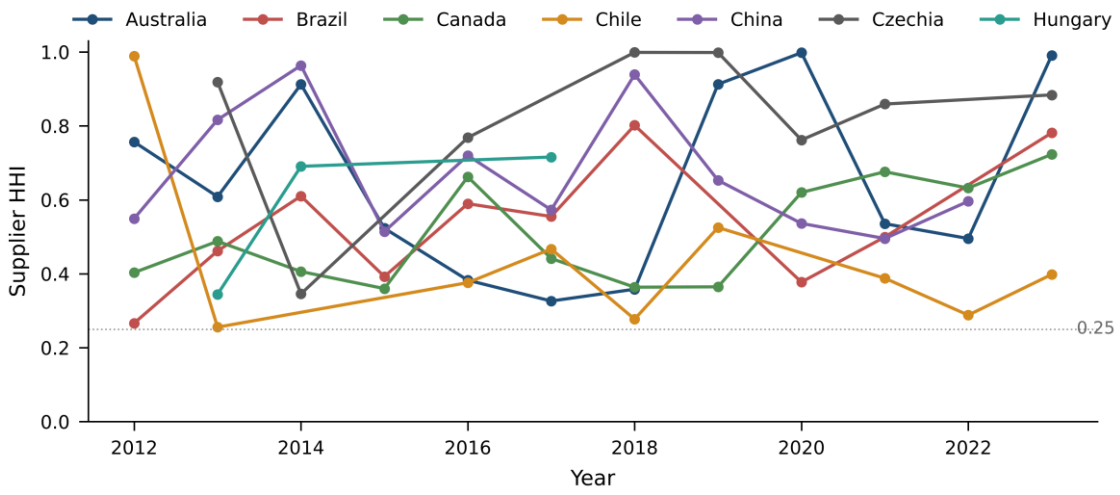


Fig. 4. Wind-equipment (HS 850231) import supplier HHI by country and year, pilot sample. The dotted line marks HHI = 0.25, the conventional threshold for high concentration.

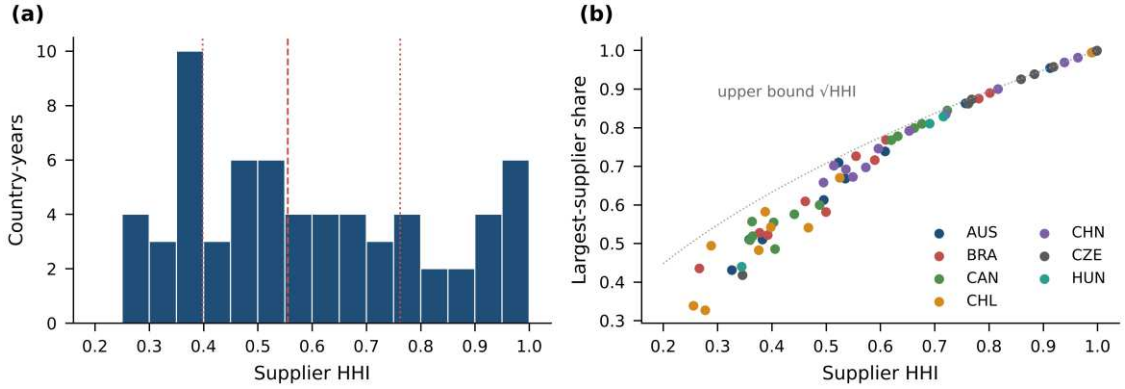


Fig. 5. Distribution of supplier concentration. (a) Histogram of HHI with quartiles (dotted) and median (dashed). (b) Largest-supplier share against HHI; the dotted curve is the bound √HHI.

4.4 Monetary conditions

The BIS monthly policy-rate series is averaged within the calendar year; years with fewer than 12 monthly observations (Japan 2013 and 2016, four months each; 2014–2015, none) are set to missing rather than averaged over a partial year. The annual change in this mean is the tightening variable. It is an observational measure of the stance, endogenous to the business cycle and to inflation, and we never describe it as a monetary shock. A tightening indicator ($\Delta r > 0.25$ pp) and the rate level are alternatives.

4.5 Control variables

GDP growth captures the demand and cash-flow conditions that move both policy rates and investment; omitting it would load cyclical variation onto the rate change. Inflation is the

proximate driver of tightening and of real financing costs. Trade openness proxies the ease of equipment importing and general exposure to external shocks. Log GDP and log population absorb scale, important because additions are measured in physical units. Log CO_2 proxies the carbon intensity of the existing system and hence the policy pressure to decarbonise. Renewable share controls for the maturity of the broader transition but ends in 2021 and is used only in robustness.

4.6 Data audit and the sample actually used

Every worksheet of the source workbook was audited before estimation (Appendix A). Capacity, controls and policy rates are essentially complete. Supplier concentration is not: at the time of writing, bilateral HS 850231 data are verified for seven countries (Australia, Canada, China, Brazil, Chile, Czechia, Hungary) and 65 country-years, because the anonymous Comtrade quota returned HTTP 403 for the remaining 175 country-years. The lagged interaction sample is 60 country-years, 7 countries, between 3 (Hungary) and 11 (Australia, Canada, China) years each. Table 3 reports the sample used by every specification. No value is imputed, interpolated or assumed to be zero.

Table 3 Sample selection and country coverage, pilot sample.

Panel A. Sample selection

Item	Value
Potential country-years (20 × 12)	240
Capacity observed	240
Supplier HHI observed	65
Lagged HHI available	60
Policy-rate change with full monthly coverage	235
Core estimation sample	60 country-years, 7 countries
Balanced?	No (3–11 years per country)

Panel B. Country coverage

Country	Years in core sample	Mean lagged HHI	Mean net additions (MW)
Australia	11 (2013–2023)	0.619	797
Brazil	9 (2013–2022)	0.506	2,279
Canada	11 (2013–2023)	0.492	981
Chile	8 (2013–2023)	0.446	401
China	11 (2013–2023)	0.669	34,573
Czechia	7 (2014–2022)	0.807	10
Hungary	3 (2014–2018)	0.584	0

Table 4 Descriptive statistics, core pilot sample (N = 60)

Variable	Mean	SD	Min	Median	Max
Net additions (MW)	7,061	15,983	−0.3	706	75,931
Capacity growth (%)	16.0	16.7	−0.1	13.0	91.0
ln capacity	8.89	2.08	5.63	9.20	13.00
Lagged HHI	0.585	0.216	0.256	0.543	0.999
Largest supplier share	0.701	0.182	0.327	0.705	0.999
Policy rate (%)	3.68	3.51	0.05	2.53	14.17
Policy-rate change (pp)	0.27	1.96	−4.25	−0.08	7.81
GDP growth (%)	2.75	3.02	−6.14	2.80	8.57
Inflation (%)	3.34	2.99	−0.23	2.44	15.10
Trade openness (% GDP)	64.5	41.2	24.3	47.5	168.4

The correlation matrix (replication file T04) shows the pairs that matter for design: the HHI and the largest supplier share correlate at 0.98, so the latter is a near-equivalent robustness measure rather than an independent one; the rate change and capacity growth are uncorrelated unconditionally (−0.01); and the rate level correlates with growth at 0.41, driven by Brazil and Chile where both are high, which is exactly the between-country variation that country fixed effects remove.

5. Econometric methodology

The guiding principle is that every estimator must be defensible for a panel of at most 20 countries and 12 years. We therefore use three estimators we can justify—two-way fixed effects with interaction, a parsimonious dynamic specification, and local projections—rather than a catalogue of methods whose asymptotics the sample cannot support.

