The Environmental Costs of Sanctions: Flaring, Venting, and Biodiversity in Venezuela*

Michele Fioretti† Kavanaugh FitzPatrick‡ Alessandro Iaria§

May 15, 2026

Abstract

Economic sanctions on oil-producing autocracies reduce not only how much oil is produced but also how each remaining barrel is produced. Combining field-level production data, satellite measurements of gas flaring and atmospheric methane, and citizen-science records of bird abundance as a proxy for local biodiversity across six Latin American countries from 2012 to 2024, we study the US sanctions on Venezuela. Per-barrel gas flaring at heavy Venezuelan fields roughly tripled relative to unsanctioned neighbours even as production fell by two-thirds, so post-period environmental damage is about three times what proportional decline would predict. Three reinforcing channels generate this gap: small clean fields exit first; surviving fields cut maintenance budgets; and surviving infrastructure runs past design capacity. Bird abundance falls in the affected cells, consistent with pollution-driven ecological harm. A November 2022 US Treasury licence that let Chevron alone resume operations at four Venezuelan fields provides a unique quasi-experimental laboratory for the maintenance channel: per-barrel pollution falls sharply at those fields, and bird abundance recovers in tandem. The same pattern of rising per-barrel emissions despite falling output appears in sanctions on Russia (post-2014/2022) and Iran (post-2018).

JEL Classification: F51, Q53, Q35, L71

Keywords: sanctions, gas flaring, methane, Venezuela, oil production, environmental externalities, Orinoco Belt, biodiversity

*Fioretti and Iaria gratefully acknowledge funding from the British Academy/Leverhulme Small Research Grants Scheme (grant number SRG2324/240698). The study was also funded by the European Research Council (ERC Starting Grant 101221025 – BALANCE). All errors are ours.

†Università Bocconi, Department of Economics, IGIER, and CEPR. E-mail: fioretti.m@unibocconi.it

‡Università Bocconi. E-mail: kavanaugh.fitzpatrick@studbocconi.it.

§University of Bristol, School of Economics, and CEPR. E-mail: alessandro.iaria@bristol.ac.uk

"Telegram from Walter Saxer in Iquitos. Apparently things are looking very good, except that the whole situation might collapse from one moment to the next."
Werner Herzog, *Conquest of the Useless*, 2004

1 Introduction

Economic sanctions are among the most widely deployed instruments of foreign policy, yet their environmental consequences remain poorly understood. When sanctions target oil-producing states, the standard policy calculus weighs geopolitical objectives against humanitarian costs. What it typically omits is that oil extraction is infrastructure-intensive: wells, pipelines, gas-handling plants, and processing facilities all require sustained investment, imported equipment, and technical expertise to operate within environmental norms. When sanctions sever the channels that sustain this infrastructure, the result may not be a clean reduction in output but a degradation of the production process itself, making each remaining barrel dirtier rather than just scarcer.

Whether sanctions raise or lower the environmental damage from oil extraction is theoretically ambiguous. Sanctions reduce the volume of oil produced, which by itself would proportionately reduce the total damage. But they may also raise the pollution generated per barrel, and if pollution per barrel rises enough, total damage need not fall and may even rise. Three forces can drive pollution per barrel up. First, the carbon intensity of crude oil production varies enormously across fields, and how production is reallocated when fields exit matters (Masnadi et al., 2018; Coulomb et al., 2026; Asker et al., 2024): when sanctions raise the cost of operating, marginal small fields shut down first and surviving production concentrates on the larger fields, which may be cleaner or dirtier per barrel depending on the cross-sectional intensity gradient: this is the *composition* channel. Second, sanctions cut the maintenance budgets at the fields that keep producing; deteriorating equipment captures less of the gas that comes up with crude oil, so more of it leaks or has to be burned off, the *maintenance* channel. Third, surviving fields are pushed harder against the limits of their gas-handling infrastructure, and more gas escapes through leaks or flaring as those limits are approached, the *strain* channel.

Each channel responds to a different sanctions-design choice. The composition channel responds to how broadly sanctions are applied, uniformly across all fields or targeted at the largest and dirtiest ones, as the G7 oil price cap on Russia attempts to do. The maintenance channel responds to whether exemptions allow some operators to keep importing skilled technicians and replacement parts, since well-maintained equipment leaks less gas: the November 2022 US Treasury licence for Chevron in Venezuela is one such exemption. The strain channel responds to whether sanctions cut off physical inputs that production requires; in Venezuela, for instance, the thickest crude needs to be diluted with imported lighter petroleum products before it can flow through pipelines, so when sanctions cut off these imports, production has to squeeze through ever-smaller pipeline capacity. What is new here is that we introduce this three-channel framework, measure each channel separately, recover an underlying damage function that links operating decisions to per-barrel

pollution, and trace its consequences through to local biodiversity.

The Venezuelan setting is what makes this identification possible. The country holds the world's largest proven oil reserves, concentrated in extra-heavy crude deposits in the Orinoco Belt that require continuous investment, imported inputs, and technical expertise to keep flowing (U.S. Energy Information Administration, 2024; Rodríguez, 2022b). The sanctions we study are public: explicit US Treasury restrictions, distinct from the private sanctions through which firms voluntarily disengage from a jurisdiction (Hart et al., 2024; Fioretti et al., 2026b). The international joint-venture partners covered by these sanctions did not expressly choose to withdraw from Venezuela and could consider re-entry, but the 2007 nationalisation of upstream JV assets creates a sovereign hold-up precedent that conditions any future investment decision (Fioretti et al., 2026a). US sanctions came in two waves with different identifying content. The first, between 2017 and 2019, severed financing (Congressional Research Service, 2024) and access to the imported naphtha diluent that heavy crude requires to flow through pipelines, generating extensive year-to-year variation in maintenance budgets and production at each Venezuelan field. The second, in November 2022, was a US Treasury licence that allowed Chevron alone to resume operations at four Venezuelan joint-ventures while sanctions remained in place at every neighbouring field. This licence is the key identifying shock: it raised maintenance spending and restored access to imported parts and foreign expertise specifically at the four covered fields, so its differential effect on pollution there separates the maintenance mechanism from any sanctions-driven co-movement of spending and pollution. Russia and Iran exhibit the qualitative three-channel pattern under their own sanctions episodes but lack a within-firm licence shock and therefore cannot recover the underlying damage function.

