

Does Carbon Co-Jump with Energy? High-Frequency Co-Jumps and Mutually Exciting Jumps in EU ETS Futures Around the 2026 Strait of Hormuz Shock

Authors

Ankit Agarwal, CFA,
Department of Management Science,
Indian Institute of Technology Kanpur, Kanpur,
Uttar Pradesh, India.
Email: ankita21@iitk.ac.in
ORCID: 0009-0006-7283-0419

Keywords: carbon markets; EU ETS; co-jumps; Hawkes processes; energy futures; Strait of Hormuz

JEL classification: C32; C58; G13; G14; Q54

Funding

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

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.

Data availability

The processed series and the full replication code reproducing all reported results are provided in Supplementary Material S1, and the price construction is documented in Appendix A of the manuscript.

Does Carbon Co-Jump with Energy? High-Frequency Co-Jumps and Mutually Exciting Jumps in EU ETS Futures Around the 2026 Strait of Hormuz Shock

Abstract

This paper asks whether, and through which channel, price jumps are transmitted between the European carbon market and the energy complex. Conceptually, we separate two mechanisms that daily data confound: a diffusive demand channel operating through fuel switching and generation margins, and a discontinuous information channel through which energy news reaches carbon within minutes. Empirically, we combine univariate jump detection, co-exceedance and common-jump tests, and bivariate Hawkes models, applied to one- and five-minute futures on EUA allowances, TTF gas, German power, Brent, and WTI over July 2025 to July 2026, a sample that brackets the 2026 Strait of Hormuz closure. Carbon almost never co-jumps with gas or power. Its only reliable discontinuous link runs through crude oil, completes within five minutes, and survives the crisis intact: cross-excitation from oil is stable across regimes, while carbon's self-excitation triples and its diffusive correlation with oil turns negative. The widely discussed 2026 decoupling was therefore real, but it was diffusive and compositional; the jump channel never closed. To our knowledge, this is the first high-frequency study of carbon-energy co-jumps and mutually exciting jumps, with direct implications for hedging carbon jump risk and for the design and surveillance of the EU ETS.

Keywords: carbon markets; EU ETS; co-jumps; Hawkes processes; energy futures; Strait of Hormuz

JEL classification: C32; C58; G13; G14; Q54

1. Introduction

A carbon allowance is a claim on the right to emit, and the textbook account of its price ties that claim tightly to the fuels whose combustion creates the demand for it. When gas prices rise relative to coal, dispatch shifts toward the dirtier fuel, emissions rise, and allowances should

appreciate; when power prices jump, the marginal abatement cost embedded in the allowance should move with them. A substantial empirical literature establishes these relationships at daily and weekly frequencies (Mansanet-Bataller, Pardo, & Valor, 2007; Chevallier, 2009; Hintermann, 2010; Bredin & Muckley, 2011; Creti, Jouvet, & Mignon, 2012), and more recent works documented the rich spillovers in prices and volatility between carbon and energy (Ji, Zhang, & Geng, 2018). The literature has not examined whether the relationship persists at the frequency at which modern price discovery actually occurs: in the seconds and minutes immediately after news arrives, when prices move discontinuously rather than drift gradually. This paper poses a simple question at that frequency. When energy jumps, does carbon jump with it?

The question bridges two literatures that remained distinct. On one side, energy economics has documented how strongly the effect of energy shocks propagates: European gas shocks transmit measurably into euro-area inflation (Chen, Wang, & Miao, 2025); regulatory interventions such as the 2022 gas price cap left TTF spikes largely unaltered (Goodell, Gurdgiev, Paltrinieri, & Piserà, 2024); and commodity-price uncertainty transmits to financial markets in a strongly regime-dependent and asymmetric (Chen, Msofe, Wang, & Chen, 2025). Conversely, financial econometrics has built sophisticated tools for detecting jumps, testing whether they arrive together, and modelling how one market's jumps excite another's. Carbon sits squarely between the two literatures, yet the intraday evidence that exists on EUA prices is entirely contained in the allowance market (Chevallier & Sévi, 2011; Conrad, Rittler, & Rotfuß, 2012; Rittler, 2012), and there is no research that addresses the question of whether the carbon-energy connection operates through jumps or not. Section 2 reviews both strands and develops our hypotheses.

The answer matters for three audiences. Compliance buyers and the growing population of financial traders in EU allowances need to know whether energy shocks propagate into their

positions instantaneously or with a lag, because only the latter can be traded against. Risk managers pricing options on EUA futures need to know whether the relevant jump risk is idiosyncratic or systematic with the energy complex, because the two imply very different hedging programmes. Policymakers, finally, read the 2026 episode, in which allowance prices fell by roughly a quarter while gas rose by roughly half and Brent traded above 130 dollars, as evidence that the carbon market had come untethered from its fundamentals. They need to know what, precisely, broke. Our answer is that nothing broke: what changed was the composition of carbon's jump risk, not the transmission mechanism.

We assemble 12 months of 1-minute and 5-minute futures data from LSEG for the five instruments that define the European carbon-energy nexus: ICE EUA allowances, TTF natural gas, German baseload power, Brent, and WTI. The sample, July 2025 through July 2026, brackets the Strait of Hormuz crisis, which began with military strikes on 28 February 2026, saw the strait closed to tanker traffic in early March, pushed Brent to an intraday peak near 138 dollars on 7 April, and wound down with a memorandum of understanding signed on 17 June. The episode comes close to an ideal natural experiment for our purposes: a large, exogenous, oil-first supply shock whose arrival date is known to the minute and that splits the sample into 170 trading days of a calm regime and 77 days of a conflict regime, over which every transmission channel can be re-estimated.

The empirical design proceeds in four stages, each standard on its own, combined here for the first time in a carbon market setting. We first detect jumps asset by asset with the Lee and Mykland (2008) statistic applied to periodicity-filtered returns following Boudt, Croux, and Laurent (2011), cross-checked with daily bipower and median-based tests (Barndorff-Nielsen & Shephard, 2004, 2006; Huang & Tauchen, 2005; Andersen, Dobrev, & Schaumburg, 2012). We then test whether

jumps coincide across assets, using co-exceedance counts in the manner of Gilder, Shackleton, and Taylor (2014), extreme cross-products following Bollerslev, Law, and Tauchen (2008), the common-versus-disjoint ratio of Jacod and Todorov (2009), and a systemic count in the spirit of Caporin, Kolokolov, and Reno (2017). Third, we ask who excites whom: we estimate bivariate mutually exciting point processes (Hawkes, 1971) on one-minute jump times, following Ait-Sahalia, Cacho-Diaz, and Laeven (2015), with inference in the spirit of Boswijk, Laeven, and Yang (2018), and we complement the likelihood machinery with a transparent nonparametric measure, the probability that a carbon jump follows an energy jump within fifteen minutes, benchmarked against a Poisson null. Finally, we exploit the crisis itself by conducting an event study around nine date escalations, as well as a difference-in-differences analysis, using the Brent-WTI pair, two claims on almost the same barrel, which serves as a placebo for what a truly coupled pair would do during a crisis. Four hypotheses, fixed before estimation and developed formally in Section 2.4, discipline the exercise: carbon should co-jump mainly with its fundamental neighbours, gas and power (H1); carbon should receive rather than send jumps (H2); energy-to-carbon transmission should have weakened during the crisis (H3); and co-jumps should have a distinctive microstructure signature (H4).

The evidence challenges our prior assumptions. Carbon essentially never co-jumps with gas: only 1 coincidence in the entire year, compared to 1.6 expected under independence. The fuel-switching channel is so firmly established in the day-to-day literature, and it does not operate at jump frequency. Carbon's only reliable high-frequency link runs through crude. Brent and carbon have 6 common 5-minute jumps versus the expected 1.05, a ratio of 5.7 with a Monte Carlo p-value of 0.002, and the Jacod-Todorov ratio for the two is 1.003, almost exactly what the theoretical value predicts for true co-jumps. The Hawkes estimates confirm the dynamic link: the

branching ratio from Brent to carbon is 0.037 in the baseline and 0.041 in the crisis, both significant in their respective regimes, while the corresponding loadings from gas and power are statistically zero under all conditions. The evidence of timings suggests the transmission within five minutes and it would be re-estimated on a five-minute event grid, the apparent excitation disappears.

H3 fails in an instructive way. The oil-to-carbon jump channel did not weaken during the Hormuz crisis, and it even became stronger with carbon-oil co-jump intensity rising nearly sixfold relative to the Brent-WTI placebo. Everything that changed happened around the channel. Carbon's own jump intensity increased by a third, and its self-excitation tripled, causing the share of carbon's jump risk attributable to energy to fall even as the absolute channel remained the same. Meanwhile, carbon's diffusive five-minute correlation with Brent reversed direction, from +0.105 to -0.097, as the demand channel took over: the crisis crushed gas-fired generation profitability, with the clean spark spread falling from +1.3 to -34 euros per megawatt hour, cutting the marginal demand for allowances precisely when oil was screaming higher. The decoupling, in short, was real but compositional and diffusive. The discontinuous information channel from crude was not destroyed, but remained intact.

Because the co-jump counts at the core of this claim are modest, we apply the wide range of robustness tests that a sceptical reader might ask of such a claim. We allow the data to identify the structural break date rather than the military timeline, and no jump series rejects a no-break null once permutation inference accounts for persistence. However, the best-fitting break for carbon-oil co-jumps lands on 20 March 2026 with a confidence interval that encompasses the invasion. We replace the binary regime with a continuous dose, the OVX oil-volatility index, and find a steep response of carbon-oil co-jumps whose Newey-West significance dissolves under the same permutation discipline, a caution we report rather than bury. We decompose realised

semicovariance following Bollerslev, Li, Patton, and Quaadvlieg (2020) and find that the concordant share of carbon-crude dependence falls from three-fifths to two-fifths across the break, providing a moment-level image of the sign flip in correlations. We measure jump-count connectedness following Diebold and Yilmaz (2012) and find the system's total connectedness rising from 17.4 to 27.4 per cent during the conflict, with carbon's dependence on shocks from the rest of the complex more than quadrupling. Finally, we run the event window against every alternative window of the same length the sample allows: only 7 per cent of placebos match the true window's carbon-oil co-jump count, while an exact power calculation shows the design could not have detected rate increases below a factor of 6.7, which is why our reading of stability rests on the coherence of many tests rather than the significance of any one.

The paper makes three contributions. First, to our knowledge, it provides the first study of co-jumps and mutually exciting jumps between carbon and the energy complex at high frequency; the intraday carbon literature to date has treated EUA jumps univariately, and the daily spillover literature cannot resolve discontinuities at all. Second, the Hormuz episode allows us to separate two transmission channels that daily data confound, a diffusive demand channel that turned negative and a discontinuous information channel that did not move, and the separation reverses the natural reading of the 2026 price action. Third, along the way, we document features of independent interest to the jump literature: gas jump intensity fell during a gas price shock of historic size, so the crisis moved gas smoothly and oil discontinuously; and calm-period carbon jumps cluster at a daily rather than a minute timescale, a separation with direct consequences for how Hawkes models of carbon should be specified.

The paper is structured as follows. Section 2 surveys the literature and develops the four hypotheses, contextualising the 2026 shock within the carbon-energy economics framework:

Section 3 details the data, the price construction for quote-driven contracts, and the descriptive evidence. Section 4 presents the econometric framework. Section 5 summarises the results in the order of the four stages. Section 6 addresses robustness, extensions, and threats to identification, and Section 7 concludes with policy implications. Appendix A records the implementation details required for replication, and Supplementary Material S1 provides the replication code, along with the derived series for each exhibit.

2. Related Literature and Hypothesis Development

2.1. Carbon-energy linkages: the lower-frequency evidence

There is a documented relationship between allowance prices and energy prices. From an equilibrium perspective, the allowance price equals the marginal cost of abatement. Historically, fuel price differentials have been most significant in the European power sector, where switching from coal to gas has been the main focus of emission-reduction efforts. At the daily level and weekly level this relationship was first validated by Mansanet-Bataller, Pardo and Valor (2007), who found that energy prices and weather conditions are the main drivers of EUA returns; Hintermann (2010) demonstrates that fuel prices explain allowance prices after accounting for the Phase I collapse; Chevallier (2009) links carbon futures to macroeconomic fundamentals; and Bredin and Muckley (2011) and Creti, Jouvét, and Mignon (2012) document an increasingly mature, fundamentals-driven equilibrium across trading phases. Subsequent research analysed the bidirectional relationship, with Ji, Zhang, and Geng (2018) documenting substantial information transmission and dynamic spillovers in prices and volatility between the carbon and energy markets.

Recent literature investigates the propagation of energy shocks under stress, and three results are directly relevant to our design. Chen, Wang, and Miao (2025) demonstrate, in a time-varying structural framework, that European gas shocks pass through to euro-area inflation, and that these impacts are mostly at short and medium horizons, making gas repricing macroeconomically material rather than a sectoral one. Goodell, Gurdgiev, Paltrinieri, and Piserà (2024) demonstrate that the EU's December 2022 gas price cap had no quantifiable impact on TTF price spikes or volatility, a cautionary result for regulatory measures intended to mitigate energy price discontinuities. Furthermore, Chen, Msofe, Wang, and Chen (2025) find that the impact of oil-price uncertainty on financial markets is very regime-dependent, particularly in volatile quantiles and much weaker in calm quantiles. Closest to our concerns, Ferrara, Karadimitropoulou, and Triantafyllou (2024) estimate oil jump-tail risk from high-frequency futures prices and demonstrate that the discontinuous component of oil, rather than its smooth volatility, carries distinct information about inflation dynamics. Together, these papers establish that energy shocks are large, hard to regulate away, transmitted asymmetrically across market states, and economically distinct in their discontinuous component. What none of them observes, because none uses intraday data on carbon, is whether the transmission into the allowance market operates through the discontinuous component of prices at all. That is the layer this paper adds.