5.1 Baseline two-way fixed-effects model

$$Y_{it} = \beta_1 \Delta r_{it} + \beta_2 \widetilde{H}_{i,t-1} + \beta_3 \left(\Delta r_{it} \times \widetilde{H}_{i,t-1}\right) + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it} \qquad (3)$$

Y is capacity growth (headline), net additions or log capacity; Δr is the annual policy-rate change in percentage points; $\tilde{H}$ is the lagged supplier HHI minus its sample mean (0.585); X contains GDP growth, inflation and trade openness in the baseline and adds log GDP, log population and log CO_2 in the full specification; μ and λ are country and year effects; ε is the error. Country effects remove time-invariant differences in resource endowment, grid structure

and support regime; year effects remove global shocks common to all countries, including the global rate cycle, so that β_1 is identified from countries tightening more or less than the cross-section average in a given year.

Interpretation follows Brambor, Clark and Golder (2006). Because $\tilde{H}$ is centred, β_1 is the association between a one-point rate rise and capacity growth at mean concentration; β_2 is the association between concentration and growth when the rate is unchanged; β_3 is the change in the rate effect per unit of HHI. β_3 is never interpreted alone. The quantity of interest is the marginal effect

$$\frac{\partial Y_{it}}{\partial \Delta r_{it}} = \beta_1 + \beta_3 \tilde{H}_{i,t-1} \qquad (4)$$

with delta-method variance $\mathrm{Var}(\beta_1) + \tilde{H}^2 \mathrm{Var}(\beta_3) + 2\tilde{H}\,\mathrm{Cov}(\beta_1, \beta_3)$, evaluated across the observed distribution of lagged HHI and plotted with a histogram of the moderator so that effects are not read outside the support (Hainmueller, Mummolo, & Xu, 2019). Confidence bands use the t distribution with $G - 1$ degrees of freedom.

5.2 Alternative dependent variables

The same specification is estimated for net additions in MW, net additions per GW of prior capacity (a scale-free flow that does not require positive capacity) and log capacity with its own lag. The MW model is reported for completeness and read with the China caveat of Section 4.2; the per-GW model is the preferred flow alternative; the level model tests whether any association survives in the stock.

5.3 Dynamic specification

Capacity expansion is plausibly persistent: pipelines, permitting and grid connection spread a financing decision over several years. We therefore estimate

$$Y_{it} = \rho Y_{i,t-1} + \beta_1 \Delta r_{it} + \beta_2 \tilde{H}_{i,t-1} + \beta_3 \left(\Delta r_{it} \times \tilde{H}_{i,t-1}\right) + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it} \qquad (5)$$

by least-squares dummy variables. The lagged dependent variable is correlated with the within-transformed error, producing the Nickell (1981) bias of order 1/T; with T ≤ 11 the bias in ρ is non-negligible and biases toward zero. We report the LSDV estimate as a bound and explicitly do not apply difference or System GMM: with 20 countries, the instrument count required to address a lagged dependent variable and potentially endogenous concentration would approach or exceed the number of clusters, weakening the Hansen test and making the asymptotic inference unreliable (Roodman, 2009). A bias-corrected LSDV estimator is the appropriate extension once the full panel is available.

5.4 Nonlinearity

Two restricted approaches ask whether the interaction changes across the concentration distribution. A quadratic interaction adds $\tilde{H}^2$ and $\Delta r \times \tilde{H}^2$. A single-threshold profile estimates the rate effect below and above each candidate HHI cut-point between the 15th and 85th percentiles and reports the residual sum of squares; following Hansen (1999) the minimum identifies the candidate threshold, but we do not report a threshold estimate or its confidence set, because the bootstrap p-value for a threshold effect is not credible with seven clusters in the pilot and would require the full 20-country sample at minimum. No threshold value is claimed.

5.5 Local projections

To trace dynamics we estimate, for horizons h = 0 to 3 years,

$$\ln C_{i,t+h} - \ln C_{i,t-1}$$
$$= \alpha_{i,h} + \lambda_{t,h} + \beta_h S_{it} + \theta_h\left(S_{it} \times \tilde{H}_{i,t-1}\right) + \Gamma_h' X_{it} + \phi_h Y_{i,t-1}$$
$$+ \varepsilon_{i,t+h} \qquad (6)$$

where C is capacity and S the shock variable (Jordà, 2005). The pipeline runs two designs. In the association design $S = \Delta r$, and the impulse responses are conditional correlations between

rate changes and subsequent capacity, not causal responses. In the shock design, which activates automatically when an identified series is merged, *S* is the annual sum of high-frequency monetary-policy surprises. Bolhuis, Das and Yao (2024) provide such a series for 21 advanced and 8 emerging economies over 2000–2022, decomposed into a pure monetary shock and a central-bank information component in the spirit of Jarociński and Karadi (2020) and Bauer and Swanson (2023); Emeksiz, Kamber, Otten and Pasricha (2026) offer an intraday alternative for 16 economies. Horizons beyond h = 3 are not estimated: with 12 years, each additional horizon removes a year from every country and the h = 3 sample already loses 27 observations relative to h = 0.

5.6 Identification

The policy-rate change is not an exogenous monetary shock. Central banks raise rates in response to inflation and growth, both of which also move investment, and the year effects remove only the common component of that response. The TWFE and association-design estimates are therefore conditional associations. We do not claim instrumental-variable identification: there is no instrument in the dataset with a defensible exclusion restriction for a national policy rate. Concentration is lagged so that it is predetermined, but it may still be endogenous to anticipated expansion (a country planning large additions may concentrate orders with one supplier), and this is noted as a limitation rather than solved. Causal language is reserved for the shock design.

5.7 Inference with few clusters

With G = 20 countries in the full design and G = 7 in the pilot, conventional cluster-robust standard errors over-reject (Cameron, Gelbach, & Miller, 2008; MacKinnon, Nielsen, & Webb, 2023). We report, for every headline coefficient, four p-values: conventional, CR1 cluster-robust with the normal reference distribution, CR1 with the t(G − 1) reference distribution, and

a restricted wild-cluster bootstrap-t. The bootstrap imposes the null, re-estimates with the coefficient of interest set to zero, resamples residuals cluster by cluster with Webb (2023) six-point weights when $G < 12$ and Rademacher weights otherwise, and compares the observed CR1 t-statistic with the bootstrap distribution of 4,999 replications. Driscoll–Kraay (1998) standard errors are also reported because they are common in this literature, but the Bartlett kernel with two lags is designed for large T and the pilot comparison shows they understate the bootstrap-consistent standard error by a factor of two to three, so they are not used for inference. Diagnostics comprise the Pesaran (2021) CD statistic on residuals, an AR(1) residual regression, a Breusch–Pagan test on within-transformed regressors and variance inflation factors. Leave-one-country-out re-estimation is reported for every specification.

5.8 Machine-learning block

The predictive block is complementary. Elastic net, random forest and gradient boosting are trained to predict capacity growth from the same regressors, with and without the interaction term, under two leakage-safe validation schemes: leave-one-country-out, which measures cross-sectional generalisation, and a chronological split training to 2019 and testing on 2020–2023, which measures temporal generalisation. Hyper-parameters are tuned by grouped k-fold cross-validation inside each training fold, so no test country or test year informs tuning. Performance is compared with the training-fold mean and with a ridge regression on year dummies. SHAP values (Lundberg & Lee, 2017) are computed on out-of-fold predictions and reported as predictive attribution, with the rank instability of each feature across folds; they are not causal effects, and the section is explicit that predictive performance below the baseline means the SHAP values describe an over-fitted model.