We link field-level production and expenditure data to satellite measures of gas flaring, atmospheric methane, and avian biodiversity. The cross-country dimension is asymmetric across outcomes: flaring and operating expenditure use the full six-country LATAM panel; methane uses Venezuelan TROPOMI cells (with a six-country robustness check in the appendix); and the biodiversity regressions use Venezuelan eBird checklists, with Colombian checklists added as a placebo control for the bird-relocation test. Venezuelan production fell to roughly one-third of its pre-sanctions level over the sanctions period, but the decline was sharply uneven: production concentrated on the largest fields in the Orinoco extra-heavy belt, whose resources are more polluting, while smaller fields and the conventional Maracaibo Basin contracted sharply across the board. Because PDVSA owns or operates most Venezuelan fields, the reshuffling happens largely within a single multi-asset firm choosing across its heterogeneous portfolio rather than between firms exiting and entering (Fioretti et al., 2025). This is the *composition channel*, the across-field margin: sanctions reshaped which fields keep producing.

The other two channels operate within fields. The *maintenance* channel: holding production fixed, the average within-Venezuela elasticity of flaring volume to per-barrel operating expenditure is small, on the order of a 3% reduction when OPEX per barrel doubles, and the post-2017 within-field slope flattens close to zero, consistent with sanctions-driven binding constraints overwhelming

the pre-sanctions maintenance margin. The November 2022 US Treasury licence sharpens this within-field test: at the four Chevron joint-ventures during the licence period, doubling maintenance spending cuts flaring per barrel by roughly 25%, an order of magnitude steeper than the average Venezuelan field, because each restored dollar buys real gas-handling parts that secondary sanctions block elsewhere.

The *strain* channel operates through a different lever. When small Venezuelan fields shut down, the surviving big heavy-oil fields had to make up the lost production. They now waste about three times as much gas per barrel as comparable heavy-oil fields in countries without sanctions. Venezuela's lighter-oil fields show no such jump. The asymmetry tracks the underlying physics: only operations whose gas handling depends on imported diluent — itself in chronically short supply since the 2019 embargo — respond to the supply-chain shock, and only in fields where post-sanctions production was sustained near pre-sanctions throughput.

The three-channel decomposition, applied to the observed surge in maintenance spending and production at Chevron fields, predicts pollution per barrel roughly halving at the four licence fields. Flaring damages local populations both directly — through respiratory and other health effects from chronic exposure to combustion pollutants (Blundell and Kokoza, 2022) — and indirectly through degradation of surrounding ecosystems. We test this prediction on local avian biodiversity. Birds are widely used as indicators of broader ecosystem health: they sit high in the food web and integrate effects on insects, plants, and small vertebrates that they consume; they are sensitive to temperature, air quality, and habitat disturbance; and their mobility means local declines reflect local conditions rather than slow demographic inertia (Şekercioğlu, 2006; Gregory and Van Strien, 2010; Eglington et al., 2012).

Bird abundance near the four Chevron fields recovers post-licence by approximately 53%, against the framework's predicted 51% decline in per-barrel pollution. The implied elasticity of bird abundance to local pollution is substantial and in the order of magnitude estimated by Liang et al. (2020) for US bird populations under the NO_x Budget Trading Program, though somewhat larger — consistent with higher baseline pollution at sanctioned Venezuelan oil sites, the simultaneous decline of multiple flaring co-pollutants, and the habitat sensitivity of tropical forest insectivores.

Combining the three channel-specific estimates, total environmental damage from Venezuelan oil extraction in 2022 to 2024 is approximately three times what the proportional decline of output alone would predict. The maintenance and strain channels together contribute essentially all of the gap (the composition channel does less than one percent of the work numerically), but composition determines where the maintenance–strain damage lands. The cross-country comparison with operating heavy-oil regions in Mexico, Ecuador, and Colombia, which face no diluent shortage and exhibit no comparable methane trend at pipeline-vicinity grid cells, confirms that the within-Venezuela pattern is sanctions-specific rather than generic to heavy-oil extraction.

The same qualitative pattern appears in Russia and Iran. We do not run our own field-level decomposition for either case; the overlapping sanctions waves on Russia and the opacity of post-2018 Iranian production data both go beyond the scope of this paper. External evidence

on production composition and methane emissions, however, is consistent with the qualitative pattern we document for Venezuela: pooling Russia's 2014, 2017–2018, and February 2022 sanctions episodes, satellite and industry sources point to flaring intensity at sanctioned Russian fields rising sharply against unsanctioned analogues. In Iran, post-2018 maximum-pressure sanctions roughly halved production at the 2020 trough, with partial recovery as dark-shipping networks routed crude to non-cooperating buyers (Fernández-Villaverde et al., 2025); satellite reanalysis by the International Energy Agency shows Iranian oil-and-gas methane emissions rising despite falling output. In all three cases, total output falls or stagnates, production concentrates on the largest dirtiest fields, and per-barrel emissions rise. What does not travel is the structural identification: only Venezuela has the within-firm licence carve-out that decomposes the regime differential into maintenance and strain.

These findings have direct implications for the design of sanctions targeting petrostates. The decomposition implies that sanctions design has two distinct objectives, not one: maximising coercive impact on the regime, and minimising the environmental damage that production reshuffling induces. The three channels operate through distinct policy levers, so these objectives are not always in conflict. A package that targets revenue from the dirtiest large fields while licensing maintenance and parts at those same fields hits the regime's revenue base hard while leaving the maintenance and strain channels quiet. The Chevron carve-out demonstrates one side of this directly, and the G7 oil price cap on Russia is a partial implementation of the other; our framework suggests how to extend both. The environmental costs we document are absent from the standard sanctions calculus, and the structural decomposition we recover names the policy levers that can blunt the environmental cost without compromising the coercive one.

Related literature. The paper sits at the intersection of work on the economic consequences of sanctions, the environmental externalities of oil and gas extraction, and the misallocation literature. The Venezuelan case has attracted considerable attention. Rodríguez (2022a), exploiting credit-access heterogeneity across Orinoco joint ventures, finds large output losses attributable to financial and oil-sector sanctions; Bahar et al. (2019) argue that aggregate counterfactuals for this idiosyncratic case are infeasible. The broader sanctions literature documents trade reductions (Afesorgbor, 2019), GDP losses through disrupted capital flows (Gutmann et al., 2023), and evasion through opaque shipping networks (Fernández-Villaverde et al., 2025); in an oil sector dependent on imported services and diluent, these channels supply the macro background to the field-level effects we document.[1] We move from aggregate output to field-level environmental performance and decompose the response into three separately identified channels rather than a single volume-or-intensity effect.