2.2. Jumps, co-jumps, and mutually exciting processes

The econometrics of discontinuities supplies the tools. Univariate jump detection builds on the bipower variation theory of Barndorff-Nielsen and Shephard (2004, 2006), implemented as daily ratio tests by Huang and Tauchen (2005), refined against noise, periodicity, and finite-sample distortions by Andersen, Bollerslev, and Dobrev (2007), Corsi, Pirino, and Reno (2010), and Andersen, Dobrev, and Schaumburg (2012), and formalised as a general testing problem by Ait-

Sahalia and Jacod (2009). At intraday frequency, the Lee and Mykland (2008) statistic identifies individual jumps, and Boudt, Croux, and Laurent (2011) demonstrate that accounting for the pronounced intraday periodicity of volatility is essential to avoid false positives at the market open. Bandi and Reno (2016) and Bibinger and Winkelmann (2015) document how price-volatility co-jumps and microstructure noise complicate inference, which justifies the conservative test sizes we adopt.

Whether jumps arrive together across assets is a separate question with its own set of literature. Bollerslev, Law, and Tauchen (2008) study co-jumps through the extremes of the cross-products and relate them to systematic risk; Jacod and Todorov (2009) furnish the formal theory for testing simultaneous against separate jump arrivals; Gilder, Shackleton, and Taylor (2014) lists co-exceedances of univariate detections and delineate co-jumps in equity markets; and Caporin, Kolokolov and Reno (2017) show that multi-asset co-jumps contain information which is not captured by the pair-wise statistics. A supplementary dynamic view treats jump times as a point process: Hawkes (1971) processes permit one market's jumps to raise the intensity of another's, Ait-Sahalia, Cacho-Diaz, and Laeven (2015) use mutually exciting jumps to model financial contagion, and Boswijk, Laeven, and Yang (2018) develop likelihood-based tests for self-excitation. Bollerslev, Li, Patton, and Quaadvlieg (2020) decompose realised covariation by the signs of the underlying returns, and Diebold and Yilmaz (2012) quantify system-wide connectedness; we employ both as extensions. In commodity settings specifically, Zhou, Bagnarosa, Gohin, Pennings, and Debie (2023) demonstrates that high-frequency price discovery within the soybean complex is governed by liquidity and microstructure frictions, a lesson our treatment of the quote-driven gas and power books incorporates. None of this machinery, however, has been brought to the carbon-energy nexus at high frequency. The intraday evidence on the

allowance market is confined to carbon itself: Chevallier and Sévi (2011) characterise the realised volatility and jump component of ECX carbon futures, Conrad, Rittler, and Rotfuß (2012) trace the minute-scale response of EUA prices to scheduled regulatory news, and Rittler (2012) measures price discovery between allowance spot and futures. The spillover literature that does reach into energy, meanwhile, operates on daily aggregates, in which jumps and diffusion are inseparable. Our contribution is to combine the full toolkit, detection, coincidence, excitation, and a natural experiment, on the five markets that define the nexus.

2.3. The 2026 Hormuz shock as an identification device

Roughly a fifth of globally traded oil and a comparable share of liquefied natural gas transit the Strait of Hormuz, so the closure that followed the strikes of 28 February 2026 was, mechanically, a shock to seaborne energy supply of a size the market had war-gamed for decades but never observed. Escalation trajectory is key to our methodology, and is established from contemporaneous reporting: strikes on 28 February, the strait declared closed on 2 March, Brent peak price of over 100 dollars on 8 March, an ultimatum on 21 March, the Brent peak near 138 on 7 April followed by a ceasefire on 8 April, a re-closure threat on 18 April, and a de-escalation culminating preliminary talks on 14 June and a signed agreement on 17 June. We designate 28 February 2026 as the structural break and 17 June as the end of the conflict regime; these nine dates comprise the event chronology of Section 5.4.

The shock propagated through European energy in two different ways. Crude repriced immediately and violently because Hormuz controls position over the supply chain. Gas repriced through the LNG channel: Qatari cargoes traverse the strait, and the marginal supply unit in the winter and shoulder seasons is maritime. TTF rose roughly half over the conflict window in our sample, from a baseline average of 31.6 euros per megawatt hour to 47.7 euros per megawatt hour.

In contrast, Brent averaged 100.1 dollars per barrel, up from 65.9 dollars previously. German baseload power followed its fuel. Carbon fell: front-December EUA futures traded near ninety-two euros in January 2026, touched the mid-sixties in March, and recovered only to the mid-seventies by June.

The direction of carbon's movement is less puzzling than it first appears. The allowance is a derived claim, and the sign of its exposure to a gas shock depends on which margin binds. If coal-to-gas switching binds, dearer gas raises emissions and allowance demand; industry estimates circulated in March 2026 put the implied lift in 2026 power-sector emissions at around 4 per cent, on the order of 20 million tonnes of additional allowance demand. If instead the shock destroys demand for power and industrial output, or crushes gas-fired generation margins to the point where the gas plant, the marginal buyer of the allowance-gas bundle, does not run, allowance demand falls. Our clean spark spread series, constructed at 50 per cent efficiency and an emission factor of 0.37 tonnes per megawatt-hour, moved from +1.3 euros in the baseline to -34.0 euros during the conflict. The bundle that transmits gas prices into carbon demand was deeply out of the money for most of the crisis, which is the fundamental reason a gas shock of historic size coincided with a falling allowance price. What daily data cannot reveal is whether this demand-side repricing also silenced the high-frequency information channel, the near-instantaneous propagation of discontinuous news from energy into carbon. Separating the two is the business of this paper.

2.4. Hypothesis development

On the information channel, the natural prior is a chain running from crude, the global shock carrier, through gas via oil-indexed contracts and LNG arbitrage, into power through the merit order, and finally into carbon through expected emissions. Located at the terminal node, carbon should receive rather than send, and its closest discontinuous neighbours are the domestic

components, gas and power, whose connection to allowance demand is mechanical. Four testable hypotheses, fixed prior to the estimation process, derive from this prior and from the microstructure literature.

Hypothesis 1 (H1). In calm markets, carbon experiences concurrent jumps with its fundamental counterparts, gas and power, more frequently than with crude. The fuel-switching mechanism directly ties allowance demand to the gas and power complex, and if the lower-frequency evidence in Section 2.1 has an intraday counterpart, sharp price movements in gas and power should provide the information relevant to allowances.

Hypothesis 2 (H2). Energy jumps lead to carbon jumps, rather than the reverse, in which carbon is the receiver. The allowance occupies the terminal node of the oil-to-gas-to-power-to-carbon sequence and originates little global news of its own outside issuance and policy events, so cross-excitation should flow from energy into carbon, and who-moved-first statistics should be biased towards the energy components.

Hypothesis 3 (H3). During the 2026 crisis, energy-to-carbon transmission declined. This is the prediction the price dynamics seemed to invite: allowance prices declined by a quarter while gas and oil were rising. One interpretation is that the intraday transmission mechanism is also dampened, an intraday counterpart of the daily decoupling visible in levels. H3 is the hypothesis that the methodology directly tested a placebo and its failure, as recorded in Section 5, is the key finding of the paper.

Hypothesis 4 (H4). Co-jumps have a unique microstructure signature, with higher volume, more trades, and wider quoted spreads relative to carbon's idiosyncratic jumps. If co-jumps lead to the arrival of market-wide information that many participants must trade on simultaneously,

while idiosyncratic jumps represent market-specific order flow, then the order-flow patterns of these two should differ, as they do in equity markets.

Each hypothesis is tested against its own placebo or benchmark. H1 against independence nulls that preserve within-day clustering, H2 against Poisson benchmarks and reverse loadings, H3 against the Brent-WTI placebo pair, and H4 against carbon's own idiosyncratic jumps. As it transpires, the pattern of rejections is more informative than the priors themselves.

3. Data and Descriptive Evidence

3.1. Instruments, fields, and sample

We retrieve intraday bars from LSEG Workspace for 5 near-term perpetual futures: ICE EUA carbon (RIC CFI2Zc1), ICE Endex TTF natural gas (TRNLTTFMc1), EEX German baseload power (TRDEBMc1), ICE Brent (LCOc1), and NYMEX WTI (CLc1). Each observation includes the open, high, low, and last trade price, trading volume, the best bid and ask quotes, and the number of price moves. We use one-minute bars for the point-process analysis, five-minute bars for the headline jump and co-jump tests, and thirty-minute bars for robustness; LSEG does not supply a fifteen-minute bar. The sample runs from 7 July 2025 to 6 July 2026. All sixteen source data files, together with full field definitions, are provided in Supplementary Material S1.

The five books differ radically in how they trade, and respecting those differences is the first substantive modelling choice. Carbon and the two crude contracts are trade-driven: last-trade prices are present in essentially every bar and move in ninety per cent or more of them. TTF and German power on the continuation composites are quote-driven and one-sided. For power, a trade price exists in only 38 per cent of five-minute bars while a bid exists in 85 per cent; for TTF, the figures are 64 and 76 per cent, and the ask side is populated barely half the time for either, which

rules out mid-quotes. We therefore price each asset from the source with the greatest moving-price coverage: the last trade for carbon, Brent, and WTI, and the bid quote for TTF and power, where bid coverage of price changes (66 and 58 per cent of bars) is roughly triple the trade coverage for power and half again as high for TTF. Pricing an asset off one side of the book embeds liquidity variation in its returns, a cost we return to in Section 6; the alternative, stale trade prints, is worse, and was responsible for near-zero detected jump counts in preliminary work.

The cleaning follows standard practice, adapted to sparsely populated books. A complete regular grid is created throughout the trading day, prices are carried forward throughout the day by at most sixty minutes, and a quotation that is more than 25 log per cent away from its own daily median is filtered out, which results in just four anomalous observations over a twelve months, one instance being a TTF bid of 5.0 printed against an ask of 32.53. Returns are first differences of log prices, with the first bar of each day dropped. Because the continuation series splice contracts between sessions, roll gaps do not affect the return series, and no additional roll adjustment is necessary. The common intraday window is selected from the data rather than imposed: we determine the time-of-day slots where at least 60 per cent of assets showed a valid price, which is 07:05 to 16:30 UTC, 114 five-minute slots per day, essentially the EUA session. A day is included in the sample for an asset if at least 50 per cent of its window bars are usable. Full details of the price construction and cleaning are given in Appendix A.1, and the complete routine, including the price-source selection rule and despiking filter, is implemented in the master script in Supplementary Material S1.

Volume and trade counts are unavailable for the TTF and power composites, so the microstructure tests of Section 5.5 condition entirely on carbon-side volume, trade counts, and quoted spreads. The regime partition allocates 170 common-panel trading days to the baseline (7

July 2025 to 27 February 2026), 77 to the conflict window (28 February to 17 June 2026), and 13 days to the short post period, which we tabulate but do not test.

3.2. Descriptive statistics

Table 1 reports the moments of five-minute returns, the correlation matrices for each regime, and the correlations of the jump indicators as defined in Section 4. Prior to any formal test, three features should be noted. First, the crude series reflects the crisis in its tails: Brent has excess kurtosis of 560 and skewness of -9.4, with large negative five-minute returns, the largest being -13.1 per cent, recorded around the ceasefire and de-escalation dates. Carbon is much less pronounced, with a kurtosis of 41 and extremes of around 3 per cent. Second, diffusive dependence between carbon and anything remains weak at this frequency. The overall correlation between carbon and Brent is -0.02, with a notable regime shift: carbon-Brent shifts from +0.105 in the baseline to -0.097 in the conflict, a sign reversal that we explain through the demand channel in Section 2. At the same time, carbon gas sits at zero in both regimes. Sampling error does not account for this: the Brent-WTI benchmark holds above 0.91 throughout. Part of the low level reflects the Epps (1979) effect for non-synchronously traded assets, which is why Panel C also reports daily-frequency correlations. At the daily horizon, carbon-Brent is slightly negative in the baseline at -0.058 and widens to -0.294 in the conflict, indicating that the two frequencies continue to agree on the direction of travel, but differ on the calm-period level, a horizon structure consistent with attenuation at five minutes. Third, and ahead of everything that follows, the correlation of five-minute jump indicators tells the story in dichotomous format: 0.654 for Brent-WTI, 0.028 for carbon-Brent, and effectively zero for carbon-gas.