6. Empirical results (pilot sample)

All estimates in this section use the 60-country-year, 7-country pilot sample. With seven clusters, no coefficient is interpretable as a population parameter; the section reports them in full because the master design requires that nothing be omitted and because they reveal what the sample can and cannot support. Every number traces to regression_tables.xlsx in the replication package.

6.1 Baseline estimates

Table 5 Two-way fixed-effects estimates, dependent variable capacity growth (100·Δln MW)

	M1	M2	M3	M4
Policy-rate change (pp)	0.531 (2.301) [0.67]	0.547 (2.276) [0.66]	1.127 (1.741) [0.61]	0.244 (1.756) [0.89]
Lagged HHI, centred		−4.077 (7.080) [0.59]	−3.200 (6.928) [0.68]	0.484 (7.392) [0.94]
Rate change × lagged HHI			4.925 (4.193) [0.23]	3.570 (5.225) [0.56]
Controls	GDP growth, inflation, openness	same	same	+ ln GDP, ln population, ln CO_2
Country and year effects	Yes	Yes	Yes	Yes
N (countries)	60 (7)	60 (7)	60 (7)	60 (7)
Within R^2	−0.117	−0.123	−0.104	−0.582

Notes. CR1 country-clustered standard errors in parentheses; restricted wild-cluster bootstrap p-values (Webb weights, B = 4,999) in brackets. Negative within R^2 arises because linearmodels reports the within fit net of absorbed effects and the regressors explain less than the fixed effects alone. The HHI is centred on 0.585, so the rate-change coefficient is the association at mean concentration.

None of the three hypotheses receives support, and the pattern of signs is informative about the sample rather than the mechanism. The rate change carries a positive point estimate at mean concentration, contrary to H1, with a bootstrap p-value above 0.6 in every column. The lagged HHI is negative in M2 and M3, in the direction of H2, but with a standard error more than twice the coefficient. The interaction is positive—the buffering sign, contrary to H3—with a bootstrap p-value of 0.23 in M3 that rises to 0.56 when scale controls are added in M4. Figs. 6 and 7 show the raw and within-country associations behind these estimates. A positive point

estimate on the rate change in a sample whose large tightening episodes (Brazil 2015 and 2021–2022, Chile 2022) coincided with large wind build-outs under auction pipelines committed years earlier is exactly the alternative outcome anticipated for H1 in Section 3.3; it is also exactly what seven countries cannot adjudicate.

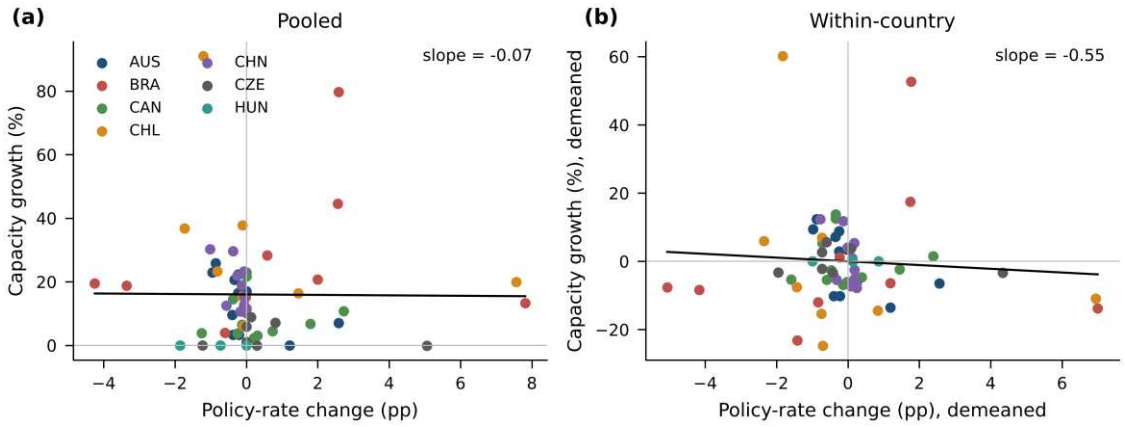


Fig. 6. Policy-rate changes and capacity growth, pilot sample. (a) Pooled observations. (b) Both variables demeaned by country. Lines are bivariate OLS fits without controls or year effects.

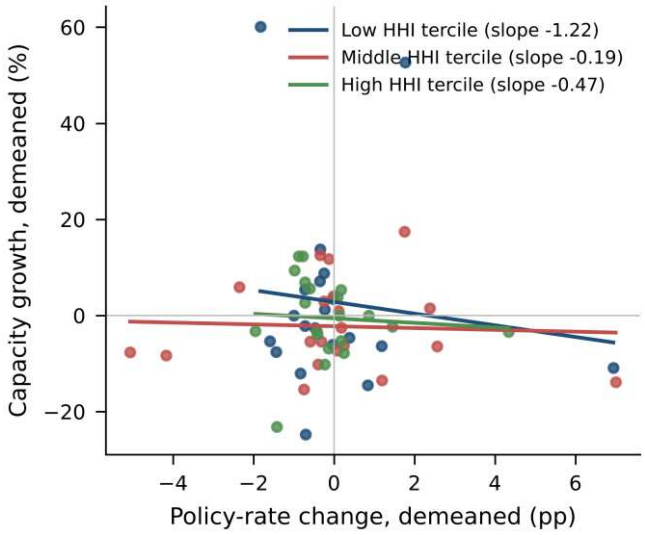


Fig. 7. Within-country association between policy-rate changes and capacity growth by tercile of lagged supplier HHI. Bivariate fits; slopes in the legend.

6.2 Alternative dependent variables

Table 6 Alternative outcomes, specification M3

Dependent variable	Rate change	Lagged HHI	Interaction	N	Within R^2
Net additions (MW)	1,018 (1,113) [0.27]	−5,371 (5,035) [0.33]	3,477 (5,047) [0.63]	60	−0.087
Net additions per prior GW	15.0 (32.4) [0.62]	−73.2 (90.5) [0.43]	82.6 (73.7) [0.22]	60	−0.082

Dependent variable	Rate change	Lagged HHI	Interaction	N	Within R²
Capacity growth (%)	1.13 (1.74) [0.61]	−3.20 (6.93) [0.68]	4.93 (4.19) [0.23]	60	−0.104
ln capacity (with lagged level)	−0.007 (0.012) [0.55]	−0.079 (0.121) [0.56]	−0.079 (0.052) [0.17]	53	0.893

The three flow measures agree in sign; the level model, which is dominated by the lagged level (coefficient not shown, $R^2 = 0.89$), flips the interaction negative. That the sign of the central coefficient depends on whether the outcome is a flow or a stock, in a sample this small, is itself evidence that the result is not yet a result.