A growing body of work, building on Madani (2020), documents environmental degradation in sanctioned countries and attributes it to constrained maintenance and increased reliance on

[1] On the humanitarian costs of Venezuelan sanctions, see Weisbrot and Sachs (2019) and the response in Hausmann and Muci (2019). The paper also speaks to the political economy of sanctions (Egorov et al., 2025; Itskhoki and Ribakova, 2024).

inefficient production (Fu et al., 2020; Sun and Wang, 2025; Khalid et al., 2024; Afesorgbor and Steinbach, 2024; Madani, 2021). None of these papers separates composition, maintenance, and strain; none identifies the maintenance channel quasi-experimentally. We do both, and we tie the result to ecological outcomes in the same panel. Our biodiversity analysis sits in a complementary literature that uses avian data to recover environmental costs of human activity (Liang et al., 2020, 2025; Katovich, 2023; Hogan et al., 2025) and that treats bird populations as integrating indicators of ecosystem health (Eglington et al., 2012; Gregory and Van Strien, 2010). Field-level production shocks have, to our knowledge, not been linked to satellite pollutants and ecological outcomes within a single decomposition.

The infrastructure-and-composition view of oil-sector externalities has clear antecedents. Coulomb et al. (2026) show that crude carbon intensity varies sharply across fields, that historical extraction patterns embody large-scale carbon misallocation, and that small policy shifts can reshape which barrels are produced. Asker et al. (2024) decompose OPEC's environmental effects into a volume effect (fewer barrels) and a composition effect (dirtier barrels at high-cost producers). Covert and Kellogg (2025) show that pipeline and processing constraints push operators toward dirtier extraction methods. We extend this perspective by identifying sanctions as another institutional shock that reshapes extraction composition, and by separating composition from maintenance (within-field operating-budget collapse) and strain (within-field capacity-constraint binding).[2]

The misallocation literature provides the conceptual frame (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Bartelsman et al., 2013; Asker et al., 2014), including a recent strand that treats the depreciation rate of capital as itself endogenous to the supply of repair, parts, and expertise (Graff, 2026). An environmental strand has applied the lens to pollution: Shapiro and Walker (2018) decompose the US manufacturing pollution decline into composition, technique, and scale and find that composition does relatively little while technique change drives most of the result; Tombe and Winter (2015) show how intensity-based environmental standards themselves distort allocation across heterogeneous firms; Fowlie et al. (2012) study how emissions markets reallocate abatement across plants. We add a sanctions-driven misallocation along the environmental margin: the production base reallocates toward fields with sharply higher per-barrel pollution, and the welfare loss runs through external damage rather than through measured TFP. The Chevron quasi-experiment recovers the marginal product of the maintenance input directly, supplying an elasticity that misallocation models typically calibrate from aggregate moments, and the three-channel framework separates what the literature usually treats as a single composition wedge into a survival margin, an input-spending margin, and a capacity-constraint margin, each mapped to a distinct policy lever.

The remainder of the paper proceeds as follows. Section 2 states the conceptual framework and

[2]A related supply-side literature studies deliberate restrictions on extraction (Harstad, 2012; Ahlvik et al., 2026; Kellogg, 2024); sanctions differ in that they restrict extraction involuntarily and without environmental intent, worsening outcomes by degrading the infrastructure that contains emissions. The flaring and methane mechanisms themselves have been studied extensively (Agerton et al., 2023; Lade and Rudik, 2020; Marks, 2022; Nathan et al., 2024; Hancock et al., 2025), and a parallel literature in finance examines how credit constraints affect pollution control (Xu and Kim, 2022; Lanteri and Rampini, 2025; Singer, 2026). Both feed directly into the maintenance channel we identify.

its three predictions. Section 3 provides the institutional and engineering context: the chronology of sanctions, the physics of extra-heavy oil, and the gas-handling chain that links naphtha supply to flaring outcomes. Section 4 describes the data sources and identification strategy. Sections 5, 7, and 6 present the empirical evidence on each channel in turn. Section 8 quantifies the combined damage ratio R, extends the analysis to Russia and Iran, and draws implications for sanctions design. Section 9 concludes. Appendix A formalises the two-field reallocation model underlying the framework; the remaining appendices document data sources, matching procedures, and robustness specifications.

2 Conceptual framework

A natural benchmark for the environmental impact of a sanctions-induced output decline is *proportional reduction*: if production falls by x%, then the damage from extraction (flaring, methane leakage, local pollutants, and the biodiversity outcomes those entail) falls by x% too. This benchmark presumes that the production technology is unchanged, the same set of fields are operated, each is run with unchanged operating intensity, and none are pushed beyond the production rates they were designed to handle. Whether the proportional benchmark holds is the central empirical question of this paper.

The production technology in oil extraction has three margins along which sanctions can move outcomes away from the proportional benchmark.

Composition. Sanctions raise the per-period fixed cost of operating any field (sanctions-evasion infrastructure, the loss of foreign expertise, premium pricing for sanctioned counterparties, deferred replacement of expatriate technical staff). Small fields with low marginal extraction costs are typically the same fields with low cumulative production capacity; once fixed costs rise, they become unprofitable first. Surviving production therefore concentrates on the larger fields, and the environmental consequence depends on whether those surviving large fields are cleaner or dirtier per barrel than the small fields that exit. In Venezuela the surviving large fields are extra-heavy and require imported diluent to flow, but they are also at the flat part of the gas-to-oil-ratio decline curve and operate large fixed gas-handling infrastructure that amortises over many barrels: empirically, large Venezuelan fields are mildly *cleaner* per barrel than the small fields that exit. The composition shift therefore does little to aggregate intensity directly; the environmental damage instead loads predominantly on the within-field maintenance and strain channels described next.

Maintenance. Per-barrel emission intensity is not a fixed engineering constant; it depends on operating expenditure broadly construed: how often leaks are repaired, how quickly worn equipment is replaced, whether associated gas is captured and re-used or simply burned, alongside the labour, contracted services, and lifting inputs that move with maintenance activity. Sanctions cut these budgets across the board (through limited access to spare parts, capital, and foreign service

contractors), raising per-barrel emission intensity *everywhere production continues*, even at fields whose total output is unchanged. The empirical work in Section 7 proxies the maintenance budget with total per-barrel OPEX, which captures maintenance spending plus the correlated operating inputs that move with it in the data.