TABLE 1 Descriptive statistics and dependence

	N	Mean (bp)	SD (bp)	Skewness	Kurtosis	Min (%)	Max (%)			
<i>Panel A: Five-minute return moments, full sample</i>										
EUA carbon	28,895	0.08	14.3	-0.02	40.6	-3.57	3.18			
TTF gas	24,163	-0.35	46.5	-0.67	104.9	-9.70	10.95			
DE power	22,783	0.36	54.4	-0.65	102.7	-15.42	9.56			
Brent	29,121	-0.10	20.9	-9.36	559.6	-13.13	2.66			
WTI	28,968	-0.02	22.4	-8.15	467.4	-13.45	3.10			
	Carbon		TTF		Power		Brent		WTI	
<i>Panel B: Five-minute return correlations, baseline (below) and conflict (above diagonal)</i>										
EUA carbon	1		0.009		0.015		-0.097		-0.090	
TTF gas	0.031		1		0.022		0.071		0.066	
DE power	0.051		-0.006		1		0.066		0.051	
Brent	0.105		0.020		0.043		1		0.915	
WTI	0.101		0.026		0.045		0.940		1	
	Carbon		TTF		Power		Brent		WTI	
<i>Panel C: Daily return correlations, baseline (below) and conflict (above diagonal)</i>										
EUA carbon	1		-0.031		0.084		-0.294		-0.252	
TTF gas	0.023		1		0.660		0.648		0.661	
DE power	0.108		0.531		1		0.416		0.383	
Brent	-0.058		0.300		0.195		1		0.911	
WTI	-0.054		0.285		0.178		0.982		1	

Note: Five-minute intraday log returns across the common window 07:05 to 16:30 UTC, 7 July 2025 to 6 July 2026, and daily close-to-close returns over the same span. Prices are last trades for carbon, Brent, and WTI and bid quotations for TTF gas and German power (Section 3.1). Means and standard deviations are in basis points; minimum and maximum values in per cent. In Panels B and C, entries below the diagonal represent the baseline regime (7 July 2025 to 27 February 2026), and entries above the diagonal represent the conflict regime (28 February to 17 June 2026). The full-sample five-minute carbon-Brent correlation is -0.022. Phi coefficient correlations of the five-minute jump detections produced by the Lee-Mykland test of Section 4.1 are 0.654 for Brent-WTI and 0.028 for carbon-Brent, with all remaining carbon pairings within 0.02 of zero. Data source: LSEG Workspace.

Figure 1 displays the standardised daily price series, the 22-day rolling correlation between daily carbon returns and each energy component, and the clean spark spread, with the onset of the

war highlighted. The visual pattern of decoupling is unambiguous at the daily level: the rolling correlation of carbon and crude is mildly positive before February, shifts to roughly -0.4 through April and May, and the spark spread collapse dates the demand destruction that resulted from the escalation. The formal analysis begins with nothing provided in the figure, which deals with price discontinuities.



FIGURE 1 Market backdrop: prices, rolling correlations, and the clean spark spread

Note: The upper panel shows the normalised prices for the five contracts in the common sample. The middle panel shows the 22-day rolling correlation between daily carbon returns with each energy component. The bottom panel depicts the clean spark spread, computed as the German power price minus the TTF gas price divided by 0.50 (the assumed efficiency), minus 0.37 times the EUA price. The shaded band shows the period of conflict (28 February through 17 June 2026), and the dotted vertical lines mark the nine event dates listed in Table 5. Average daily by regime: carbon 78.4 baseline versus 74.2 conflict; Brent 65.9 versus 100.1; TTF 31.6 versus 47.7; spark spread +1.3 versus -34.0

4. Econometric Framework

4.1. Univariate jump detection

Let $p(t)$ denote the log price of an asset and $r(i)$ its i -th within-day return on a grid of M bars per day. Intraday volatility in these markets follows a pronounced deterministic pattern and, if left untreated, inflates jump detections at the open and starves them mid-session. We first filter returns using the Boudt, Croux, and Laurent (2011) periodicity factor. A daily bipower volatility standardises returns, a robust scale is computed for each time-of-day slot s as the cross-day median of absolute standardised returns, and the resulting factor $f(s)$ is normalised so that its squared values average one:

$$f(s) = \text{med}_d \left| \frac{r_{d,s}}{\sigma_d} \right|, \quad M^{-1} \sum_s f(s)^2 = 1, \quad (1)$$

with σ_d the square root of day- d realised bipower variation. Filtered returns are $\tilde{r}(i) = r(i)/f(s(i))$. On these, we apply the Lee and Mykland (2008) statistic. The local volatility estimate is a backward bipower average over a window of $K = 20$ filtered returns,

$$\hat{\sigma}_i^2 = (K - 2)^{-1} \sum_{j=i-K+2}^{i-1} |\tilde{r}_j| |\tilde{r}_{j-1}|, \quad L(i) = \frac{\tilde{r}_i}{\hat{\sigma}_i}, \quad (2)$$

and bar i hosts a jump when $|L(i)|$ exceeds the Gumbel threshold $C(n) + S(n)\beta^*$, with n the median number of bars per day and $\beta^* = -\log(-\log(1 - \alpha))$ at $\alpha = 0.01$, where

$$C(n) = \frac{(2 \log n)^{1/2}}{c} - \frac{\log \pi + \log \log n}{2c (2 \log n)^{1/2}}, \quad S(n) = \frac{1}{c (2 \log n)^{1/2}}, \quad c = (2/\pi)^{1/2}.$$

The effective threshold is above about $|L| = 5.17$, with $n \sim 114$. Two procedural safeguards protect the quote-driven series. Null products of absolute returns, the signatures of prices propagated through inactive bars, are removed from the calculation in equation (2), and a local estimate is formed only when at least eight valid products remain; otherwise, the bar is not tested. Without the safeguard, the dormant interval depresses the local measure and yield false detections when the first genuine revision of the quote, which is exactly what a naive implementation produces for the power series. As validation tests, we run the daily ratio tests of Barndorff-Nielsen and Shephard (2004, 2006) in the Huang and Tauchen (2005) form and the median-truncated variant of Andersen, Dobrev, and Schaumburg (2012), both derived directly from the filtered returns; the noise-robustness concerns raised by Bandi and Reno (2016) and Bibinger and Winkelmann (2015) justify the stringent one per cent size. Detection settings, resampling parameters, and the Monte Carlo protocol are reported in Appendix A.2.

4.2. Co-jump tests

We use the co-exceedance count of Gilder, Shackleton, and Taylor (2014) as our workhorse. For assets a and b with jump indicators $J_a(i)$ and $J_b(i)$ on N common bars, the statistic is the number of bars flagged by both,

$$CJ(a,b) = \sum_i J_a(i) J_b(i), \quad E[CJ] = N p_a p_b, \quad (3)$$

where p_a and p_b are the marginal jump rates. We report an exact binomial p-value against the independence benchmark and, because jump indicators are serially and diurnally clustered, a Monte Carlo p-value that preserves both: within each day, one asset's indicator vector is circularly shifted by a uniform random offset, destroying cross-asset synchrony while leaving each day's internal clustering intact, and the co-count distribution is built from five hundred replications. The same day-preserving shift generates the null for a systemic count, the number of bars in which at

least three of the five assets jump simultaneously, in the spirit of the panel co-jump analysis of Caporin, Kolokolov, and Reno (2017). Two further statistics triangulate. Following Bollerslev, Law, and Tauchen (2008), we count extreme standardised cross-products $r_a(i) r_b(i)$ that exceed their 99.9th percentile under the shifted null. Following Jacod and Todorov (2009), we compute the common-versus-disjoint ratio.

$$\Phi(k) = \frac{V(k\Delta)}{V(\Delta)}, \quad V(\Delta) = \sum_i [r_a(i) r_b(i)]^2, \quad (4)$$

where the numerator aggregates non-overlapping blocks of k bars within days. When the two assets jump at the same time, the dominant terms of V are invariant to the block size, and Φ tends to 1; when jumps are disjoint, each product pairs a jump with a diffusion, and Φ tends to k . The full asymptotic variance theory of Jacod and Todorov (2009) and Caporin, Kolokolov, and Reno (2017) is not implemented in available software and hand-coding it invites error, so we state plainly that our inference is by resampling: day-block bootstrap intervals for Φ and a desynchronised null, built from the same within-day shifts, for the test that jumps are disjoint. The point statistics are the canonical ones; only the critical values are simulated.

4.3. Mutually exciting jumps

Coincidence within a bar is one notion of connection; excitation across bars is another, and for the receiver hypothesis, it is the sharper one. Let $t_1 < t_2 < \dots$ denote the one-minute jump times of a carbon-energy pair, mapped to business time as the cumulative count of trading minutes, so that overnight gaps do not stretch the kernel. We model the pair as a bivariate Hawkes (1971) process with exponential kernels,

$$\lambda_m(t) = \mu_m + \sum_{k: t_k < t} \alpha_{m, c(k)} e^{-\beta_m(t - t_k)}, \quad m \in \{\text{carbon}, \text{energy}\}, \quad (5)$$

where $c(k)$ is the component of the k -th event. The object of interest is the branching matrix $n = \alpha/\beta$, whose (m, j) entry counts the expected number of type- m events directly triggered by one type- j event; $n(\text{carbon}, \text{energy})$ is the transmission mass of H2 and H3. We estimate by maximum likelihood over the exact log-likelihood with its exponential-kernel compensator, using the recursive intensity representation, and we test the null of no energy-to-carbon excitation by a likelihood ratio against the model with $\alpha(\text{carbon}, \text{energy}) = 0$. Restricted and unrestricted fits warm-start one another, which guarantees a non-negative likelihood ratio and, in practice, removes two convergence artefacts produced by a single-start optimiser. Our approach follows the mutually exciting jump framework developed by Ait-Sahalia, Cacho-Diaz, and Laeven (2015) for return contagion, with inference in the spirit of Boswijk, Laeven, and Yang's (2018) likelihood-ratio approach. The likelihood, the multi-start estimation protocol and the construction of the likelihood-ratio tests are detailed in Appendix A.3.

One specification issue is substantive enough to preview here. Estimated freely, the decay β is weakly identified in the calm regime, because baseline carbon jumps cluster at the scale of days, several detections on a news day, quiet weeks between, and a single exponential kernel gladly latches onto that slow timescale, absorbing the clustering into near-permanent excitation and leaving minute-scale cross-effects on a flat ridge of the likelihood. The instability is informative, and we report it, but the hypotheses concern minute-scale transmission. The primary specification, therefore, fixes the kernel at the microstructure scale, $\beta = 20$ per hour for both components, a half-life of 2.1 minutes, and re-estimates baselines and loadings regime-by-regime on identical kernels; a grid of $\beta \in \{10, 20, 40\}$ accompanies it as a sensitivity check. Likelihood-ratio tests remain valid for nested comparisons at every fixed β . In parallel with the parametric machinery, we report a statistic that requires no optimiser: the share of energy jumps followed by a carbon jump within

fifteen minutes, against the Poisson benchmark $1 - \exp(-\lambda_{cm})$ with λ_c the carbon event rate, with permutation p-values from circular shifts of the carbon event times, together with simple who-moved-first tallies within the same window. To complete the picture, we present lead-lag cross-correlograms of one-minute returns, following Hayashi and Yoshida (2005).

4.4. *The crisis as an experiment*

The final stage is to use the Hormuz timeline directly. Around each of the nine dated escalations, we take a window of two trading days on either side and compare carbon jump intensity, volume, and quoted spreads against same-regime non-event days, an important design point since comparison with the calm baseline would just be restating the regime shift; empirical percentiles yield two-sided p-values adjusted across the nine events by the Benjamini-Hochberg step-up. The regime comparison itself is a difference-in-differences on daily co-jump counts,

$$E[CJ_{p,d}] = \exp(a + bTreated_p + gConflict_d + \delta Treated_p \times Conflict_d), \quad (6)$$

estimated by Poisson regression with standard errors clustered by day. The treated pair is carbon with its energy partners, in two variants, gas plus power and crude, and the placebo is Brent-WTI, two claims on nearly the same barrel, whose coupling should respond to a supply shock, if anything, by tightening. The interaction delta measures whether carbon's coupling moved differently from the placebo's, which is the cleanest available statement of H3. Liu, Patton, and Sheppard (2015) motivate our preference for simple, loss-based regime comparisons over more complex multivariate volatility machinery; we deliberately keep dynamic conditional correlation and vector autoregressive apparatus out of a paper about discontinuities.

5. Results

5.1. *The anatomy of jumps, asset by asset*

Table 2 reports detections at the five-minute frequency by asset and regime. The counts are economically sensible at the level: carbon jumps on roughly two days in three in the baseline, with 0.665 detections per day and a mean absolute size of 0.51 per cent, figures in the neighbourhood of what daily-frequency studies of EUA discontinuities would lead one to expect. The information is in the regime transition. In the conflict, carbon's intensity is raised to 0.908 per day, and its jump magnitude is increased to 0.77 per cent, which is three times as frequent and half as large. Crude transforms: Brent makes average jumps of 0.536 to 0.896 times jumps per day, the mean absolute jump nearly triples to 1.33 per cent, and 64 per cent of conflict-period jumps are negative, the characteristic of ceasefire relief and de-escalation cascades rather than escalation.

Gas presents a surprise. TTF jump intensity declines during the crisis, from 1.21 to 0.914 per day, while the gas price level rises by roughly half. The crisis moved gas violently but smoothly: a persistent repricing of the LNG supply curve expressed through drift and diffusion rather than discrete revisions. Oil, on the other hand, responded to events. The asymmetry anticipates the co-jump, since an asset that stops jumping cannot co-jump, and cautions against reading level moves as evidence of discontinuities. Figure 2 traces the weekly intensities, making the disparity readily apparent. Power occupies an intermediate position between the two regimes, with approximately one jump per day in both regimes and a magnitude near 2 per cent, but its bid-quote pricing means that some detections capture discrete book adjustments rather than trades, a concern we discuss in Section 6.

A complementary daily analysis comes to a similar conclusion. Panel B implements the Barndorff-Nielsen and Shephard ratio test day by day: the proportion of carbon trading days exhibiting a significant jump rises from 34.7 to 46.1 per cent across the break, and the share of carbon's realised

variance due to the discontinuous component rises from 10.4 to 14.3 per cent, so approximately one seventh of carbon's crisis variance arrived discontinuously. The gas and power rows register much higher numbers, above-80-per-cent jump days, but those legs are priced from one side of the book, and their daily statistics absorb quote-driven noise; we read Panel B for the trade-priced contracts and rely on the guarded intraday methodology everywhere else.