6.3 Marginal effects across the concentration distribution

Table 7 Marginal effect of a 1 pp rate increase on capacity growth, by lagged HHI

Percentile of lagged HHI	HHI	Marginal effect (pp)	SE	95% CI, t(6)	p
10th	0.346	−0.05	2.62	[−6.46, 6.35]	0.99
25th	0.391	0.17	2.44	[−5.80, 6.15]	0.95
50th	0.543	0.92	1.89	[−3.70, 5.53]	0.64
75th	0.729	1.84	1.31	[−1.36, 5.03]	0.21
90th	0.921	2.78	1.03	[0.25, 5.31]	0.04

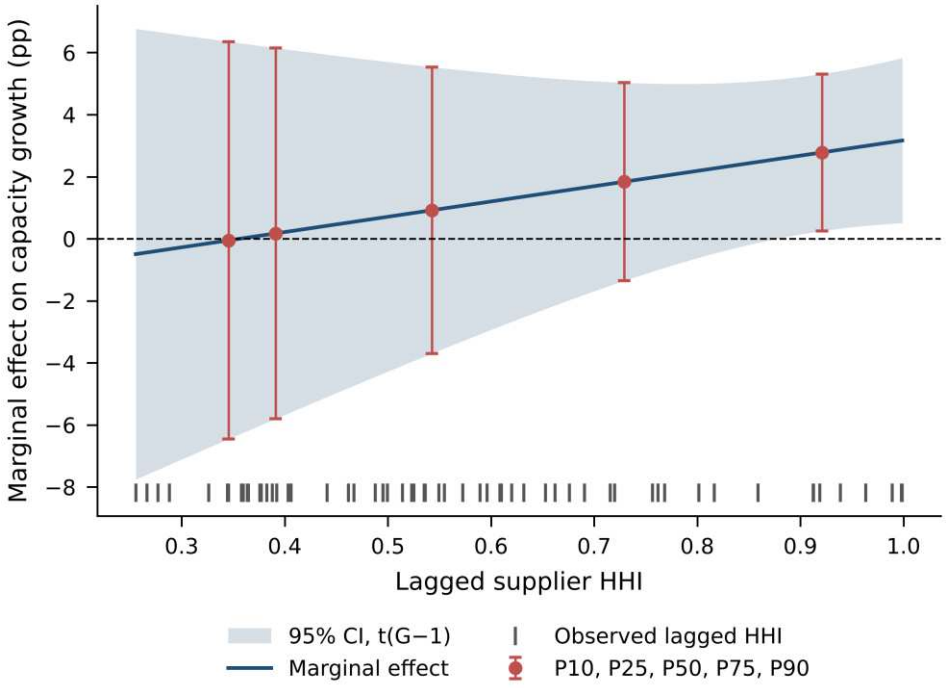


Fig. 8. Marginal effect of a one-percentage-point policy-rate increase on capacity growth across lagged supplier HHI (specification M3). Shaded band: 95% confidence interval using t(G−1) critical values; points: effects at the 10th, 25th, 50th, 75th and 90th percentiles; ticks: observed values of the moderator.

Fig. 8 plots the full profile against the moderator's histogram. The marginal effect is indistinguishable from zero across the lower three-quarters of the distribution and becomes

positive, with a t(6) p-value of 0.04, only above an HHI of about 0.9—a region of nine country-years (Australia 2015, 2020–2021; Chile 2013; China 2015, 2019; Czechia 2014, 2019–2020). Read literally this is the lock-in channel: in the most concentrated country-years, tightening coincides with faster expansion. Read responsibly it is nine observations from four countries, and the 95 per cent band at the 90th percentile barely clears zero under a reference distribution with six degrees of freedom. We report it because the design requires marginal effects; we do not interpret it.

6.4 Dynamics

The LSDV coefficient on lagged capacity growth is 0.054 (SE 0.088), and on lagged additions per GW −0.008 (SE 0.085): no detectable persistence at annual frequency in the pilot, which is consistent with growth rates being close to white noise around a country mean once the stock is large, and with the Nickell bias pushing ρ toward zero. The interaction in the dynamic model (5.66, SE 4.14) is indistinguishable from the static one. The association-design local projections (Fig. 9) show the cumulative association between a rate change and log capacity rising from 0.006 at h = 0 to 0.086 at h = 3, with every confidence band spanning zero and the h = 3 sample reduced to 33 observations. No dynamic claim is made.

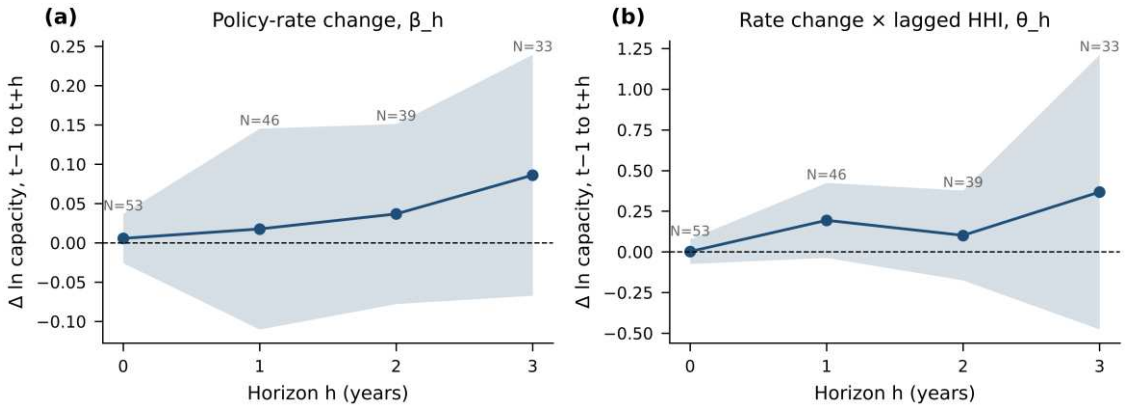


Fig. 9. Local projections (association design) of cumulative log capacity on (a) the policy-rate change and (b) its interaction with lagged HHI, horizons h = 0–3. Shaded bands are 95% CR1 intervals; N at each horizon shown above the band. These are conditional associations, not responses to identified shocks.

6.5 Nonlinearity

The quadratic interaction terms (lagged HHI² = −7.6, SE 38.8; rate change × HHI² = −41.7, SE 48.1) are uninformative. The SSR profile across candidate thresholds (Fig. 10) has its minimum at the lowest candidate cut-point (HHI = 0.363, the 15th percentile), where only nine observations fall in the low regime, and the difference in the rate effect across regimes (11.7, cluster p = 0.25) is not distinguishable from zero at any candidate; in line with Section 5.4, no threshold is reported.

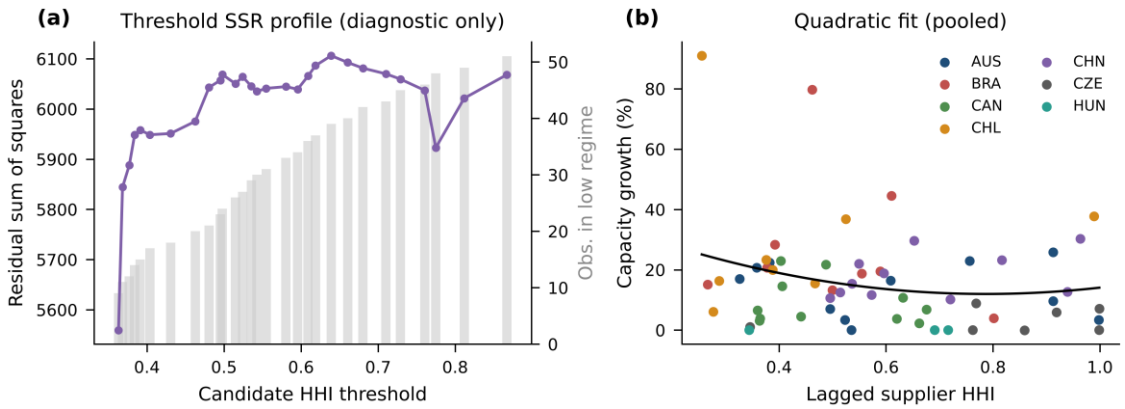


Fig. 10. Nonlinearity diagnostics. (a) Residual sum of squares across candidate single-threshold cut-points, with the number of observations in the low regime (bars); no threshold is estimated. (b) Pooled quadratic fit of capacity growth on lagged HHI.