Strain. Each oil field is built for a particular planned output level. The pipes, compressors, and gas-handling equipment that move the co-produced gas safely to market or back into the reservoir have a fixed capacity, and they have a safety margin built in for surges and breakdowns. As small fields exit and their production share is absorbed by surviving large fields, those large fields are pushed closer to their design capacity, the safety margin shrinks, and a greater share of the associated gas ends up flared or vented rather than captured. In the Venezuelan setting, the binding constraint is triggered through an imported-input bottleneck: the diluent shortage caused by the 2019 naphtha embargo means heavy crude piles up at wellheads, sales pipelines stall, and the gas-handling chain is forced past its operating margin even when nominal production has fallen. The damage per barrel rises faster than the production on each surviving field rises, formally, the damage function is convex in field-level output.[3]

Theoretical framework. These three margins are exhaustive: they map one-to-one onto the three things one needs to know to compute total environmental damage, which fields operate, how each is run, and how much each produces. Appendix A formalises these channels in a stylised two-field model and derives the threshold at which surviving production concentrates entirely on the dirty field; the model delivers a damage function of the form

$$D \;=\; \underbrace{\sum_{i \in \text{operating}}}_{\substack{\text{(i) } \textit{composition}\text{:}\\ \text{which fields operate}}} \Bigg[\underbrace{\alpha_i(m_i)}_{\substack{\text{(ii) } \textit{maintenance}\text{:}\\ \text{per-barrel intensity}\\ \text{at field } i}} \cdot q_i \;+\; \underbrace{\tfrac{\beta_i}{2}\, q_i^2}_{\substack{\text{(iii) } \textit{strain}\text{:}\\ \text{convex damage as } q_i\\ \text{approaches capacity}}} \Bigg]. \tag{1}$$

The three channels are now visible as three distinct mathematical objects: the *set* of operating fields (composition), the per-barrel intensity $\alpha_i(m_i)$ as a falling function of the maintenance budget (maintenance), and the convex function $\frac{\beta_i}{2} q_i^2$ that captures the rise in damage as field i's output q_i approaches its design capacity (strain), with $\beta_i \geq 0$. The proportional-decline benchmark obtains when sanctions shift only the overall scale of output without changing any of these three: the set of operating fields, the maintenance budget at any field, or the proximity of any field to its capacity. Any departure from the benchmark must operate through one of these three arguments.

[3]The maintenance and strain channels are conceptually distinct: maintenance is about *how* a given barrel is produced (with what budget for leak prevention and gas capture), strain is about *how much* a given field is asked to produce (relative to its design capacity). They are nonetheless quantitatively difficult to disentangle in outcome measures (methane concentrations, bird abundance) that respond to both. Identifying them separately requires research designs that hold one fixed while varying the other; we exploit such designs in Section 7 (the Chevron license shock changes the maintenance budget at approximately fixed production) and Section 6 (the production-share shock raises strain at approximately fixed maintenance).

In what follows, we first describe the background, providing information on both the institutions and the engineering, as well as the data. Then, Sections 5, 7, and 6 identify the composition, maintenance, and strain channels respectively, each through a research design that holds the other arguments of eq. (1) fixed. Section 8 combines the three estimates to quantify how much the actual environmental damage exceeds the proportional-decline benchmark, and Section 8.4 discusses implications for sanctions design.

3 Background: sanctions and the physics of extra-heavy oil

This section provides two pieces of context the empirical analysis will draw on. Subsection 3.1 summarises the sanctions chronology in Venezuela and the analogous episodes in Russia and Iran. Subsection 3.2 describes what the proportional-decline counterfactual looks like in a well-functioning oil field, so the deviations we measure later have a baseline to be measured against, and walks through the physical steps by which a cut in imported diluent supply translates into flaring and methane releases at Venezuelan extra-heavy oil fields. This second subsection is essential because the sanctions–environment link runs through a specific physical mechanism that is not generic to oil extraction; understanding it clarifies why our channel decomposition has the shape it does and why Venezuela is, in this respect, a particularly informative case.

3.1 Sanctions chronology

Venezuela holds the world's largest proven oil reserves and produces from two main regions with distinct crude characteristics. The *Maracaibo Basin*, in the country's west, produces conventional medium-density crude that flows through pipelines without modification. The *Orinoco Belt*, in the country's east, produces extra-heavy crude so viscous at surface temperatures that it cannot move through pipelines without first being blended with imported light hydrocarbons (most commonly naphtha). Smaller volumes come from Barinas in the south-west, eastern conventional fields, and a thin offshore base. Before the 2017–2019 sanctions, Orinoco accounted for roughly 51 % of national crude output and Maracaibo for roughly 30 % (U.S. Energy Information Administration, 2024; Monaldi et al., 2021). The asymmetry in crude characteristics, and in particular the Orinoco Belt's dependence on imported diluent, is the physical reason for the sanctions-impact patterns this paper documents.

The United States imposed sectoral financial sanctions on Venezuela in August 2017 (Executive Order 13808), restricting US-person dealings in new debt and equity issued by the Government of Venezuela and PDVSA. In January 2019 the Treasury's Office of Foreign Assets Control (OFAC) designated PDVSA itself as a Specially Designated National, severing PDVSA from the US dollar clearing system and from US-sourced inputs, most consequentially, naphtha and other light hydrocarbon diluents that Orinoco extra-heavy crude requires for transport. A series of secondary sanctions in 2020 expanded the perimeter to non-US counterparties dealing in Venezuelan crude. Aggregate Venezuelan production fell from a 2016 peak of 2.28 million barrels per day (Rystad

UCube) to a 2020 trough of 0.58 before stabilising near 0.78 in 2023–24. Within-country, the decline was heterogeneous: Maracaibo Basin output (medium-API conventional crude, 11–19° API) collapsed across all field-size classes to 15–30% of 2016 levels, while Orinoco extra-heavy production (8–10° API) was retained much more at the largest fields (46% of 2016 output) and collapsed at the small and medium fields (7–8%). In November 2022 OFAC issued General License 41, permitting Chevron specifically to resume operations at four Venezuelan joint-ventures (PetroBoscan, PetroIndependencia, PetroIndependiente, and PetroPiar) under restrictions on revenue distribution to PDVSA. This license is the source of the quasi-experimental variation we use in Section 7.[4]