TABLE 2 Jump detections by asset and regime, five-minute frequency

Asset	Regime	Days	Jumps	Per day	Mean size (%)	Negative (%)
<i>Panel A: Lee-Mykland detections at the one per cent level</i>						
EUA carbon	Baseline	167	111	0.665	0.513	47.7
	Conflict	76	69	0.908	0.773	37.7
TTF gas	Baseline	162	196	1.210	1.379	56.6
	Conflict	58	53	0.914	1.293	60.4
DE power	Baseline	135	126	0.933	2.006	49.2
	Conflict	72	78	1.083	1.803	61.5
Brent	Baseline	168	90	0.536	0.495	55.6
	Conflict	77	69	0.896	1.328	63.8
WTI	Baseline	167	91	0.545	0.553	61.5
	Conflict	77	65	0.844	1.384	60.0

Asset	Jump days, base (%)	Jump days, conf (%)	RV share, base (%)	RV share, conf (%)
<i>Panel B: Daily BNS jump statistics</i>				
EUA carbon	34.7	46.1	10.4	14.3
TTF gas	79.0	82.8	40.1	48.1
DE power	85.2	93.1	60.4	67.9
Brent	25.0	22.1	8.7	9.2
WTI	29.9	24.7	11.0	9.6

Note: Panel A: Lee-Mykland detections at the one per cent level (periodicity-filtered five-minute returns - Section 4.1), broken by regime. Days counts the asset's valid trading days in the common panel; jumps the number of flagged bars; mean |size| the average absolute return at flagged bars in per cent; negative the share of flagged returns below zero. The 13-day post period after 17 June 2026 is omitted for brevity: carbon records 9 detections there, gas records 2, power records none, and each crude series records 4. Full-sample totals across all regimes are 189 (carbon), 251 (gas), 204 (power), 163 (Brent), and 160 (WTI); at the one-minute frequency used for the point-process analysis, the corresponding totals are 484, 480, 307, 387, and 386. Panel B: daily ratio tests of Barndorff-Nielsen and Shephard (2004, 2006) on the same filtered returns; the jump days are the share of days on which the test rejects at five per cent, and the RV share is the mean share of realised variance attributable to jumps, computed as the mean of $\max(0, 1 \text{ minus bipower over realised variance})$. The gas and power rows inherit bid-quote noise and should be interpreted with the caveats of Section 6.




FIGURE 2 Jump intensity through time

Note: Weekly Lee-Mykland detection counts at the 5-minute frequency, by contract. The shaded band indicates the conflict regime, and the dotted lines represent nine event dates. Carbon and crude intensities increase across the break; TTF gas, whose price level increases by about half, jumps less often during the conflict; and the German power detections arrive in concentrated episodes on both sides of the break.

5.2. Do jumps coincide? Co-exceedances and common-jump ratios

Table 3 is the basis for the cross-sectional evidence. Panel A displays the full-sample co-exceedance counts with binomial. Monte Carlo p-values, Panel B splits carbon pairs by regime,

and Panel C shows the Jacod-Todorov ratios. Initiate with the diagnostic check. Brent and WTI, the pair, any credible analysis must identify, display 106 5-minute co-jumps against 0.89 expected value under independence, the ratio of 119 corresponds to p-values that are close to 0 under both nulls. The procedure can find actual links if present.

By that standard, carbon's row reads nearly as a ranking of unexpected findings. Carbon and gas experience coincide exactly once in twelve months against 1.59 expected: fewer coincidences than chance would produce, the co-exceedance equivalent of a zero. Whatever link between allowances to their fundamental fuel operates at the horizons where the daily literature found it, not within 5-minute bars. Carbon-power is modestly but genuinely positive: 5 coincidences against 1.3 expected, a ratio of 3.8, with a Monte Carlo p-value of 0.004. Carbon's strongest relationship is with crude: six coincidences with Brent (1.05 expected, binomial p-value 0.0008) and four with WTI (ratio 3.9). Three or more of the five assets jump together in eleven intervals across the year, against 0.1 expected ($p = 0.002$); systemic episodes exist, but they are oil-centred, and carbon joins only occasionally. The dominant feature of carbon is its isolation: 178 of its 189 jumps, 94 per cent, coincide with nothing. Figure 3 shows the same year interval by interval, depicting the structure transparently: there are five rows of jump time intervals, with connectors on the shared intervals, the oil pair tightly connected together at the bottom, and carbon's row above, nearly untouched. Clustered arrivals do affect the picture: widening the coincidence window by one adjacent bar raises the carbon-Brent tally from 6 to 11, and the carbon-gas tally from 1 to 4, so the ranking does not stem from exact-interval alignment.

Panel B includes the regime dimension along with the early departures from H3. During the baseline, carbon's meagre coincidences are with power, four of them, roughly three per hundred days at 5.6 times chance, consistent with domestic policy and auction news hitting the German

power curve and the allowance together; carbon-crude is a single coincidence on each contract. The conflict window reverses the pattern. Carbon-Brent rises to 5.26 co-jumps per hundred days at 7.3 times chance, and carbon-WTI to 3.95, while the power link falls back to a single event, and gas manages its lone coincidence of the year. Carbon switched co-jump partners mid-sample, from the domestic legs to the global shock carrier, and the switch runs in exactly the opposite direction from the attenuation H3 predicted. Figure 4 shows both regimes as heat maps; the Brent-WTI cell setting the colour scale at ratios the carbon pairs never reach, and the eye immediately picks up the partner switch.

Panel C corroborates a statistic based on returns rather than indicators. For carbon-Brent, the ratio Φ sits at 1.003 at $k = 2$ and 0.976 at $k = 3$, almost exactly the value of one that theory assigns to genuinely common discontinuities, and the desynchronised null, centred near k at 1.96 and 3.09 as it should be, is rejected at $p = 0.09$ and 0.06. For carbon-gas, Φ lands at 1.60 against a shifted null of 1.43, nowhere near rejection: whatever the single coincidence was, the return-based statistic sees no common-jump structure. Carbon-power falls between, directionally common but insignificant. The extreme cross-product counts in the spirit of Bollerslev, Law, and Tauchen (2008), reported in the note, are informative only for the gas and power pairs, where they agree with the null result; for the crude pairs the statistic is swamped by oil's own tail, a single thirteen per cent oil return generating extremes under any alignment, and we set it aside in favour of the count and ratio evidence.

Panel D restates the crude result in a conditional-probability form, making it directly useful for traders. In the baseline, a Brent jump elevates the chance that carbon jumps in the same five-minute interval to 1.1 per cent relative to an unconditional rate of 0.6; in the conflict, the conditional probability is 5.8 per cent relative to 0.8, a multiplier of roughly 7 times with a Wilson

interval of [2.3, 14.0] per cent. The corresponding gas figures are zero in the baseline, no joint bar in 196 gas jumps, and 1.9 per cent in the conflict, with intervals that encompass the unconditional rate.

TABLE 3 Co-jump tests

Pair	N bars	Co-jumps	±1 bar	Expected	Ratio	p (binomial)	p (MC)
<i>Panel A: Co-exceedance counts, full sample</i>							
Carbon-TTF	26,220	1	4	1.59	0.6	0.796	0.794
Carbon-Power	25,080	5	9	1.30	3.8	0.011	0.004
Carbon-Brent	29,184	6	11	1.05	5.7	<0.001	0.002
Carbon-WTI	29,070	4	7	1.03	3.9	0.021	0.036
TTF-Power	22,344	1	3	1.91	0.5	0.851	0.904
TTF-Brent	26,220	4	9	1.38	2.9	0.051	0.086
Power-Brent	25,080	4	9	1.11	3.6	0.027	0.048
Brent-WTI	29,298	106	110	0.89	119.1	<0.001	0.002

Pair	Regime	Co-jumps	±1 bar	Per 100 days	Ratio to independence
<i>Panel B: Carbon pairs by regime</i>					
Carbon-TTF	Baseline	0	2	0.00	0.00
	Conflict	1	2	1.72	2.31
Carbon-Power	Baseline	4	8	2.96	5.62
	Conflict	1	1	1.39	1.64
Carbon-Brent	Baseline	1	3	0.60	1.93
	Conflict	4	7	5.26	7.28
Carbon-WTI	Baseline	1	2	0.60	1.89
	Conflict	3	5	3.95	5.80

Pair	k	Φ	Bootstrap 95% CI	Disjoint null mean	p (reject disjoint)
<i>Panel C: Jacod-Todorov common-versus-disjoint ratios</i>					
Carbon-TTF	2	1.602	[1.07, 2.17]	1.43	0.751
	3	1.646	[1.18, 2.29]	1.82	0.517
Carbon-Power	2	1.362	[1.08, 1.59]	1.52	0.348
	3	1.578	[1.24, 2.09]	2.14	0.164
Carbon-Brent	2	1.003	[0.99, 2.00]	1.96	0.095
	3	0.976	[0.95, 3.14]	3.09	0.065

Regime	Pair	Energy jumps	Joint	P(carbon energy) (%)	95% CI (%)
<i>Panel D: Conditional co- jump probabilities</i>					
Baseline	TTF to carbon	196	0	0.0	[0.0, 1.9]
Baseline	Power to carbon	126	4	3.2	[1.2, 7.9]
Baseline	Brent to carbon	89	1	1.1	[0.2, 6.1]
Baseline	WTI to carbon	90	1	1.1	[0.2, 6.0]
Conflict	TTF to carbon	53	1	1.9	[0.3, 9.9]
Conflict	Power to carbon	78	1	1.3	[0.2, 6.9]
Conflict	Brent to carbon	69	4	5.8	[2.3, 14.0]
Conflict	WTI to carbon	65	3	4.6	[1.6, 12.7]

Note: Panel A: co-exceedances of one per cent Lee-Mykland flags on common 5-minute bars, full sample; expected is N times the product of the marginal jump rates under independence; the ± 1 bar column widens the partner's indicator by one adjacent bar before counting; clustered arrivals are merged into episodes. To save space, the omitted TTF-WTI and Power-WTI rows show 3 co-jumps each (8 and 5 within one bar) at ratios of 2.2 and 2.9, respectively. The systemic count of bars with $\geq 3/5$ asset jumps is 11, against a Monte Carlo mean of 0.1 ($p = 0.002$). Comparisons at the upper extreme percentile (99.9th percentile) are informative only for the gas/power pairs, where the upper tail of gas is shifted to a mean of 12.5, and the upper tail of power is shifted to a mean of 11.4 ($p = 0.19$ and 0.37 , respectively); for the crude pairs, the tails are swamped by oil's own extremes. Panel B: rates are per 100 panel days, with 170 panel days in the baseline and 77 in the conflict. Panel C: $\Phi = V(k\Delta)/V(\Delta)$ as in equation (4) (approaching 1 under common jumps and k under disjoint jumps); intervals are day-block bootstrap, and the last column tests the disjoint null constructed from within-day circular shifts. Panel D: conditional probabilities that carbon jumps in the same bar as the named energy leg, with Wilson 95 per cent intervals; unconditional carbon jump rates are 0.6 per cent of bars in the baseline and 0.8 per cent in the conflict.




FIGURE 3 The jump map: every detection, every coincidence

Note: Every five-minute Lee-Mykland detection in the sample, one row per contract, point size proportional to the absolute return and capped at three per cent. Vertical connectors chain the assets involved in a shared bar; orange chains mark systemic bars in which at least three of the five contracts jump. The shaded band represents the conflict regime, and the dotted lines indicate the nine event dates in Table 5. The dense Brent-WTI stitching at the bottom against the near-empty connectors reaching carbon's row summarises the paper's cross-sectional result in one picture.




FIGURE 4 Co-jump dependence by regime

Note: Each cell reports the ratio of observed to independence-expected co-exceedances on the pair's common five-minute bars within the regime, with the observed count in parentheses beneath; the colour scale is capped at 20, which the Brent-WTI cells, at ratios of 149 and 80, exceed. Diagonal cells are blank. The carbon column's mass sits on power in the baseline and moves to the crude pair in the conflict.

5.3. Who excites whom: Hawkes estimates and lead-lag

Table 4 moves from coincidence to excitation to causation in the Granger sense, for one-minute jump times, where the resolution is sharpest: 484 carbon events, 480 gas, 307 power, 387 Brent, and 386 WTI. Panel A reports the fixed-kernel bivariate estimates, branching ratios at $\beta = 20$ per hour with likelihood-ratio tests of zero energy-to-carbon excitation; Panel B reports the nonparametric fifteen-minute lifts; Panel C reports the who-moved-first tallies. Before the primary numbers, the free-kernel diagnostics deserve a paragraph, because they contain a finding. Left unrestricted, the estimated decay for carbon in the baseline converges toward zero, half-lives of a day or more, with apparent self-excitation around 0.47: the kernel is measuring the fact that calm-period carbon jumps arrive in day-scale clusters, several on an auction or policy day, silence for weeks. In the conflict regime, the free decay snaps to 27-35 per hour (on the minute scale) for

619 every pair. Carbon's clustering timescale itself changed with the crisis, from days to minutes, and
620 a single exponential kernel cannot represent both at once, which is precisely why the regime
621 comparison must hold the kernel fixed.