7. Machine-learning analysis (pilot sample)

The predictive block asks two questions the econometric block cannot: whether any flexible function of the regressors predicts capacity growth out of sample, and whether adding the rate-change × concentration interaction improves prediction. In the pilot the answer to both is no, and the section's value is in showing what disciplined validation does to a sample of 60.

Table 8 Out-of-sample predictive performance, dependent variable capacity growth (%)

Validation	Features	Model	RMSE	MAE	R² (out of sample)	n test
Leave-one-country-out	without interaction	Training-fold mean	17.42	11.77	−0.105	60
Leave-one-country-out	without interaction	Ridge + year effects	16.35	12.05	0.027	60
Leave-one-country-out	without interaction	Elastic net	18.79	13.63	−0.285	60
Leave-one-country-out	without interaction	Random forest	18.94	12.59	−0.306	60

Validation	Features	Model	RMSE	MAE	R^2 (out of sample)	n test
Leave-one-country-out	without interaction	XGBoost	19.82	13.25	−0.430	60
Leave-one-country-out	with interaction	Elastic net	18.73	13.54	−0.277	60
Leave-one-country-out	with interaction	Random forest	18.71	12.49	−0.274	60
Leave-one-country-out	with interaction	XGBoost	19.66	13.05	−0.408	60
Chronological (train ≤ 2019, test ≥ 2020)	without interaction	Training mean	11.93	10.21	−0.306	20
Chronological	without interaction	Elastic net	12.65	11.00	−0.469	20
Chronological	without interaction	Random forest	13.21	11.64	−0.601	20
Chronological	with interaction	Elastic net	12.57	10.84	−0.450	20
Chronological	with interaction	XGBoost	15.72	13.30	−1.267	20

Notes. Hyper-parameters tuned by grouped 5-fold cross-validation within each training fold. Out-of-sample R^2 is $1 - \mathrm{SSE}/\mathrm{SST}$ on the held-out observations, so values below zero mean the model predicts worse than the held-out mean.

Every learner is beaten by the training-fold mean under both validation schemes, and the only model with a positive out-of-sample R^2 (0.027) is the ridge regression with year effects—essentially the econometric model. Adding the interaction (Fig. 11) changes RMSE by less than 0.3 points in every row, in both directions. This is the expected behaviour of flexible learners on 60 observations with heavy country heterogeneity: the random forest and gradient boosting fit country-specific patterns that do not transfer to the held-out country. The result is not a finding about the mechanism; it is a demonstration that the interaction cannot be validated predictively at this sample size, and that any SHAP ranking from these models describes over-fitting.

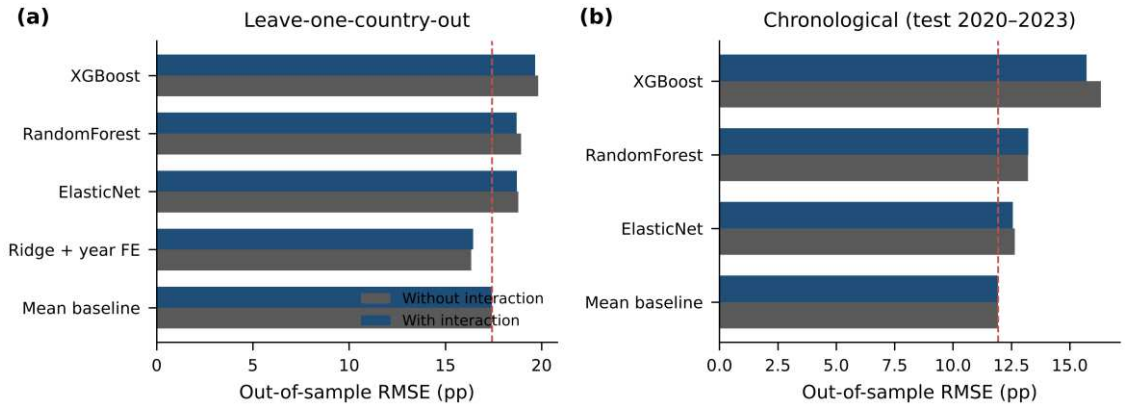


Fig. 11. Out-of-sample RMSE by learner, with and without the rate-change × HHI interaction. (a) Leave-one-country-out. (b) Chronological split, train ≤ 2019, test 2020–2023. Dashed line: training-mean baseline.

For completeness, Table 9 reports SHAP attribution from the gradient-boosting model with the interaction, computed on leave-one-country-out predictions, together with the standard deviation of each feature's importance rank across the seven folds.

Table 9 SHAP predictive attribution (XGBoost, leave-one-country-out)

Feature	Mean absolute SHAP (pp)	Rank SD across folds
Policy-rate level	3.34	1.5
GDP growth	2.46	1.9
ln CO_2	2.01	3.3
Trade openness	1.82	1.2
Inflation	0.95	1.7
Lagged HHI	0.79	1.6
ln population	0.75	2.8
Policy-rate change	0.73	2.7
Rate change × HHI	0.67	2.4
ln GDP	0.63	2.6

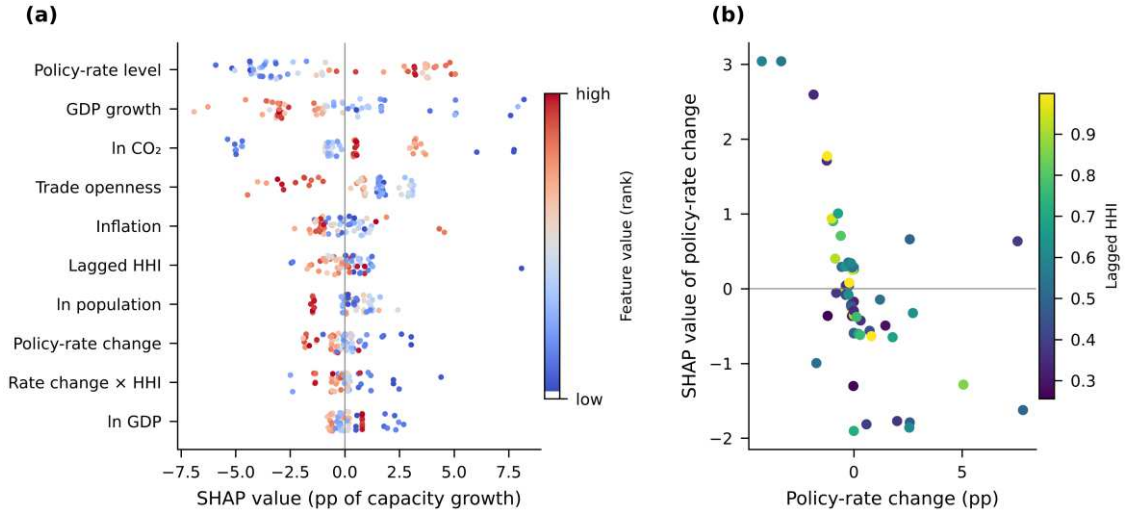


Fig. 12. SHAP attribution for the gradient-boosting model with interaction, leave-one-country-out predictions. (a) Beeswarm by feature, coloured by feature-value rank. (b) SHAP value of the policy-rate change against its value, coloured by lagged HHI. Predictive attribution only.

The rate level ranks first and the two monetary-change variables rank near the bottom, with rank standard deviations of 2.4–2.7 positions across folds—the interaction is sometimes third and sometimes tenth depending on which country is held out. Fig. 12 displays the full attribution. In a model that does not beat the mean, these attributions are a description of the fit, not of the data-generating process, and they are shown only to document the instability the design anticipated.