3.2 The proportional-decline benchmark and why Venezuela violates it

What should environmental outcomes do when production falls in a well-managed oil field? Two facts pin down the benchmark. First, flaring is engineered to be small or zero in steady state: associated gas is captured for sales, reinjected, or used as on-site fuel, and flared only as a safety relief (World Bank, 2024, 2025). In our LATAM panel, pre-sanctions flaring intensity at non-Venezuelan fields ranges from 14 to 52 m^3/bbl across countries against a Venezuelan pre-sanctions mean of $\sim$196 m^3/bbl (Table 1, computed from Rystad UCube and World Bank VIIRS, 15 km flare-allocation radius); the World Bank Global Gas Flaring Reduction Partnership's global inventory shows best-in-class producers operate well below 5 m^3/bbl (Agerton et al., 2023; International Energy Agency, 2024). Second, methane leakage scales roughly proportionally with the amount of equipment in service and its age and maintenance status (Marks, 2022; Plant et al., 2022).[5] Under proportional decline, flaring intensity stays near its pre-period level and methane falls in line with activity.

This benchmark fails in Venezuela because of how extra-heavy crude is produced. Orinoco crude is too viscous to flow through pipelines without 20–40 % blending with imported naphtha or other light hydrocarbons (Monaldi et al., 2021; U.S. Energy Information Administration, 2024); the January 2019 PDVSA designation severed the US naphtha supply, and each lifted barrel now has to wait for thinning at the gathering centre. Compounding this, heavy-oil wells using cyclic steam stimulation or electric submersible pumping cannot be cleanly throttled without permanent damage from wax and asphaltene precipitation and pump seizure. The associated gas, produced at near-fixed gas-oil ratios with the crude, has nowhere to go when diluent supply collapses: sales pipelines require diluted crude flowing through them, reinjection compressors require parts, and upgraders require crude feed. The only remaining gas sinks are the flare and the vent. The asymmetry is operationally clean: Maracaibo conventional grades flow without diluent and can be shut down cleanly; Orinoco extra-heavy fields cannot, and continue producing under sanctions only with deteriorated operations. The full five-step physical chain is in Appendix B; the three

[4] Russia and Iran have undergone broadly similar sanctions episodes, with qualitatively similar production-composition and methane patterns, which we discuss as external-validity checks in Section 8.3.

[5] The clearest empirical test of this proportionality is the COVID-19 oil-price shock: when US oil-and-gas activity contracted in spring 2020, satellite-observed methane emissions fell by a roughly proportional amount, with the shock fully reversing as activity recovered (Lyon et al., 2021).

channels of Section 2 (composition, maintenance, strain) are all consequences of this physical structure.

4 Data and identification strategy

The paper draws on five datasets. Field-level production, expenditure, and ownership come from Rystad Energy's UCube; satellite flaring is from the World Bank Global Gas Flaring Reduction Partnership's VIIRS products; atmospheric methane is from the Sentinel-5P TROPOMI L3 OFFL bias-corrected XCH_4 retrievals; biodiversity is from eBird citizen-science checklists; and pipeline networks are from a unioned OpenStreetMap and Global Energy Monitor compilation. We describe each in turn, then summarise the spatial-matching procedure that links them and the identification strategy that exploits their joint variation.

4.1 Field-level production: Rystad UCube

The Rystad UCube provides field-level monthly observations on crude production, capital expenditure, operating expenditure, ownership shares, crude grade (API gravity, sulfur content), reservoir type, and operational status, for a global panel of upstream oil-and-gas assets.[6] We restrict to the six-country Latin American panel (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, which yields approximately 1,874 fields and 11,983 field-year observations. Within Venezuela, 236 fields are tracked. Reporting density is essentially complete on tracked fields throughout the post-sanctions period: among Venezuelan fields with non-zero 2016 production, the share with non-zero post-2017 reporting in any year is 97% to 100%, indistinguishable from Argentina, Brazil, Colombia, Ecuador, or Mexico over the same window. The number of tracked Venezuelan fields itself drifts modestly downward (from 223 in 2016 to 160 in 2024), reflecting decommissioning rather than reporting attrition. We treat the dataset as the closest available approximation to a complete frame of upstream Venezuelan operations and discuss alternative coverage assumptions in Appendix C.4.

For the composition channel we aggregate to the field-year level and exploit two cross-sectional dimensions: (i) regional groups within Venezuela (Orinoco East, Maracaibo, Barinas, Other) corre-

[6]The UCube is built bottom-up from public regulatory filings (SEC 10-K/20-F, ANP-Brazil, ANH-Colombia, CNH-Mexico, Argentina's Secretaría de Energía monthly tablas dinámicas), national-oil-company operational and financial reports, operator and joint-venture-partner disclosures (annual reports, investor presentations, technical conferences), OPEC's Monthly Oil Market Report (which publishes secondary-source production estimates for member states), industry trade press, and Rystad's own analyst-network direct outreach to operators; field-level production, CAPEX, and OPEX are reconciled with country-level totals, and geological and reservoir attributes are sourced from regulatory filings and operator technical disclosures. For sanctioned or otherwise opaque jurisdictions (Venezuela, Iran, Russia), Rystad triangulates across multiple independent sources, most importantly OPEC secondary-source production figures, joint-venture-partner disclosures (Chevron, Repsol, Eni publish Venezuela exposure in their consolidated reports), tanker-tracking volumes (Kpler, Vortexa), refinery throughput data, and satellite-derived infrastructure indicators. The natural concern that PDVSA's curtailed public disclosures after 2017–2019 might erode field-year coverage is addressed directly: the number of tracked Venezuelan fields drifts gradually from 223 in 2016 to 160 in 2024 with no abrupt break at the sanctions boundaries, and the share of tracked Venezuelan fields with strictly positive reported OPEX remains 97–100% post-2017, indistinguishable from Argentina, Brazil, Colombia, Ecuador, and Mexico over the same window.

sponding to crude type and operational technology; (ii) binary large/small indicators based on the country pre-period median of OOIP, of mean 2012–2016 production, and of the Rystad field-size band. The descriptive within-region OOIP-tercile gradient (Figure G2) identifies a sub-channel of composition: even within heavy-oil operations, sanctions reallocate output from small to large fields.

Country-level descriptive statistics, including pre-period flaring intensities, production volumes, and ownership composition, are reported in Table 1.