622 On identical two-minute-half-life kernels, the loadings are straightforward. Gas excites carbon
623 nowhere: the branching ratio from TTF to carbon is 0.000 in the baseline and 0.008 in the conflict,
624 with likelihood ratios of 0 and 0.43. Power registers the identical pattern, 0.000 and 0.013, never
625 attaining statistical significance. Crude diverges considerably. One Brent jump increases the
626 expected number of carbon jumps in the subsequent minutes by 0.037 in the baseline (likelihood
627 ratio 4.86, $p = 0.028$) and by 0.041 during the conflict (likelihood ratio 8.26, $p = 0.004$). WTI
628 registers 0.030 and 0.035 with p -values of 0.022 and 0.013. The 95 per cent intervals and the delta-
629 method applied to the log loading with the kernel kept fixed show that zero is excluded in all crude
630 specification, [0.011, 0.119] and [0.016, 0.102] for Brent and [0.009, 0.099] and [0.012, 0.099] for
631 WTI, while the gas and power intervals either cluster at zero bound or remain wide open on no
632 information. The channel operates in tranquil markets, and through the most severe oil shock in
633 decades, its magnitude has not contracted by a single basis point. In Panel D, the comparison is
634 made formal: the change in the log loading across the break is 0.10 and 0.15 for Brent and WTI,
635 respectively, with z -statistics of 0.13 and 0.19 and p -values of 0.89 and 0.85. No attenuation, by
636 construction of the test built to find it. Figure 5 plots the loadings with their intervals and traces
637 the likelihood ratio is traced on the kernel grid. All reverse loadings (carbon-to-energy) are smaller
638 than against oil, ranging from 0.003 to 0.031, in line with the receiver hypothesis. While carbon's
639 self-excitation of the fixed kernel triples across the break, from 0.017 to 0.052-0.060, the
640 parametric expression of the intensity rise in Table 2 and the quantity that dilutes energy's relative
641 share of carbon's jump risk even as the absolute channel holds.

Panel B says the same thing without a likelihood. In the baseline, 10.2% of Brent jumps followed by a carbon jump within 15 minutes, which is 2.37 times the Poisson benchmark (permutation $p = 0.003$); WTI lifts 1.94. Gas increases by 0.84, which is below 1. During the conflict, the crude lifts do not change significantly as they are 1.73 and 1.63 but decline by a quarter, matching with stable branching ratios: the benchmark itself rises with the occurrence rate of carbon. Consequently, an unchanged absolute channel yields a reduced relative increase within a more active background. The two panels collectively represent the compositional decoupling result in condensed form. Figure 6 depicts the same phenomenon, independent of any temporal window specification, the probability that carbon jumps within m minutes of an energy jump for m from one to thirty: the crude curves detach from their Poisson benchmarks within five minutes in both regimes, gas hugs its benchmark, and the compression of the conflict-period lift is visible as the benchmark itself rises. Panel C establishes directionality. Within fifteen-minute windows, Brent shifts first against carbon 39 times and second 21, a ratio of 1.86; power exhibits 1.54, WTI 1.30, and gas 1.17. Every ratio exceeds one. Figure 7 finalizes the picture with one-minute cross-correlograms, whose asymmetry at small positive lags for the crude pairs align with the tallies.

TABLE 4 Mutually exciting jumps at the one-minute frequency

Pair	Regime	N(c)	N(e)	n(c,c)	n(c,e)	95% CI	n(e,c)	n(e,e)	LR	p
<i>Panel A:</i>										
<i>Fixed-kernel</i>										
<i>branching</i>										
<i>ratios, $\beta = 20$</i>										
<i>per hour</i>										
Carbon-TTF	Baseline	278	358	0.017	0.000	at bound	0.003	0.373	0.00	1.000
	Conflict	184	115	0.060	0.008	[0.000, 0.283]	0.008	0.273	0.43	0.514
Carbon-Power	Baseline	278	156	0.017	0.000	at bound	0.021	0.293	0.00	1.000
	Conflict	184	151	0.060	0.013	[0.002, 0.109]	0.000	0.257	1.30	0.255
Carbon-Brent	Baseline	278	157	0.016	0.037	[0.011, 0.119]	0.007	0.030	4.86	0.028
	Conflict	184	214	0.052	0.041	[0.016, 0.102]	0.031	0.073	8.26	0.004
Carbon-WTI	Baseline	278	168	0.017	0.030	[0.009, 0.099]	0.006	0.061	5.23	0.022
	Conflict	184	207	0.053	0.035	[0.012, 0.099]	0.055	0.102	6.12	0.013

658

Pair	Regime	N(e)	N(c)	P(hit)	Poisson base	Lift	p (perm)
<i>Panel B: Probability of a carbon jump following an energy jump within 15 minutes</i>							
TTF to carbon	Baseline	358	278	0.036	0.043	0.84	0.681
	Conflict	115	184	0.070	0.062	1.12	0.405
Power to carbon	Baseline	156	278	0.071	0.043	1.64	0.136
	Conflict	151	184	0.060	0.062	0.96	0.532
Brent to carbon	Baseline	157	278	0.102	0.043	2.37	0.003
	Conflict	214	184	0.107	0.062	1.73	0.003
WTI to carbon	Baseline	168	278	0.083	0.043	1.94	0.017
	Conflict	207	184	0.101	0.062	1.63	0.023

659

Pair	Energy first	Carbon first	Ratio
<i>Panel C: Who moves first within 15 minutes, full sample</i>			
Carbon-TTF	21	18	1.17
Carbon-Power	20	13	1.54
Carbon-Brent	39	21	1.86
Carbon-WTI	35	27	1.30

660

Pair	n(c,e) base	n(c,e) conflict	$\Delta \log n$	z	p
<i>Panel D: Regime contrast in the oil, gas, and power loadings</i>					
Carbon-TTF	0.000	0.008	at bound		
Carbon-Power	0.000	0.013	at bound		
Carbon-Brent	0.037	0.041	0.103	0.13	0.893
Carbon-WTI	0.030	0.035	0.154	0.19	0.848

661 *Note:* One minute Lee-Mykland jump timings in business time (aggregate minutes of trading in the common window). Panel
662 A: bivariate exponential Hawkes estimates where the decay parameters are held fixed at $\beta = 20$ per hour (half-life 2.1
663 minutes) for both components, $n(m, j) = \alpha(m, j)/\beta$ denotes the expected number of type-m events triggered by one type-j
664 event, with c denoting carbon and e the energy leg; the 1 d.o.f. A likelihood ratio test for the equality of α (carbon, energy)
665 = 0 is performed; cells marked at bound sites are numerically set to zero; restricted and full fits warm-start each other. Kernel
666 sensitivity at $\beta = 10$ and 40 per hour: at $\beta = 10$ all crude specifications strengthen (Brent LR 8.89 and 9.07, WTI 6.46 and
667 6.84, p between 0.003 and 0.011); at $\beta = 40$ all conflict cells hold (Brent p = 0.011, WTI p = 0.028) while baseline Brent
668 fades (LR 0.85, p = 0.357) and baseline WTI approaches marginal significance p = 0.051; gas and power never load at any
669 kernel. N(c) and N(e) count events by regime, not including the 13-day post-period. Panel B: P(hit) is the share of energy
670 jumps followed by a carbon jump within fifteen minutes; the Poisson base is $1 - \exp(-\lambda m)$; p-values are from three hundred
671 circular-shift permutations of carbon event times in the regime. Panel C tallies ordered pairs of opposite-type events within
672 15 minutes across the entire sample. Panel D tests the change in log loading across regimes using a z-statistic based on delta-
673 method standard errors; the gas and power rows are degenerate at the zero bound.




FIGURE 5 Hawkes loadings and kernel sensitivity

Note: The upper panel plots the branching ratio $n(\text{energy to carbon})$ at $\beta = 20$ per hour with delta-method 95 per cent intervals, by regime; the axis is truncated at 0.12, which clips the upper limits of the wide, uninformative gas and power conflict intervals ($[0.000, 0.283]$ and $[0.002, 0.109]$). The lower panel traces the likelihood ratio statistic for the null of no energy-to-carbon excitation across the kernel grid $\beta \in \{10, 20, 40\}$ per hour; the horizontal line marks the 5% critical value of 3.84.




FIGURE 6 Nonparametric excitation curves

Note: The probability that a carbon jump follows an energy jump within m minutes, for m from one to thirty, is computed in business time within each regime. Dashed lines are the Poisson benchmarks $1 - \exp(-\lambda m)$ at the regime's carbon event rate. The crude curves separate from their benchmarks inside five minutes; the gas curves do not.




FIGURE 7 Lead-lag cross-correlograms of one-minute returns

Note: Cross-correlations of one-minute returns of each energy leg against carbon at leads and lags of up to twenty minutes over the full sample; positive lags mean the energy return leads the carbon return. Dotted lines indicate approximate 5% bands. The large negative Brent and WTI bars, with a lag of minus one minute, are the ceasefire-window reversals visible in the tails of Table 1.

5.4. The crisis as treatment: event study and difference-in-differences

Table 5 examines the nine dated escalations against a counterfactual of conflict-regime days outside every event window. The design point bears repeating: measured against the calm baseline, every crisis week looks extreme, which restates the regime shift rather than testing the events. Against same-regime non-event days, carbon jump intensity in the windows ranges from 0.33 to 1.67 detections per day against a reference of 0.88, percentiles between the 41st and 77th, and no window survives the false-discovery adjustment. Trading volume is the one margin with event-specific texture, sitting near the 96th percentile around the onset and closure of the strait and near the 93rd when Brent passed \$100, then normalising, while quoted spreads are elevated regime-

wide. The reading is that carbon's crisis repricing was distributed across the regime rather than concentrated at the named dates, exactly what the tripled self-excitation of Table 4 implies: a market digesting a persistent shock through its own clustered dynamics rather than reacting bar by bar to headlines. A Poisson test of window co-jump counts against the regime rate agrees, returning one at every event, and Figure 8 places each window against the reference band.

TABLE 5 Event study around the Hormuz timeline

Event	Date	Co-jumps	Jumps/day	Pct (jumps)	Pct (volume)	Pct (spread)	p (FDR)
Onset of strikes	28 Feb	0	0.67	0.41	0.96	0.91	0.82
Strait closed	2 Mar	0	0.67	0.41	0.96	0.95	0.82
Brent above 100	8 Mar	0	0.33	0.41	0.93	0.91	0.82
Ultimatum	21 Mar	1	1.67	0.77	0.96	1.00	0.70
Brent peak	7 Apr	0	1.25	0.77	0.14	0.61	0.70
Ceasefire	8 Apr	0	1.00	0.77	0.13	0.61	0.70
Re-closure scare	18 Apr	0	1.33	0.77	0.68	0.61	0.70
MoU talks	14 Jun	0	1.00	0.77	0.57	0.61	0.70
MoU signed	17 Jun	1	1.20	0.77	0.57	0.61	0.70

Note: Windows of two trading days either side of each event date in 2026. Co-jumps count carbon-energy co-exceedances within the window; jumps/day is the mean carbon detections per day, compared with 0.88 on conflict-regime days outside all windows. Percentile columns locate the window mean for carbon jump intensity, carbon volume, and the carbon relative quoted spread within the distribution of the same quantities on reference days; p (FDR) applies the Benjamini-Hochberg adjustment to two-sided empirical p-values on the jump-intensity percentile across the nine events. Co-jump counts are too sparse for percentile testing and are evaluated instead by an exact Poisson test of the window count against the regime co-jump rate, which returns $p = 1$ at every event.




FIGURE 8 Event windows against the conflict-regime reference

Note: Window means over two trading days either side of each event date, plotted against the interquartile band (shaded) and median (dashed) of the same quantities on conflict-regime days outside all windows. Carbon volume is the total number of contracts accumulated daily, in millions. Volume sits far above the band around the onset, the closure of the strait, and the ultimatum, then normalises; jump intensity never leaves the reference range, matching Table 5.

Table 6 delivers the formal verdict on H3, and the cleanest statement is the exact one in Panel B. Under the placebo trend, that is, scaling by the Brent-WTI pair's own change across the break, the baseline carbon-crude co-jump rate implies 0.64 expected co-jumps during the conflict; four occurred, with an exact Poisson p-value of 0.0042. The identical framework for carbon against gas and power expects 2.56 and observes 2, with p is unity: there remained nothing to attenuate, and nothing did. The Poisson regressions in Panel A are appropriately humble given the small counts. With the treated pair defined as carbon against gas and power, the channel the hypothesis had in mind, the interaction was -0.248, but the p-value of 0.78, and the treated main effect of -2.77

records the rarity of such co-jumps in any regime. Redefining treatment as carbon against crude flips the sign and the story: the interaction of 1.832 indicates that carbon-oil co-jump intensity rose by a factor of about 6.2 relative to the placebo's own crisis-period change, with $p = 0.095$. All estimates in the paper now converge on the same direction. The naive form of H3 is rejected; conditional on jumping with energy at all, carbon coupled more tightly to the shock carrier during the shock, not less. Figure 9 shows the raw daily series behind the regressions.

TABLE 6 Difference-in-differences in daily co-jump counts, Poisson

	Treated: carbon-(gas+power)	Treated: carbon-crude
<i>Panel A: Poisson regression, cluster-robust</i>		
Intercept	-0.977 (0.125)	-0.977 (0.125)
Treated	-2.773 (0.522)	-4.159 (0.997)
Conflict	0.347 (0.195)	0.347 (0.195)
Treated $\times$ Conflict	-0.248 (0.886)	1.832 (1.099)
p-value, interaction	0.780	0.095
Rate ratio, interaction	0.78	6.24
<i>Panel B: Exact small-sample test</i>		
Observed conflict co-jumps	2	4
Expected under placebo trend	2.56	0.64
p, exact Poisson	1.000	0.0042

Note: Poisson regressions of daily pair co-jump counts on a treated indicator, a conflict indicator, and their interaction, as in equation (6), over baseline and conflict days. The placebo pair in both columns is Brent-WTI; the treated pair is carbon against gas and power combined in the first column and carbon against Brent or WTI in the second. Standard errors, in parentheses, are cluster-robust by day. The interaction coefficient measures the change in the treated pair's co-jump intensity across the regime break relative to the placebo's change; its exponential is the rate ratio. Panel B: the expected conflict-period co-jump count for the treated pair equals its baseline daily rate scaled by the placebo pair's own conflict-to-baseline rate ratio and the number of conflict days; the p-value is exact Poisson against that expectation.