Two design features of the block carry over unchanged to the full sample. Leave-one-country-out remains the right scheme because the policy question is whether the mechanism generalises across countries, and the chronological split remains the right complement because the 2022–2023 tightening is the out-of-sample event of interest. With roughly 200 observations and 20 folds, the comparison between learners and the ridge-with-fixed-effects baseline becomes informative; with 60 it is not.

8. Robustness and sensitivity (pilot sample)

8.1 Inference comparison

Table 10 Alternative inference for specification M3

Coefficient	Estimate	p, conventional	p, CR1 normal	p, CR1 t(6)	p, Driscoll–Kraay	p, wild-cluster bootstrap
Rate change	1.127	0.525	0.521	0.541	0.224	0.612
Lagged HHI	−3.200	0.742	0.647	0.660	0.605	0.678
Interaction	4.925	0.460	0.248	0.285	0.154	0.225

Notes. Driscoll–Kraay standard errors are 0.91 (rate change) and 3.38 (interaction) against CR1 values of 1.74 and 4.19; the bootstrap-consistent p-values are closer to the CR1 t(6) values than to any other column.

The table, and Fig. 13, make the methodological point the full-sample paper will rely on. Driscoll–Kraay standard errors are half the size of the cluster-robust ones and would, at a more generous sample, manufacture significance that the cluster structure does not support; with $T \leq$

11 the kernel has nothing to average over. The wild-cluster bootstrap and the t(G − 1) correction agree closely, and both are reported for every coefficient in the replication file.

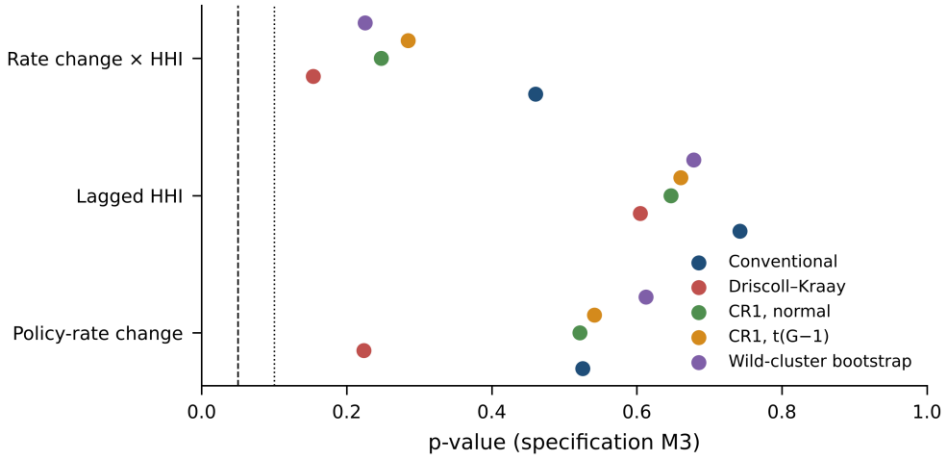


Fig. 13. p-values for the three M3 coefficients under five inference methods. Dashed and dotted lines mark 0.05 and 0.10.

8.2 Specification variants

Table 11 Interaction coefficient under alternative specifications

Variant	Interaction	SE	p, CR1	p, bootstrap	N
Baseline (M3)	4.93	4.19	0.25	0.23	60
Rate level instead of change	1.33	2.05	0.52	—	60
Tightening dummy ($\Delta r > 0.25$ pp)	33.5	23.9	0.17	—	60
Largest supplier share instead of HHI	8.22	5.81	0.17	—	60
HHI lagged two years	−9.31	6.25	0.15	—	53
Exclude 2020–2021	3.06	11.9	0.80	0.84	50
Exclude 2022–2023	6.09	17.1	0.73	0.68	50
Add renewable share	3.13	15.9	0.85	0.86	50
No controls	7.78	3.14	0.02	0.31	60
Countries with ≥ 9 years only	−21.2	13.3	0.13	0.23	42

The interaction keeps its positive sign across measures of tightening and concentration and across period exclusions, but with standard errors that grow to three times the coefficient once any year block is dropped. Two rows deserve comment. The no-controls specification is the only one with a conventional cluster p-value below 0.05 (0.02); its bootstrap p-value is 0.31. A referee would be right to treat the former as the small-cluster over-rejection that Section 5.7 warns against, and the comparison is reported for that reason. The balanced-subsample row,

which drops Czechia and Hungary, reverses the sign to −21.2: with the two smallest wind markets removed, the remaining five countries display the amplification sign of H3. That the central coefficient changes sign when two countries contributing ten observations are removed is the clearest single statement of why the pilot cannot support inference. In the contemporaneous-HHI row the rate change itself becomes large and significant (5.31, bootstrap p = 0.01); because concentration measured in the same year as the outcome is plausibly a consequence of the order book rather than a predetermined condition, this is reported as a diagnostic of simultaneity and not used.

8.3 Leave-one-country-out

Table 12 Specification M3 excluding one country at a time

Excluded	Rate change (SE)	Interaction (SE)	N
Australia	1.08 (1.75)	4.00 (3.56)	49
Brazil	−3.39 (2.83)	3.29 (2.36)	51
Canada	1.34 (2.05)	5.49 (3.30)	49
Chile	2.29 (1.02)	−1.44 (6.83)	52
China	1.25 (2.17)	4.34 (5.18)	49
Czechia	0.93 (1.02)	9.24 (15.5)	53
Hungary	1.04 (1.64)	6.25 (4.00)	57

Dropping Brazil flips the rate-change coefficient negative (−3.39), which is the H1 sign; Brazil's 2021–2022 tightening of more than 7 pp coincided with the completion of a large auction pipeline. Dropping Chile flips the interaction negative. Fig. 14 shows these estimates with their intervals; every interval contains zero and most contain each other.

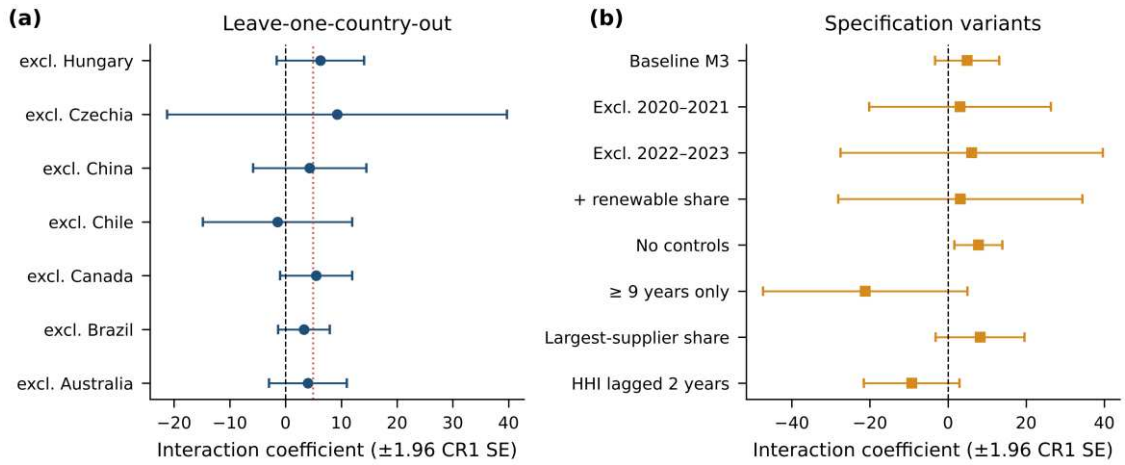


Fig. 14. Stability of the interaction coefficient. (a) Leave-one-country-out; dotted line is the full-sample estimate. (b) Specification variants measured in the same units as the baseline. Bars are ±1.96 CR1 standard errors.