Table 1: Descriptive Statistics

	Argentina	Brazil	Colombia	Ecuador	Mexico	Venezuela
Panel (a): Geology & Reservoir						
Water depth (m)	520.8 (540.3)	134.0 (161.6)	410.3 (829.1)	336.8 (388.5)	-146.1 (808.3)	111.0 (173.4)
Reservoir depth (m)	1872.1 (767.7)	1427.8 (736.8)	2446.3 (1235.9)	2572.5 (650.7)	3219.1 (1902.0)	2293.2 (1137.6)
Porosity (%)	7.6 (8.6)	2.8 (6.4)	6.5 (8.8)	18.2 (3.2)	4.2 (7.0)	9.8 (12.7)
Temperature (°C)	60.4 (19.2)	50.3 (18.5)	74.2 (27.2)	78.9 (16.7)	88.1 (37.5)	72.3 (36.1)
OOIP (million bbl)	106.9 (529.5)	14.7 (127.7)	63.8 (460.8)	224.8 (764.0)	130.8 (840.6)	1016.6 (5248.7)
Recovery factor (%)	35.1 (13.9)	31.7 (20.5)	28.5 (12.3)	24.7 (7.8)	45.7 (236.3)	32.1 (14.7)
Panel (b): Production & Economics (pre-period mean)						
Production (kbbl/d)	0.4 (1.7)	0.1 (0.9)	0.6 (5.7)	2.0 (7.1)	1.1 (12.3)	3.9 (20.3)
OPEX (M USD)	4.8 (21.2)	1.3 (10.8)	4.8 (42.3)	4.7 (16.2)	6.0 (52.4)	18.2 (87.8)
CAPEX (M USD)	6.9 (32.6)	1.0 (6.3)	5.0 (35.2)	9.1 (34.9)	6.1 (45.5)	14.0 (62.9)
NOC ownership share (%)	25.1 (40.8)	23.9 (42.0)	20.4 (36.7)	38.8 (47.7)	24.7 (42.7)	51.5 (41.3)
Panel (c): Field Counts						
Number of fields	1,313	1,502	1,455	265	2,132	666
Share onshore (%)	91.2%	100.0%	95.3%	95.8%	66.5%	79.9%
Panel (d): Flaring (pre-period mean, matched fields)						
Flaring intensity, 7.5 km (m^3/bbl)	48.2 (209.6)	15.3 (32.7)	25.2 (116.6)	19.3 (142.6)	63.8 (263.3)	310.8 (633.6)
Flaring intensity, 15 km (m^3/bbl)	36.3 (192.4)	28.8 (64.0)	17.6 (97.6)	14.5 (120.8)	52.3 (235.7)	195.6 (548.1)
Flaring volume, 7.5 km (M m^3)	3.4 (8.9)	6.4 (17.9)	9.9 (14.9)	22.0 (40.1)	41.8 (106.4)	110.6 (277.4)
Flaring volume, 15 km (M m^3)	2.3 (7.0)	4.5 (14.2)	7.0 (13.4)	15.8 (35.9)	30.7 (90.8)	85.7 (274.2)

Notes: Panel (a) reports field-level geological characteristics from Rystad Energy. Panel (b) reports pre-period (2012–2016) field-year averages. Standard deviations in parentheses. NOC ownership share is the top-owner share when the top owner is the national oil company (PDVSA, YPF, Petrobras, Ecopetrol, EP Petroecuador, Pemex), zero otherwise. Panel (c) counts all fields in the Rystad upstream database for each country. Panel (d) conditions on fields with positive flaring (matched subsample) during the pre-period. Brazil includes onshore fields only; Mexico includes all fields.

4.2 Satellite flaring: World Bank VIIRS

The Visible Infrared Imaging Radiometer Suite (VIIRS) night-time fire product detects high-temperature combustion sources globally at roughly 750 m resolution. The World Bank Global Gas Flaring Reduction Partnership processes these detections into annual flare-volume estimates by source location. We use both the country- year aggregates and the field-allocated point estimates. The field-level allocation is a two-stage procedure: Stage 1 fuzzy Jaro–Winkler name matching links flare-site records to Rystad field names; Stage 2 distributes the remaining flare-site volume across nearby Rystad fields by a production-weighted inverse-distance gravity scheme, $w_i = \text{Production}_i/(d_i + 100\,\text{m})^2$, normalised within each flare-year and applied within a candidate

radius of 7.5, 10, or 15 km. We report all three thresholds with the 15 km specification. Appendix C documents the algorithm in detail. Flaring intensity (m^3 flared per barrel of crude lifted) is then the field-allocated annual flare volume divided by the same field's Rystad annual production.

VIIRS captures combustion (flares) but not non-combustion methane release (vents and leaks). It therefore underrepresents the maintenance and strain channels' methane component. We use VIIRS for the main flaring outcome; methane is captured separately through TROPOMI.

4.3 Atmospheric methane: TROPOMI

The Sentinel-5P TROPOMI satellite provides daily column-averaged dry-air methane mixing ratios (XCH_4, in parts per billion) retrieved from the satellite's footprint, with native pixels of roughly 5.5×7 km. We use the operational bias-corrected gridded product, in which raw satellite readings are calibrated against ground-based reference stations to remove systematic errors that depend on surface brightness, aerosols, and atmospheric conditions, and observations are aggregated to monthly means at $0.05^\circ \times 0.05^\circ$ (≈ 5.5 km) resolution.

Analysis grid. For the regression analysis we lay a regular 10×10 km grid (defined on a UTM projection and reprojected to WGS 84 for centroid extraction) over the six-country LATAM region for cross-country robustness work, and restrict the body strain regression (eq. 5) to the 1,280 Venezuelan cells that have at least one Rystad oil field within 15 km of the cell centroid. Each *cell* is one 100 km^2 square of this analysis grid; it is the unit of observation throughout the strain section. Each cell's monthly XCH_4 value is the mean of all quality-filtered 0.05° TROPOMI pixels whose centroids fall within the cell boundary, so cells with no qualifying TROPOMI pixel in a given month are missing for that month. Cells are stable in location across months; the cell-month panel is therefore an unbalanced cell-by-time panel. The full six-country grid contains 2,588,253 cell-month observations spanning 72,439 unique analysis cells over November 2018 through December 2024 and is used in the appendix cross-country extension (Appendix E.6). Quality filtering retains 97% of raw retrievals; methane values are clipped to $[1700, 2100]$ ppb to remove sensor artifacts. Online Appendix E documents the TROPOMI instrument, our use of the Level 3 (rather than Level 2) product, the cell-grid construction, and a side-by-side comparison with the atmospheric-inversion methodology of Nathan et al. (2024).