Figure 10 shows the distribution bar by bar, with the eleven co-jump observations marked individually along the axis.

TABLE 7 Order flow at co-jump versus idiosyncratic carbon jumps

Jump type	N	Abnormal volume	Abnormal trades	Abnormal spread
Co-jump with energy	11	7.26	5.47	1.37
Idiosyncratic	178	4.30	3.46	1.06
Wilcoxon p-value		0.204		0.339

Note: Carbon-side order flow at the five-minute bars flagged as jumps, split by whether any energy asset jumped in the same bar. Each measure is the median across bars of the ratio of the bar's value to the median value for that time-of-day slot over the full sample: accumulated volume, number of price moves, and the relative quoted spread (ask minus bid over the

midpoint). Wilcoxon rank-sum tests compare the co-jump and idiosyncratic distributions of abnormal volume and abnormal spread. Volume and trade counts are unavailable for the TTF and power composites, so all measures are reported on the carbon side.




FIGURE 10 The microstructure signature of co-jumps

Note: Empirical distribution functions of carbon-side abnormal volume and abnormal relative spread at flagged five-minute bars, split by co-jump versus idiosyncratic status; each measure is the ratio of the bar's value to the median for its time-of-day slot, plotted on a log axis. Tick marks along the base locate the eleven co-jump bars. The volume distribution shifts right for co-jumps; the spread distributions largely overlap, matching the Wilcoxon results of Table 7.

6. Robustness, Extensions, and Threats to Identification

A referee should worry about six things in a design like this, and we take them in turn. The first is timing resolution. Re-estimated on five-minute event times with the same fixed kernel, the Brent and WTI likelihood ratios collapse to roughly 0.3 with branching ratios near 0.01, because a kernel with a two-minute half-life has nothing to work with once event times are rounded to five minutes. Far from undermining the result, the collapse locates it: the transmission operates inside

a five-minute window, which is why the co-exceedance tests, which ask about the same bar, and the Hawkes tests, which ask about the following minutes, agree. Table 8 presents the thirty-minute co-exceedance counts for completeness: entries ranging from zero to two events, directional pattern preserved for the crude pairs during the conflict, and insufficient power to detect anything, exactly what one expects when rare minute-level events are bifurcated into twenty-interval days.

TABLE 8 Detector robustness: sampling frequency and test size

Pair	Regime	Carbon jumps	Co-jumps	Per 100 days		
<i>Panel A: Thirty-minute frequency, carbon co-exceedances by regime</i>						
Carbon-TTF	Baseline	24	0	0.00		
	Conflict	17	0	0.00		
Carbon-Power	Baseline	17	2	1.45		
	Conflict	16	0	0.00		
Carbon-Brent	Baseline	24	1	0.60		
	Conflict	17	1	1.32		
Carbon-WTI	Baseline	24	0	0.00		
	Conflict	w17	1	1.32		
α (%)	Carbon	TTF	Power	Brent	WTI	Carbon-oil co-jumps
<i>Panel B: Detector size sensitivity, five-minute frequency</i>						
0.5	156	227	188	137	133	5
1.0	189	251	204	163	160	6
5.0	314	323	239	277	266	11

Note: Panel A: Lee-Mykland detections and carbon-energy co-exceedances at the 30-minute frequency, constructed identically to Table 3, Panel B; carbon jump counts vary slightly across pairs because each pair retains only its common valid bars. LSEG does not furnish a 15-minute grid. At 20 bars per day, the tests retain little power for events this rare, and the panel is reported for direction. Panel B: the complete five-minute detection and co-exceedance pipeline rerun at test sizes of one-half, one, and 5 per cent; the final column counts bars in which carbon co-jumps together with Brent or WTI.

The test size is the related issue, and this question is dealt with in Panel B of Table 8. Rerunning the detector at half, one, and five per cent moves carbon's tally from 156 to 189 and then to 314, with the crude counts scaling proportionately, and the carbon-oil co-jump count moves accordingly, 5 to 6, and to 11, monotonically aligned with the population of detected events. The

coupling is not a knife-edge property of the one per cent threshold; it scales the way genuine common discontinuities must, and threshold artefacts do not.

The second worry is that bid-quote pricing for gas and power can manufacture or destroy results. Two observations discipline it. The null results for those legs are not artefacts of missing data: on their chosen sources, gas and power generate 480 and 307 one-minute detections, comparable to carbon's 484, so the inputs to the co-jump and Hawkes stages are rich; the legs jump, they do not jump with carbon. Moreover, the favourable results involve only trade-priced series on both sides, carbon against Brent and WTI, where quote artefacts cannot enter. One-sided quotes do embed liquidity variation in the gas and power returns, which plausibly inflates their univariate jump counts. However, there is no channel through which they could suppress a genuine carbon-gas coincidence, since a real common shock would move the bid within the bar. The Hawkes results have four pairs across two regimes; the crude cells at $p = 0.004$ to 0.028 survive a factor-of-eight correction at the ten per cent level, and the conflict-period Brent cell at five per cent. The event examination incorporates its own false-discovery control and documents null results.

Fifth, the fixed kernel. Freezing β at 20 per hour is a modelling choice, and the honest defence has two parts. The free-kernel likelihood is flat along a ridge, trading slow decay for high self-mass in the calm regime, so per-regime free estimates are not comparable; we report them and note that the free decays during the conflict, 27 to 35 per hour, bracket the choice. Appendix A.3 documents the multi-start protocol that eliminates the associated optimiser artefacts. The β grid then completes the analysis. At $\beta = 10$, a four-minute half-life, all crude cell strengthens, with the baseline Brent likelihood ratio attaining 8.89 at $p = 0.003$, the most decisive statistic in the paper. At $\beta = 40$, a one-minute half-life, all four conflict cells sustain significance while the baseline

805 Brent fades to $p = 0.357$ and the baseline WTI to $p = 0.051$; the lower panel of Figure 5 traces this
806 pattern. The interpretation is substantive rather than delicate: during the calm period, the
807 transmission lasts a few minutes, a window that a one-minute kernel is too impatient to catch,
808 while during the crisis period transmission is tight enough that there is no kernel on the grid that
809 misses it. The crude loadings are stable across the grid; gas and power have null at every β ; and
810 the nonparametric lifts, which involve no kernel at all, replicate the parametric ranking for all cells:
811 6th, small counts, and the difference-in-differences. 11 carbon-energy co-jumps are a thin base for
812 an asymptotic Poisson interaction, which is why the exact construction of Table 6, Panel B, leads
813 our reading, and why the rejection of H3 rests on coherence across the exact test, the stable Hawkes
814 loadings with their formal contrasts, the regime-switched partner pattern of Table 3, and the
815 mechanism evidence of the spark spread. Episode counting points to a similar process: at a
816 tolerance of one adjacent bar, the carbon-Brent tally nearly doubles to 11 whereas carbon-gas
817 reaches only 4, so the ordering survives any plausible definition of coincidence. One design feature
818 , that is worth restating as a strength, is that returns span within-day intervals and the continuation
819 series links across sessions; contract rolls cannot contaminate any statistic in the paper, and the
820 crisis dates stem from a military timeline that no market participant selected, making reverse
821 causality from carbon into the treatment unavailable.

822 **Allowing the data to break.** Our regime division inherits its dates from a military timeline, which
823 provides the strength of design, because it is exogenous, but its vulnerability if the market reprices
824 before or after the tanks move. Table 9 lets each jump series be given its own break, for which a
825 sup-LR statistic for a single change in Poisson mean, in the spirit of Andrews (1993), is scanned
826 over the interior 70 per cent of the sample, with its CUSUM as a distribution-free companion and
827 p-values from 999 circular day-shift permutations that preserve each series' serial dependence. No

series rejects the no-break null at the 10 per cent level; the smallest p-value in the table is 0.19. That is the true headline, and it is good news for the paper: jump intensities did not lurch discontinuously at any single date, as Section 5 suggests for a channel that intensifies without being rewired. The point estimates are still useful for the information. The best-fitting break for carbon-oil co-jumps is on 20 March 2026, 20 days after the strikes, with a 90 per cent confidence set from 12 February to 10 April that comfortably brackets the onset; carbon's own jumps prefer 19 January, when the pre-war escalation began moving oil; and the OVX level, included as a yardstick, dates its break to 17-18 February with a two-day confidence set, precisely when implied volatility started pricing the impending closure. Figure 11 illustrates the sup-LR trajectories and demonstrates the carbon-oil measures rising steadily into the conflict window rather than spiking at a single announcement, the signature of a process whose frequency ramped with the news flow. The CUSUM trajectory for carbon jumps recounts the parallel narrative from a different angle: a long negative drift through the tranquil autumn, a reversal at the January escalation, with a gradual ascent through the conflict that never exceeds the 5 per cent threshold.

TABLE 9 Endogenous break tests on daily jump intensities

Series	Break (sup-LR)	sup-LR	p (perm)	90% set	Break (CUSUM)	CUSUM	p (perm)	Days from onset
Carbon jumps	19 Jan 2026	5.77	0.695	5 Jan-20 Mar 2026	19 Jan 2026	1.16	0.525	-40
Carbon-oil co-jumps	20 Mar 2026	7.71	0.331	12 Feb-10 Apr 2026	20 Mar 2026	1.35	0.190	+20
Brent-WTI co-jumps	9 Apr 2026	3.69	0.780	8 Sep 2025-1 May 2026	9 Apr 2026	0.85	0.802	+40
OVX (level)	18 Feb 2026	1516.6	0.261	17-18 Feb 2026	17 Feb 2026	6.38	0.196	-10

Note: Sup-LR tests for a single break in the Poisson mean of daily counts (levels for OVX), scanned over the interior 70 per cent of the 260-day sample; CUSUM is the standardised cumulative-sum statistic. P-values are from 999 circular day-shift permutations preserving serial dependence. The 90 per cent set inverts the likelihood ratio (chi-square with 1 degree of freedom). Days from onset measures the sup-LR break date against 28 February 2026; negative values precede the strikes. No series rejects the no-break null at ten per cent.




FIGURE 11 Endogenous break tests

Note: Upper panel: sup-LR paths for a single break in the Poisson mean of daily carbon jump counts and carbon-oil co-jump counts, scanned over the interior seventy per cent of the sample; the dotted vertical line marks 28 February 2026 and the shaded band the conflict regime. Lower panel: standardised CUSUM of carbon jump counts with five per cent bands (dashed). No statistic breaches its permutation critical value; the carbon-oil path peaks on 20 March 2026.

A continuous dose. The binary regime discards information: within the conflict, oil stress waxed and waned, and if crude is genuinely the carrier of carbon's discontinuous news, co-jump intensity should track the day-to-day intensity of oil uncertainty, not merely its regime average. Table 10 regresses daily counts on the log of the CBOE crude-oil volatility index in a Poisson QMLE with Newey-West standard errors on the scores and then, because OVX is itself violently persistent, recomputes each slope's p-value from 499 circular shifts of the regressor. The dose-response is steep exactly where the paper says it should be: a ten per cent increase in OVX is associated with 4.8 per cent more carbon jumps and 20.4 per cent more carbon-oil co-jumps, while the gas-plus-power slope is negative and the Brent-WTI placebo flat. Figure 12 plots the decile means against

the fitted curves, and the separation between the rising oil channel and the flat-to-falling gas channel is visible without statistics. The inference deserves the same candour as everywhere else in the paper. The Newey-West p-values, 0.010 for carbon jumps and 0.035 for the co-jump slope, overstate the evidence: under circular permutation, the former weakens to 0.086, and the latter dissolves to 0.52, because a slowly moving regressor can fit any persistent count series in twelve months of data. We therefore read Table 10 as a consistency check that the cross-sectional ordering survives a continuous treatment, not as independent identification; the ordering across pairs, oil positive, gas negative, placebo flat, is the finding, and it is not one a spurious-persistence story naturally generates. The OVX series used here is included in Supplementary Material S1.

TABLE 10 Oil-volatility dose-response of daily jump activity

Specification	Coef. on log OVX NW s.e.		p (NW)	p (perm)	+10% OVX
Carbon jumps	0.488	0.189	0.010	0.086	+4.8%
Carbon-oil co-jumps	1.944	0.921	0.035	0.521	+20.4%
Carbon-oil, activity controls	1.111	1.066	0.297	0.549	+11.2%
Carbon-(gas+power) co-jumps	-0.533	1.017	0.600	0.601	-5.0%
Brent-WTI co-jumps (placebo)	0.297	0.290	0.305	0.577	+2.9%

Note: Poisson QMLE of daily counts on log OVX over the baseline and conflict regimes (247 days); OVX is the CBOE crude-oil implied-volatility index, forward-filled across non-overlapping holidays. Newey-West standard errors on the scores use 4 lags. Permutation p-values recompute the slope under 499 circular shifts of the regressor, preserving both series' autocorrelation. The activity specification adds carbon and crude jump counts as controls. The final column converts the coefficient to the percentage change in intensity per 10% rise in OVX.




FIGURE 12 Oil-volatility dose-response

Note: Daily co-jump counts averaged within OVX deciles (points, log scale) against fitted Poisson curves from Table 10 for the carbon-oil pair (solid) and the carbon-(gas+power) pair (dashed). The rising oil channel and flat-to-falling gas channel reproduce the paper's cross-sectional ordering under a continuous treatment; caveats for inference are in the text.