8.4 Diagnostics

The Pesaran CD statistic on M3 residuals is −1.92 (p = 0.055), computed over pairwise overlaps of at least four years; weak evidence of residual cross-sectional dependence that year effects do not absorb, to be re-examined with 20 countries. The AR(1) coefficient on residuals is 0.03 (p = 0.83), so serial correlation is not a concern at annual frequency in this sample. Breusch–Pagan on within-transformed regressors gives 4.87 (p = 0.56). Variance inflation factors are all below 4 (rate change 3.7, inflation 2.9, remaining regressors below 1.7), so the interaction does not induce collinearity once the moderator is centred.

8.5 What the sensitivity analysis establishes

The results are not stable; they are not meant to be. The analysis establishes three things for the full-sample design: that the bootstrap and t(G − 1) inference agree and Driscoll–Kraay does not; that the interaction sign depends on the inclusion of the two smallest markets, so the full sample must report results with and without small-base countries; and that Brazil's 2021–2022 episode is influential enough to warrant explicit treatment, most naturally through the identified-shock design that separates its policy response from the committed pipeline.

9. Discussion and policy implications

9.1 What the framework and the pilot together imply

The MSTM predicts that the monetary footprint on wind expansion is conditional on sourcing structure, and that the sign of the conditioning is theoretically open: cost-pass-through and delay channels amplify, lock-in buffers. The pilot does not resolve the sign. Its positive interaction with a bootstrap p-value of 0.23, reversing under a balanced subsample, is precisely the indeterminacy the theory allows, measured in a sample too small to adjudicate. What the pilot does establish is procedural: the HHI can be constructed from bilateral trade with full reconciliation; the marginal-effect profile can be estimated and bounded; the identified-shock

extension has an available data source; and the inference tools that will matter at G = 20 behave as the literature predicts at G = 7.

9.2 Policy implications supported by the design, not yet by the results

Because the empirical results are diagnostic, we separate what the framework implies in principle from what would require the full-sample evidence.

Central banks. Nothing here supports changing the monetary stance to favour wind, and we do not recommend it; price stability is the mandate and the carbon footprint of tightening, where it exists, is an externality for fiscal and industrial policy to address. The relevant implication is informational: if H3 holds in the full sample, the sectoral incidence of a rate rise varies with the importing country's supplier structure, which is a legitimate input to the distributional analysis central banks already publish, and a reason to include equipment-sourcing indicators in climate-related financial-stability monitoring.

Energy ministries and procurement authorities. The lagged-HHI design assumes concentration is predetermined relative to expansion, but Section 8.2 shows that contemporaneous concentration behaves differently, which is consistent with concentration being partly chosen through the auction and procurement process. Auction design that scores supplier diversity, or that permits re-phasing of delivery when financing conditions change, is the policy lever the framework points to. Whether it is worth its cost in higher equipment prices is the trade-off the full sample is designed to quantify.

Renewable developers. The delay channel implies that the value of supplier optionality is highest during tightening cycles; multi-sourcing of turbine sets, even at a premium, functions as a hedge against the carrying cost of committed capital. The pilot cannot say how large the premium can be before the hedge stops paying.

Trade policymakers. The distinction between concentration and dependence matters directly. Tariff or local-content measures that raise import dependence on domestic production

may *increase* supplier concentration if the domestic industry is itself a single producer. Eschmann and Jochem (2025) show that expansion does not mechanically raise import dependence; the open question is whether it raises concentration.

Wind-equipment manufacturers. A manufacturer whose customers are concentrated on it faces, under the pass-through channel, the same financing-cycle exposure in reverse: order books that thin when importers' rates rise. Diversifying the customer base across monetary cycles is the manufacturer's counterpart to the importer's supplier diversification.

Supply-chain diversification strategies. The framework's central trade-off is between the price advantage of concentration on the lowest-cost supplier and the resilience value of alternatives when financial conditions tighten. The IEA (2023) recommends diversification on security grounds; the MSTM adds a financial-cycle rationale that is quantifiable, and the full-sample marginal-effect profile (Fig. 8) is the quantity that would make the recommendation evidence-based.

9.3 Broader recommendations requiring additional evidence

Three extensions would sharpen all of the above and are flagged rather than asserted. Bilateral exchange-rate exposure, weighted by supplier share, would test the pass-through channel directly. Order-to-commissioning lead times from project databases would test the delay channel. Component-level concentration (towers, blades, nacelles under their own HS headings) would test whether turbine-set concentration is a sufficient statistic for sourcing exposure. None is in the present data, and the policy statements above are correspondingly conditional.

10. Conclusion

This paper asks whether the structure of a country's wind-equipment sourcing conditions how monetary tightening transmits to wind-capacity expansion. It proposes a Monetary–

Supply-Chain Transmission Mechanism in which the financing channel of policy rates and the supply-chain channel of supplier concentration enter the same irreversible investment decision, so that concentration can amplify or buffer the financing shock; it constructs a supplier Herfindahl index from reconciled bilateral trade in wind-powered generating sets for 20 non-euro economies over 2012–2023; and it specifies an estimation and inference framework—two-way fixed effects with centred interactions and marginal-effect profiles, a bounded dynamic specification, local projections with a switch to an identified-shock design, restricted wild-cluster bootstrap inference, and a leakage-safe machine-learning block—calibrated to what twenty countries and twelve years can support.

The empirical evidence reported here comes from a seven-country pilot and is diagnostic. The interaction between rate changes and lagged concentration is positive, imprecise, and reverses sign when the two smallest markets are excluded; the rate-change coefficient reverses when Brazil is excluded; no machine-learning model predicts out of sample; and Driscoll–Kraay inference, had it been used, would have understated uncertainty by a factor of two to three. None of the four hypotheses is supported, and none is rejected. The pilot's contribution is to show that the design works, to identify the two data additions that make it decisive—bilateral HS 850231 imports for the remaining thirteen countries, freely available through a registered Comtrade account, and the IMF high-frequency monetary-shock series—and to fix in advance the specifications, inference and validation rules under which the full-sample results will be read, so that they cannot be selected after the fact.

The limitations are stated throughout and bear repeating: policy-rate changes are not shocks; concentration is predetermined but not exogenous; turbine-set concentration is not component concentration; the sample is purposive and excludes the euro area. Within those limits the question is well posed, the measurement is sound, and the machinery is reproducible.

Whether supplier diversification is a hedge against the financing cycle in the energy transition is an answerable question, and this paper sets out how to answer it.