Cell-to-pipeline distance. For each cell we compute the great-circle distance from the cell centroid to the nearest pipeline segment in the OSM-plus-GEM combined network. We then call a cell "near pipelines" at distance band $d \in \{10, 20, 50, 100\}$ km if that nearest-pipeline distance is below d. The 20 km band is our favourite specification; the others are reported as sensitivity. The cell-area on which pollution is integrated is the 100 km^2 square itself; d controls only which cells are counted as near-pipeline, not the area over which methane is averaged at each cell. Pipeline-vicinity methane responds to both compressor-station leakage and corridor venting, providing the spatial gradient our triple-difference exploits.

4.4 Biodiversity outcomes: eBird

eBird is a citizen-science platform run by the Cornell Lab of Ornithology. The unit of submission is a *checklist*: a record of all bird species an observer (or small group) detected on a single survey, together with the survey's start time, duration, distance travelled, GPS coordinates, and an explicit "complete" flag indicating that the observer reported every species seen and heard, rather than only highlights. Each checklist is therefore one spatially geo-referenced bird-survey observation, with effort attributes (duration, distance, traveling indicator) used as controls. We tag each checklist to the same 10×10 km analysis cell as the TROPOMI panel by point-in-polygon assignment of the checklist coordinate, so a "cell" has the same meaning across pollution and biodiversity outcomes. For the LATAM panel we obtain approximately 11.5 million checklists, 2012–2024. We construct two types of outcome: *species richness* (the count of distinct species reported on the checklist, a per-checklist measure of community breadth) and *abundance* (the total individuals counted across species on the checklist, a per-checklist measure of community size); for the species-level guild decomposition in Section 7.2 we restrict to the 265 species with ≥ 5 checklist detections in the Chevron event-study window. The distribution of checklists per observer is checked in Appendix H: $\geq$2-checklist observers account for 99.8% of the OOIP-gradient sample and 98.9% of the Chevron quasi-experiment sample, supporting the inclusion of observer fixed effects in the main specifications.

Birds integrate multiple environmental stressors at small spatial scales: their abundance responds to flaring-driven thermal disturbance, methane and associated VOC emissions, particulate deposition from incomplete combustion, and habitat disturbance from operational activity. We use bird outcomes as the integrating measure that captures the joint local externality of all three channels at the ecologically relevant spatial scale, while flaring (VIIRS) and methane (TROPOMI) provide channel-specific quantitative anchors.

4.5 Pipeline network: OSM + GEM

The OpenStreetMap (OSM) global pipeline layer is community-edited and includes both major trunk lines and gathering systems. The Global Energy Monitor (GEM) Oil and Gas Pipelines Database is curated and emphasises operational major-trunk infrastructure. Each has known gaps. We union the two sources, deduplicating overlapping segments within 200 m, and clip to the six-country panel boundaries. Total combined coverage by country: Argentina 36,302 km, Brazil 30,145 km, Colombia 10,736 km, Ecuador 3,331 km, Mexico 31,189 km, Venezuela 11,600 km. We benchmark coverage against the CIA World Factbook's most recently published country totals (sum of crude-oil, oil-products, and natural-gas pipeline kilometres, 2013–2017 vintages).[7] We

[7] Combined OSM+GEM kilometres against the CIA Factbook total: Argentina 36,302 vs. 39,809 (91%), Brazil 30,145 vs. 21,896 (138%; OSM includes gathering systems the Factbook excludes), Colombia 10,736 vs. 15,711 (68%), Ecuador 3,331 vs. 3,179 (105%), Mexico 31,189 vs. 35,429 (88%), Venezuela 11,600 vs. 14,669 (79%). The comparison is approximate because the Factbook reports long-distance transmission pipelines (the trunk lines that move crude, refined products, and natural gas between basins, refineries, and export terminals) while OSM and GEM also include local gas-distribution networks (the city-level pipes that deliver natural gas from trunk-line city-gates to households and businesses) and field

report sensitivity to the OSM-only specification throughout (which removes the GEM curation but yields the same qualitative pattern).

4.6 Spatial matching

Each VIIRS flare detection is allocated to nearby Rystad fields via the two-stage name-plus-production-weighted-distance algorithm described in Section 4.2 and detailed in Appendix C. Each TROPOMI grid cell is tagged "near oil field" if any Rystad field falls within 15 km of the cell centroid, and "near pipeline" if any pipeline segment falls within d km. eBird checklists are tagged with the same proximity codes at the checklist coordinate. Field-pair distance thresholds are varied as a robustness check; the substantive findings are stable across all three. Diagnostic checks on match quality, multi-field overlap, and unmatched-flare residual are in Appendix C.

4.7 Identification overview

Each of the three channels is identified by a different empirical strategy that exploits a different source of variation in the joint panel.

Composition (Section 5). A difference-in-differences on field-level Venezuelan production (Q_{ift}) against the five-country no-sanctions LATAM control, 2012–2024, with field and year fixed effects and SE clustered at the field level. The within-Venezuela cross-sectional source of variation is the regional partition (Orinoco East / Maracaibo / Barinas / Other) and the OOIP-tercile partition within each region. Composition-driven changes in average per-barrel intensity are recoverable from this design.

Maintenance (Section 7). The November 2022 Chevron OFAC General License 41 generates an unanticipated, sharply dated, partial restoration of foreign maintenance budgets at four specific Venezuelan joint-venture fields. We compare these fields, and the cells around them, to otherwise comparable Venezuelan fields and cells not covered by the license, in particular to cells near other Western-major joint ventures. Cell-level outcomes (flaring intensity, methane, bird abundance) are estimated in event-study form in semi-annual bins around the licence date, with cell, year-month, and observer fixed effects.