Semicovariance: where the directional shift resides. Section 5 demonstrated the five-minute carbon-Brent correlation, revising from +0.105 to -0.097. Concurrently, the co-jump channel remained intact and characterised the pair as exhibiting demand-channel diffusive behaviour, in contrast to information-channel jumps. Realised semicovariance makes that interpretation testable at the level of moments. Following Bollerslev, Li, Patton, and Quaadvlieg (2020), we decompose each day's carbon-crude covariation into aligned flows, both returns positive or both negative, and discordant flows, opposite signs, on the same filtered five-minute grid every other exhibit uses. Table 11 reports the regime means. In the baseline, concordant flows dominate: the concordant share of gross covariation is 0.60 for both crude pairs, an ordinary positive relationship. In the conflict, the mixed components explode, from -0.35 to -1.99 (times ten thousand) for carbon-Brent,

and the aligned proportion diminishes to 0.41; the daily period shift in the discordant-minus-concordant gap carries t-statistics of -3.41 and -3.48. Total covariation becomes negative, representing the directional reversal shown in Figure 1. The decomposition, however, reveals what the correlation coefficient cannot: the crisis did not switch off the carbon-oil dependence; it redirected that dependence through opposite-signed flows, oil rising on supply news while carbon declined on the demand consequences, at the very moments when the co-jump tests demonstrate both markets responding jointly to the same discontinuities. Panel B repeats the decomposition on daily close-to-close returns and finds the same rotation with sharper statistics, so the result is not an artefact of the intraday grid.

TABLE 11 Realised semicovariance of carbon-crude pairs by regime

Pair	Regime	P+	N-	M+	M-	Conc.	Disc.	Conc. share	t (shift)
Panel A: intraday five-minute									
Carbon-Brent	Baseline	0.27	0.27	-0.17	-0.18	0.53	-0.35	0.60	-3.41
Carbon-Brent	Conflict	0.71	0.65	-1.27	-0.71	1.36	-1.99	0.41	
Carbon-WTI	Baseline	0.28	0.28	-0.18	-0.19	0.56	-0.37	0.60	-3.48
Carbon-WTI	Conflict	0.75	0.69	-1.31	-0.75	1.44	-2.06	0.41	
Panel B: daily close-to-close									
Carbon-Brent	Baseline	0.30	0.35	-0.43	-0.37	0.65	-0.79	0.45	-4.39
Carbon-Brent	Conflict	1.15	0.48	-2.72	-1.85	1.64	-4.57	0.26	
Carbon-WTI	Baseline	0.32	0.37	-0.45	-0.39	0.70	-0.84	0.45	-4.23
Carbon-WTI	Conflict	1.30	0.55	-2.45	-2.02	1.85	-4.47	0.29	

Note: Realised semicovariance components from Bollerslev, Li, Patton, and Quaadvlieg (2020), daily means by regime, scaled by 10,000. P+ accumulates products of jointly positive five-minute returns, N- jointly negative, M+ and M- the mixed-sign cross-products; concordant is P+ plus N-, discordant M+ plus M-. The concordant share divides concordant by the sum of absolute concordant and discordant flows. The t-statistic tests the baseline-to-conflict shift in the daily discordant-minus-concordant gap (unequal variances); $p < 0.01$ in all four comparisons. Panel A uses the paper's filtered five-minute grid; Panel B, daily close-to-close returns from the same continuation series.

Connectedness of the jump system. The pairwise tests ask whether two assets jump together; a complementary question is how much of each asset's jump activity is explained by shocks to the others, treated as a system. Table 12 applies the Diebold-Yilmaz (2012) decomposition, a VAR(2) on the five daily jump-count series with generalised forecast-error variance shares at a ten-day

horizon. The full-sample matrix takes the form by every other exhibit predicts: Brent and WTI form a bloc, with each accounting for about 31 per cent of the other's variance, whereas carbon is about self-contained at 96.2 per cent own-share, its largest external influence, 1.6 per cent, arriving from Brent, and gas explaining essentially nothing of carbon at 0.6 per cent. Split by regime, the system tightens exactly when the narrative says it should: total connectedness rises from 17.4 per cent in the baseline to 27.4 per cent in the conflict, and carbon's dependence on shocks from the rest of the complex more than quadruples, from 3.0 to 13.7 per cent. Figure 13 shows the total measure through a 60-day window and adds the texture the split conceals: connectedness dips in the phoney-war weeks after the strikes, when gas and power stopped jumping while crude went up, then surges to its sample maximum in mid-April as the de-escalation cascade synchronised the complex, and decays back toward baseline as the MoU approached. Because the measure runs on jump counts rather than returns, it is immune to the volatility spillovers that dominate return-based connectedness estimates and can be read as pure discontinuity contagion.

TABLE 12 Jump-count connectedness (Diebold-Yilmaz), full sample

From \ To	Carbon	TTF	Power	Brent	WTI	FROM others
Carbon	96.2	0.6	0.3	1.6	1.3	3.8
TTF gas	0.1	92.4	3.4	2.7	1.3	7.6
DE power	1.3	3.1	95.0	0.3	0.2	5.0
Brent	1.4	1.3	0.2	65.6	31.5	34.4
WTI	1.3	0.6	0.4	31.0	66.8	33.2
TO others	4.0	5.6	4.3	35.6	34.4	Total: 16.8

Note: Generalised forecast-error variance decomposition of Diebold and Yilmaz (2012) on daily jump counts, VAR(2), ten-day horizon, full sample. Entry (i, j) is the per cent of row asset i's variance attributable to shocks in column asset j; FROM-others sums the off-diagonal row, TO-others the off-diagonal column, and the corner reports total connectedness, the mean FROM-others share. By regime, total connectedness is 17.4 per cent in the baseline, and 27.4 per cent in the conflict, and carbon's FROM-others share rises from 3.0 to 13.7 per cent. Figure 13 reports the rolling sixty-day version.




FIGURE 13 Rolling jump connectedness

Note: Total Diebold-Yilmaz connectedness of the five daily jump-count series, VAR(2), ten-day horizon, sixty-day rolling window, in per cent. The shaded band marks the conflict regime and the dotted line the onset. Connectedness dips in the phoney-war weeks, surges to its sample maximum during the April de-escalation cascade, and decays toward baseline as the June MoU approaches.

What the design could and could not see. Two calculations close the section by honestly sizing the event-study inference, and they are the right place to end because they discipline every claim that came before. First, a placebo grid: the conflict window is 77 trading days, and the baseline-plus-conflict sample admits 171 windows of that length, of which 170 are placebos: Table 13, Panel A counts carbon-oil co-jumps in everyone. The true window's four events are matched or beaten by 7.1 per cent of placebos, agreeably close to the 9.6 per cent that the exact Poisson test in Table 6, based on a different construction, delivered. At the same time, the carbon-gas-plus-power count sits at the 48th percentile, indistinguishable from any window whatever. Figure 14 shows the placebo distribution with the true window marked, and the picture matches the statistics:

the oil channel's conflict concentration is unusual but not unprecedented in a sample that also contained the autumn 2025 escalation, which produced a cluster of its own. Second, exact power: under the placebo-trend null of Table 6, which expects 0.64 carbon-oil co-jumps in the window, a five per cent one-sided Poisson test needs the true rate to rise by a factor of 6.7 before it detects the increase with eighty per cent probability, and it can never detect a decrease, because the expected count is too small for any downward departure to register. Panel B tabulates the full power curve. The asymmetry is the point: the design could reject H3's attenuation story only through the coherence of many instruments, which is how Section 5 proceeded, and the affirmative finding that intensity rose survives precisely because it recurs across the exact test, the placebo grid, the dose-response ordering, the semicovariance rotation, and the connectedness surge, five constructions with five different failure modes.

TABLE 13 Placebo-window grid and exact power of the event-study design

Series	Window	True	Placebo mean	95th pct	Max	p (placebo)	Windows
Panel A: placebo grid							
Carbon-oil co-jumps	77 d	4	1.03	4	4	0.071	170
Carbon-(gas+power) co-jumps	77 d	2	1.77	3	3	0.482	170
Panel B: exact power ($\mu_0 = 0.641$)							
Rate ratio ρ	0.25	0.5	1	2	4	6.7 (MDE)	10
Power, detect increase	0.001	0.004	0.027	0.139	0.472	0.800	0.954
Power, detect decrease	0	0	0	0	0	0	0

Note: Panel A slides a 77-day window across the 247 baseline-and-conflict days and counts co-jumps in each; p is the share of the 170 placebo windows with a count at least the true window's. Panel B computes the exact power of the one-sided five per cent Poisson test of Table 6, Panel B, whose null expects $\mu_0 = 0.641$ events; the minimum detectable effect at eighty per cent power is a rate ratio of 6.7 (power at ρ of 6 is 0.738 and at 8 is 0.885). No decrease is detectable at any ratio because the null expectation is below one event. Figure 14 plots both panels.




FIGURE 14 Placebo distribution and exact power

Note: Left: carbon-oil co-jump counts across all 77-day placebo windows; the orange line marks the true conflict window's count of four, matched or beaten by 7.1 per cent of placebos. Right: exact power of the five per cent one-sided Poisson test against rate ratios from one quarter to ten (log scale); the dashed line marks eighty per cent power, reached at a ratio of 6.7. No decrease is detectable at any ratio.

Finally, there are limitations we cannot engineer away. We lack API2 coal, so the dark spread and the full switching calculus sit outside the data; the level-two book beyond best bid and offer is unavailable on these composites, so depth-based liquidity measures are not possible; and twelve months is one macro cycle with one compliance-calendar rhythm, which an extended sample will only partially remedy. The extension battery inherits one further limit worth naming: with six oil-channel events, its instruments can corroborate and bound, but not independently identify, and we have sized every claim in this section accordingly. None of these limitations touches the paper's central contrasts, which live within the sample and within the bar.

7. Conclusions and Policy Implications

This paper asked whether carbon jumps with energy, and the answer has a shape nobody quite predicted. Not with gas, ever, at any point in a year that included the largest gas supply scare in a decade: one coincidence against 1.6 expected, a Hawkes loading of zero, and a common-jump ratio that cannot tell the pair from independence. Modestly with power during a tranquil period, the domestic news influences both. And genuinely, repeatedly, within minutes, with crude oil, the asset furthest from the allowance in textbook fundamental relationships and closest to it in the flow of global discontinuous news. That channel carried a branching ratio close to 0.04 before the Strait of Hormuz closed and again after it reopened, with more statistical robustness during the crisis than previously.

The 2026 decoupling, accordingly, was real and encompassed two elements at once. It operated as diffusive behaviour: the demand channel turned carbon's smooth co-movement with oil into a negative, as collapsing spark spreads eroded the marginal allowance purchaser. Moreover, it was compositional: carbon's own jump intensity rose by a third, and its self-excitation tripled, so energy's share of carbon's discontinuity risk shrank while the pipe itself stayed open. The market did not detach from its fundamentals in 2026; it repriced one fundamental, expected fuel-switching demand, while continuing to process another, global energy news, at the speed it always had.

For market participants, the distinction is operational rather than philosophical. Hedging carbon jump risk with gas was never available at this frequency, whatever daily correlations may have suggested; hedging the oil-borne component remained available straight through the crisis; and the growing idiosyncratic component, priced or not, is hedgeable with nothing but carbon itself. Risk managers pricing EUA options should treat carbon jump risk as a mixture with regime-

1015 dependent weights: a stable, modest systematic component loading on crude, and a self-exciting
1016 idiosyncratic component that dominates under stress and clusters at timescales that migrate from
1017 days in calm markets to minutes in turbulent ones.

1018 The policy implications are sharper in the current regulatory conjuncture than they would have
1019 been even a year ago. The EU ETS is entering its most consequential redesign since Phase III: the
1020 carbon border adjustment mechanism moved to its definitive regime in January 2026, extending
1021 exposure to EUA price discontinuities to importers who must purchase certificates priced off the
1022 allowance market; free allocation for CBAM sectors is being withdrawn on a schedule that runs to
1023 2034; the extension of emissions trading to buildings and road transport (ETS2) has been deferred
1024 to 2028 amid concerns about price acceptability, with a soft price ceiling and a dedicated stability
1025 reserve built into its design; the Union has adopted a binding 2040 target of a 90 per cent net
1026 emissions reduction; and the Commission's July 2026 review proposes, among other changes,
1027 halving the Market Stability Reserve intake rate from 2028. Our results speak to this agenda in
1028 three ways. First, quantity-based stability mechanisms such as the MSR are calibrated to the level
1029 of the allowance surplus and operate at annual frequency; they are structurally silent about jump
1030 risk, which our evidence identifies as the margin along which crisis repricing actually occurs and
1031 which policy announcements themselves generate. Second, surveillance frameworks that read a
1032 declining allowance price during an energy shock as an indication of speculative distortion or
1033 impaired price discovery risk, mistaking the diffusive demand channel for a market failure: in the
1034 most severe energy shock of our sample, the information transmission from global energy news
1035 into carbon remained fully intact, the signature of a market that is functioning properly, not one
1036 that is damaged. Third, the composition finding bears on market structure. ESMA's 2026 carbon
1037 market report records that futures account for roughly three quarters of on-venue EUA trading and

1038 that financial institutions hold the majority of derivative holdings (European Securities and
1039 Markets Authority, 2026), so a concentrated set of financial participants intermediates the liquidity
1040 that absorbs jump risk; and because carbon's jump risk is predominantly idiosyncratic and self-
1041 exciting, circuit breakers or position limits keyed to co-movement with energy would bind on the
1042 wrong margin, while mechanisms that dampen self-excited cascades (transparent auction
1043 calendars, predictable policy communication, and pre-announced review timetables) address the
1044 component that actually grew.