CRediT authorship contribution statement

[Author One]: Conceptualization, Methodology. [Author Two]: Software, Formal analysis, Data curation, Writing – original draft. [Author Three]: Validation. [Author Four]: Writing – review & editing. [Author Three]: Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

All source data are public: UN Comtrade (bilateral HS 850231 imports), IRENA (installed capacity, archived via Zenodo record 10952917), BIS (central-bank policy rates) and the World Bank WDI. The replication package—download, construction, estimation, machine-learning and figure scripts, with the processed panel—will be deposited in a public repository [DOI/URL to be added] on acceptance.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used Chatgpt in order to correct the grammatical errors. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Baldwin, R., & Freeman, R. (2022). Risks and global supply chains: What we know and what we need to know. *Annual Review of Economics, 14*, 153–180. https://doi.org/10.1146/annurev-economics-051420-113737

Bauer, M. D., & Swanson, E. T. (2023). An alternative explanation for the "Fed information effect." *American Economic Review, 113*(3), 664–700. https://doi.org/10.1257/aer.20201220

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica, 77*(3), 623–685. https://doi.org/10.3982/ECTA6248

Bolhuis, M. A., Das, S., & Yao, B. (2024). *A new dataset of high-frequency monetary policy shocks* (IMF Working Paper No. 2024/224). International Monetary Fund. https://doi.org/10.5089/9798400290930.001

Brambor, T., Clark, W. R., & Golder, M. (2006). Understanding interaction models: Improving empirical analyses. *Political Analysis, 14*(1), 63–82. https://doi.org/10.1093/pan/mpi014

Cai, Y., & Wu, Y. (2021). Time-varying interactions between geopolitical risks and renewable energy consumption. *International Review of Economics & Finance, 74*, 116–137. [DOI to be confirmed]

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics, 90*(3), 414–427. https://doi.org/10.1162/rest.90.3.414

Chen, S.-S., & Lin, T.-Y. (2024). Monetary policy and renewable energy production. *Energy Economics, 132*, Article 107495. https://doi.org/10.1016/j.eneco.2024.107495

Cloyne, J., Ferreira, C., Froemel, M., & Surico, P. (2023). Monetary policy, corporate finance, and investment. *Journal of the European Economic Association, 21*(6), 2586–2634. https://doi.org/10.1093/jeea/jvad009

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics, 80*(4), 549–560. https://doi.org/10.1162/003465398557825

Egli, F., Steffen, B., & Schmidt, T. S. (2018). A dynamic analysis of financing conditions for renewable energy technologies. *Nature Energy, 3*(12), 1084–1092. https://doi.org/10.1038/s41560-018-0277-y

Emeksiz, E. O., Kamber, G., Otten, J., & Pasricha, G. K. (2026). *Still packing a punch: Monetary policy transmission in a new cross-country high-frequency dataset* (IMF Working Paper No. 2026/018). International Monetary Fund. https://doi.org/10.5089/9798229038898.001

Eschmann, J., & Jochem, P. (2025). Assessing import dependencies in the accelerating energy transition: A structural gravity model analysis. *Energy Policy, 205*, Article 114697. https://doi.org/10.1016/j.enpol.2025.114697

Grossman, G. M., Helpman, E., & Lhuillier, H. (2023). Supply chain resilience: Should policy promote international diversification or reshoring? *Journal of Political Economy, 131*(12), 3462–3496. https://doi.org/10.1086/725173

Hainmueller, J., Mummolo, J., & Xu, Y. (2019). How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice. *Political Analysis, 27*(2), 163–192. https://doi.org/10.1017/pan.2018.46

Hansen, B. E. (1999). Threshold effects in non-dynamic panels: Estimation, testing, and inference. *Journal of Econometrics, 93*(2), 345–368. https://doi.org/10.1016/S0304-4076(99)00025-1

Hashmi, S. M., Syed, Q. R., & Inglesi-Lotz, R. (2022). Monetary and energy policy interlinkages: The case of renewable energy in the US. *Renewable Energy, 201*(1), 141–147. [DOI to be confirmed]

Hirth, L., & Steckel, J. C. (2016). The role of capital costs in decarbonizing the electricity sector. *Environmental Research Letters, 11*(11), Article 114010. https://doi.org/10.1088/1748-9326/11/11/114010

International Energy Agency. (2023). *Energy Technology Perspectives 2023*. IEA. https://www.iea.org/reports/energy-technology-perspectives-2023

Jarociński, M., & Karadi, P. (2020). Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics, 12*(2), 1–43. https://doi.org/10.1257/mac.20180090

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review, 95*(1), 161–182. https://doi.org/10.1257/0002828053828518

Lucchetta, M. (2025). Monetary trade-offs: Stability vs. renewables. *Economics Letters, 257*. [article number to be confirmed]

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, 4765–4774.

MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics, 232*(2), 272–299. https://doi.org/10.1016/j.jeconom.2022.04.001

Mbassi, C. M., Samba, C. M., & Elomo Zogo, T. (2025). Does monetary policy fuel energy consumption across the world? Focus on inflation targeting. *Energy Economics, 145*. [article number to be confirmed]

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica, 49*(6), 1417–1426. https://doi.org/10.2307/1911408

Ottonello, P., & Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica, 88*(6), 2473–2502. https://doi.org/10.3982/ECTA15949

Pesaran, M. H. (2021). General diagnostic tests for cross-sectional dependence in panels. *Empirical Economics, 60*, 13–50. https://doi.org/10.1007/s00181-020-01875-7

Roodman, D. (2009). A note on the theme of too many instruments. *Oxford Bulletin of Economics and Statistics, 71*(1), 135–158. https://doi.org/10.1111/j.1468-0084.2008.00542.x

Schmidt, T. S., Steffen, B., Egli, F., Pahle, M., Tietjen, O., & Edenhofer, O. (2019). Adverse effects of rising interest rates on sustainable energy transitions. *Nature Sustainability, 2*(9), 879–885. https://doi.org/10.1038/s41893-019-0375-2

Serebriakova, A., Polzin, F., & Sanders, M. (2026). Monetary policy and energy installation: Implications for the European green transition. *Energy Economics, 155*, Article 109159. [DOI to be confirmed from publisher page]

Steffen, B. (2020). Estimating the cost of capital for renewable energy projects. *Energy Economics, 88*, Article 104783. https://doi.org/10.1016/j.eneco.2020.104783

Sun, C., Khan, A., Liu, Y., & Lei, N. (2022). An analysis of the impact of fiscal and monetary policy fluctuations on the disaggregated level renewable energy generation in the G7 countries. *Renewable Energy, 189*, 1154–1165. [DOI to be confirmed]

Ulug, M., & Andrei, R. (2025). Energy transition under twin shocks: Geopolitical and macrofinancial risks. *Energy Economics, 151*. [article number to be confirmed]

Webb, M. D. (2023). Reworking wild bootstrap-based inference for clustered errors. *Canadian Journal of Economics, 56*(3), 839–858. https://doi.org/10.1111/caje.12661

Appendix A. Data audit and sample-selection statement

Check	Result
Duplicate country-year keys	0
HHI recomputed from bilateral squared shares vs panel value	max absolute error 1.1×10^{-16}
Partner-sum / world-total reconciliation, 65 usable rows	1.000 in every case (tolerance 2%)
HHI outside [0, 1]	0
100·ln(C_t/C_{t−1}) recomputed vs panel value	identical
Lagged HHI recomputed without look-ahead vs panel value	identical
Imputed, interpolated or assumed-zero values	0
Policy-rate years with fewer than 12 months	Japan 2013, 2016 (4 months); 2014–2015 (none) → set to missing
Renewable-share coverage	ends 2021 for all countries
Negative net additions	9 country-years, all ≤ 21 MW in absolute value
Country-years with verified HS 850231 trade	65 of 240 (7 countries); 175 blocked by Comtrade anonymous quota (HTTP 403)

Source vintages: Comtrade downloaded 28 October 2026; IRENA capacity via archived CSV (Zenodo record 10952917, April 2024 vintage), to be checked against the current IRENASTAT release for 2022–2023 revisions; WDI and BIS from prior official downloads logged in the workbook's Source_Log sheet. The replication package (wind_mstm_pipeline_v1.zip) contains the download script that completes the trade panel with identical provenance fields, the shock-aggregation script for the IMF series, and one command that regenerates every table and figure.