1045 The contribution distilled represents the first high-frequency anatomy of carbon's
1046 discontinuities against the energy complex, identified through a natural experiment sufficiently
1047 separating the diffusive channel from the discontinuous one, and verified with the dual-
1048 mechanism, dual-analysis rigour that a result this policy-adjacent deserves. No constituents of
1049 jump system broke at a date the data could find; the dependence rotated, in the semicovariance
1050 decomposition, from aligned to opposing flows exactly as a demand-channel repricing overlaid on
1051 an undisrupted information channel requires; the system's discontinuity contagion rose by ten
1052 percentage points while carbon's share of it quadrupled from a modest baseline; and the placebo
1053 and power calculations bound what a 77-day window can prove, so that no single p-value is asked
1054 to carry the thesis alone. Two conclusions extend beyond the current investigation. The migration
1055 of carbon's jump clustering from a daily timescale in calm markets to a minute timescale under
1056 stress implies that single-kernel Hawkes specifications are structurally insufficient for this asset
1057 class; regime-dependent or multi-timescale kernels are the natural next step. Moreover, the
1058 observation that a 40 per cent gas repricing arrived almost without gas jumps, while oil traded
1059 headline by headline, hints at a division of labour between diffusion and jumps across commodities
1060 that the transmission literature, built on daily data, lacks the resolution to observe. Extended

samples and the upcoming compliance cycles will show how much of what we document is the signature of one extraordinary spring, and how much is the permanent microstructure of the carbon-energy nexus.

References

- Ait-Sahalia, Y., Cacho-Diaz, J., & Laeven, R. J. A. (2015). Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics*, 117(3), 585-606.
- Ait-Sahalia, Y., & Jacod, J. (2009). Testing for jumps in a discretely observed process. *Annals of Statistics*, 37(1), 184-222.
- Andersen, T. G., Bollerslev, T., & Dobrev, D. (2007). No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and i.i.d. noise: Theory and testable distributional implications. *Journal of Econometrics*, 138(1), 125-180.
- Andersen, T. G., Dobrev, D., & Schaumburg, E. (2012). Jump-robust volatility estimation using nearest neighbour truncation. *Journal of Econometrics*, 169(1), 75-93.
- Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61(4), 821-856.
- Bandi, F. M., & Reno, R. (2016). Price and volatility co-jumps. *Journal of Financial Economics*, 119(1), 107-146.
- Barndorff-Nielsen, O. E., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2(1), 1-37.
- Barndorff-Nielsen, O. E., & Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, 4(1), 1-30.

- 1082 Bibinger, M., & Winkelmann, L. (2015). Econometrics of co-jumps in high-frequency data with noise.
1083 Journal of Econometrics, 184(2), 361-378.
- 1084 Bollerslev, T., Law, T. H., & Tauchen, G. (2008). Risk, jumps, and diversification. Journal of Econometrics,
1085 144(1), 234-256.
- 1086 Bollerslev, T., Li, J., Patton, A. J., & Quaadvlieg, R. (2020). Realized semicovariances.
1087 Econometrica, 88(4), 1515-1551.
- 1088 Boswijk, H. P., Laeven, R. J. A., & Yang, X. (2018). Testing for self-excitation in jumps. Journal of
1089 Econometrics, 203(2), 256-266.
- 1090 Boudt, K., Croux, C., & Laurent, S. (2011). Robust estimation of intraweek periodicity in volatility and
1091 jump detection. Journal of Empirical Finance, 18(2), 353-367.
- 1092 Bredin, D., & Muckley, C. (2011). An emerging equilibrium in the EU emissions trading scheme. Energy
1093 Economics, 33(2), 353-362.
- 1094 Caporin, M., Kolokolov, A., & Reno, R. (2017). Systemic co-jumps. Journal of Financial Economics,
1095 126(3), 563-591.
- 1096 Chen, Y., Msofe, Z. A., Wang, C., & Chen, M. (2025). Oil price uncertainty, exchange rate volatility, and
1097 African stock markets: A nonparametric quantile-on-quantile analysis. International Review of
1098 Financial Analysis, 105, 104385.
- 1099 Chen, Y., Wang, C., & Miao, J. (2025). Marginal impact of European gas shocks on EU inflation variations:
1100 Evidence from a TVP-SVAR-SV approach. Economic Analysis and Policy, 88, 1525-1543.
- 1101 Chevallier, J. (2009). Carbon futures and macroeconomic fundamentals. Energy Economics, 31(4), 614-
1102 625.
- 1103 Chevallier, J., & Sévi, B. (2011). On the realized volatility of the ECX CO2 emissions 2008 futures
1104 contract: Distribution, dynamics and forecasting. Annals of Finance, 7(1), 1-29.

- 1105 Corsi, F., Pirino, D., & Reno, R. (2010). Threshold bipower variation and the impact of jumps on volatility
1106 forecasting. *Journal of Econometrics*, 159(2), 276-288.
- 1107 Conrad, C., Rittler, D., & Rotfuß, W. (2012). Modeling and explaining the dynamics of European
1108 Union Allowance prices at high-frequency. *Energy Economics*, 34(1), 316-326.
- 1109 Creti, A., Jouvet, P.-A., & Mignon, V. (2012). Carbon price drivers: Phase I versus Phase II equilibrium?
1110 *Energy Economics*, 34(1), 327-334.
- 1111 Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional
1112 measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57-66.
- 1113 Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical*
1114 *Association*, 74(366), 291-298.
- 1115 European Securities and Markets Authority. (2026). EU carbon markets 2026 (ESMA market report). Paris:
1116 ESMA.
- 1117 Ferrara, L., Karadimitropoulou, A., & Triantafyllou, A. (2024). Oil jump tail risk as a driver of inflation
1118 dynamics. *Journal of Commodity Markets*, 36, 100434.
- 1119 Gilder, D., Shackleton, M. B., & Taylor, S. J. (2014). Cojumps in stock prices: Empirical evidence. *Journal*
1120 *of Banking and Finance*, 40, 443-459.
- 1121 Goodell, J. W., Gurdgiev, C., Paltrinieri, A., & Piserà, S. (2024). Do price caps assist monetary authorities
1122 to control inflation? Examining the impact of the natural gas price cap on TTF spikes. *Energy*
1123 *Economics*, 131, 107359.
- 1124 Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*,
1125 58(1), 83-90.
- 1126 Hayashi, T., & Yoshida, N. (2005). On covariance estimation of non-synchronously observed diffusion
1127 processes. *Bernoulli*, 11(2), 359-379.

- 1128 Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental*
1129 *Economics and Management*, 59(1), 43-56.
- 1130 Huang, X., & Tauchen, G. (2005). The relative contribution of jumps to total price variance. *Journal of*
1131 *Financial Econometrics*, 3(4), 456–499.
- 1132 Jacod, J., & Todorov, V. (2009). Testing for common arrivals of jumps for discretely observed
1133 multidimensional processes. *Annals of Statistics*, 37(4), 1792-1838.
- 1134 Ji, Q., Zhang, D., & Geng, J. (2018). Information linkage, dynamic spillovers in prices and volatility
1135 between the carbon and energy markets. *Journal of Cleaner Production*, 198, 972–978.
- 1136 Lee, S. S., & Mykland, P. A. (2008). Jumps in financial markets: A new nonparametric test and jump
1137 dynamics. *Review of Financial Studies*, 21(6), 2535-2563.
- 1138 Liu, L. Y., Patton, A. J., & Sheppard, K. (2015). Does anything beat 5-minute RV? A comparison of realized
1139 measures across multiple asset classes. *Journal of Econometrics*, 187(1), 293-311.
- 1140 Mansanet-Bataller, M., Pardo, A., & Valor, E. (2007). CO₂ prices, energy and weather. *Energy Journal*,
1141 28(3), 73-92.
- 1142 Rittler, D. (2012). Price discovery and volatility spillovers in the European Union emissions trading
1143 scheme: A high-frequency analysis. *Journal of Banking and Finance*, 36(3), 774-785.
- 1144 Zhou, X., Bagnarosa, G., Gohin, A., Pennings, J. M. E., & Debie, P. (2023). Microstructure and high-
1145 frequency price discovery in the soybean complex. *Journal of Commodity Markets*, 30, 100314.

1150 **Supplementary Material**

1151

1152 Supplementary material associated with this article can be found in the online version.

1153 Supplementary Material S1. Data files and replication code (ZIP, 14 MB). Contains 16 Excel

1154 workbooks of LSEG Workspace intraday bar data (/data) and the master R script reproducing all

1155 reported results (/code). The data comprise front-month continuation futures for ICE EUA carbon

1156 (CFI2Zc1), ICE Endex TTF gas (TRNLTTFMc1), EEX German baseload power (TRDEBMc1),

1157 ICE Brent (LCOc1) and NYMEX WTI (CLc1) at 1-min, 5-min and 30-min frequencies over 7

1158 July 2025 to 6 July 2026, plus the daily CBOE crude oil volatility index (OVX) used in Section 6.

1159

1160 **Appendix A. Implementation Details**

1161 ***A.1. Price construction and cleaning***

1162 Each asset-frequency file from LSEG carries a timestamp and the fields listed in Section 3.1.

1163 Timestamps are parsed in UTC, and weekend rows are eliminated. Prices are placed on a complete

1164 regular grid of bar-start times covering 04:00 to 19:55 UTC for every observed date, comfortably

1165 bracketing any admissible common window. The price source is selected per asset and per

1166 frequency by unconditional moving-price coverage, the share of grid bars whose price differs from

1167 the preceding bar, with one guard: any series whose bars carry a trade price at least ninety per cent

1168 of the time is locked to last trades at every frequency, which covers carbon, Brent, and WTI. For

1169 the one-sided books the source with the greatest moving coverage wins unless trades come within

1170 fifteen per cent of it; at five minutes this selects the bid for TTF (58 per cent moving against 41

1171 for trades) and power (66 against 18), with mid-quotes infeasible on books whose ask side is

1172 populated barely half the time, while at thirty minutes TTF's trade coverage is dense enough that
1173 last trades win. The chosen series is despiked against its own daily median at a threshold of 25 log
1174 per cent, removing 4 points from the sample, then carried forward within the day for at most 60
1175 minutes, or 12 5-minute bars. The common window maintains time-of-day slots in which every
1176 asset has a valid, filled price on at least 60% of days, covering 07:05 to 16:30 UTC. Returns are
1177 within-day log differences of the filled series; the first return of each day is discarded, which also
1178 removes contract-roll splices, and an asset-day is entered when at least half of its window bars
1179 carry returns. Regimes are determined by date: baseline before 28 February 2026, conflict until 17
1180 June 2026, post thereafter.

1181 ***A.2. Detection and resampling settings***

1182 The periodicity factor of equation (1) is based on the cross-day median of absolute returns
1183 standardised by daily bipower volatility, slot by slot, floor away from zero and normalised to unit
1184 mean square. The Lee-Mykland window is $K = 20$, with no products removed and a minimum of
1185 8 valid products; the test size is 1%, with n set to the median number of bars per day, yielding
1186 thresholds of 5.16 to 5.18 across assets. The detector is rerun at sizes of 1/2, 1, and 5 per cent for
1187 Table 8, and co-exceedances are also counted at a tolerance of 1 adjacent bar for the episode
1188 columns of Table 3. Daily BNS ratio tests with tripower quarticity and the median-truncated ADS
1189 variant are run as cross-checks through the high-frequency package. All Monte Carlo inference
1190 uses day-preserving circular shifts, five hundred replications for co-exceedance and systemic
1191 counts, two hundred for the extreme cross-product and disjoint-null distributions, three hundred
1192 permutations for the fifteen-minute lifts, and two hundred day-block bootstrap draws for the Phi
1193 intervals, with a fixed seed.

1194 *A.3. Hawkes' likelihood and estimation protocol*

1195 For the bivariate exponential model of equation (5) with parameters $\theta = (\mu, \alpha, \beta)$,
1196 the log-likelihood over $[0, T]$ is

$$1197 \quad \ell(\theta) = \sum_k \log \lambda_{c(k)}(t_k) - \sum_m [\mu_m T + \sum_j \frac{\alpha_{m,j}}{\beta_m} \sum_{k: c(k)=j} (1 - e^{-\beta_m(T-t_k)})], \quad (A1)$$

1198 evaluated by the standard recursion for the decayed event sums. Event times are one-minute jump
1199 timestamps mapped to cumulative trading minutes within each regime, with ties broken by a
1200 fraction-of-a-second offset. Parameters are estimated in logarithms using quasi-Newton
1201 optimisation from three dispersed starting points; in the fixed-kernel specification, β is held at the
1202 grid value, and only baselines and loadings are estimated. For each likelihood ratio test, the
1203 restricted model, $\alpha(\text{carbon}, \text{energy}) = 0$, is estimated first; its solution seeds the unrestricted model
1204 with the freed loading initialised near zero, and the unrestricted solution symmetrically re-seeds
1205 the restricted fit; the better of each pair is retained. Delta-method standard errors for the loadings
1206 are read from the numerical Hessian of the fixed-kernel fit on the log scale. The protocol
1207 guarantees a non-negative likelihood ratio and eliminates the two convergence artefacts, one
1208 spuriously significant and one spuriously negative, that single-start estimation produced in
1209 development. Free-kernel estimates reported in the text use the same machinery with β
1210 unrestricted. The five-variate system is estimable in principle with the same recursion; with the
1211 event counts of this sample, we judged the pairwise models to be an honest specification and left
1212 the system fit as an option in the replication code provided in Supplementary Material S1.