Strain (Section 6). The body strain specification (equation (5)) is a within-Venezuela cell + year-month fixed-effects regression of log XCH_4 on the log spot price of naphtha interacted with a near-pipeline indicator, restricted to the Venezuelan Orinoco cell-month sample. TROPOMI Level 3 coverage begins in November 2018, so the time variation we identify is post-2018 high-frequency

gathering lines (the small-diameter pipes that collect crude and associated gas from individual wells and feed them into trunk lines). Brazil's 138 % coverage ratio illustrates the gap: OSM picks up the country's extensive gas-distribution network, which the Factbook treats as a separate category. We treat coverage above two-thirds of the Factbook total as adequate for the spatial-matching analysis that follows.

variation in naphtha price rather than a pre/post comparison around the 2017 sanctions. The appendix runs the analogous regression on the full six-country LATAM grid as a triple-difference (Venezuela versus LATAM controls $\times$ near-pipe versus far-from-pipe $\times$ time trend) at four distance bands $d \in \{10, 20, 50, 100\}$ km; the within-Venezuela near-pipeline trend remains positive and significant across all of them.

4.8 Ownership and operator categories

For the maintenance and strain analyses we partition Venezuelan fields by operator into six bins: Chevron joint-ventures (4 fields, the License 41 set: PetroBoscan, PetroIndependencia, PetroIndependiente, and PetroPiar); other Western majors (5 fields, primarily Shell-Maraven legacy assets); Chinese state operators (8 fields, CNPC and Sinopec); Russian operators (11 fields, primarily Rosneft-Junta legacy); PDVSA full ownership (78 fields); and other operators (38 fields, mostly small Western JVs and private Venezuelan firms). The Chevron quasi-experiment is a *within-Venezuela between-operator* comparison: we contrast Chevron-near cells to Western-JV-near cells in the same post-licence window, so the licence-induced OPEX restoration at the four covered fields is identified off variation across otherwise similar Venezuelan operators rather than off within-firm Chevron variation across countries.

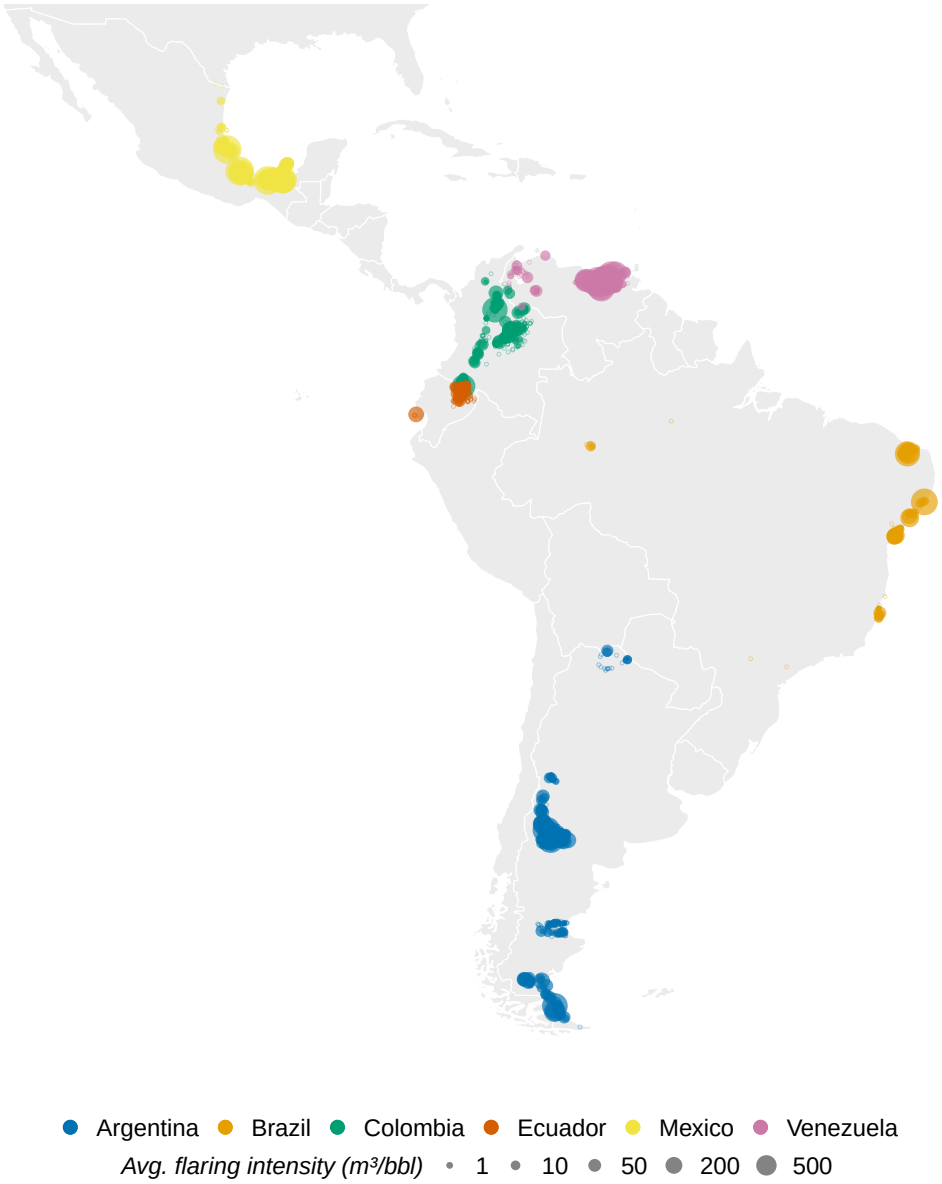
Spatial distribution of fields and flaring. Figure 1 maps the spatial distribution of the 1,874 Rystad-tracked oil fields and their average annual flaring intensities across the six-country panel. Figure 2 traces the country-level evolution of total flaring volume, flaring intensity, production, and operating expenditure over 2012–2024, with shaded periods marking the Brent-price super-cycle (gas-flaring intensity is dimensionally detached from oil-price dynamics in every country except Venezuela). The two figures together establish the descriptive setup that the empirical sections will exploit: Venezuela's pre-existing high flaring intensity, the post-2017 sharp deviation from the non-sanctioned Latin American trajectory, and the asymmetry between the Orinoco and Maracaibo bases that the composition channel will formalise.[8]

5 The composition channel

The composition channel predicts that when sanctions raise the fixed cost of operating, the surviving production base concentrates on assets that can absorb those fixed costs: large fields, heavy crude operations whose gas-handling equipment cannot easily be turned down, and operators with international maintenance access. We test the three margins in turn against the cross-country LATAM panel, using Argentina, Brazil onshore, Colombia, Ecuador, and Mexico as unsanctioned

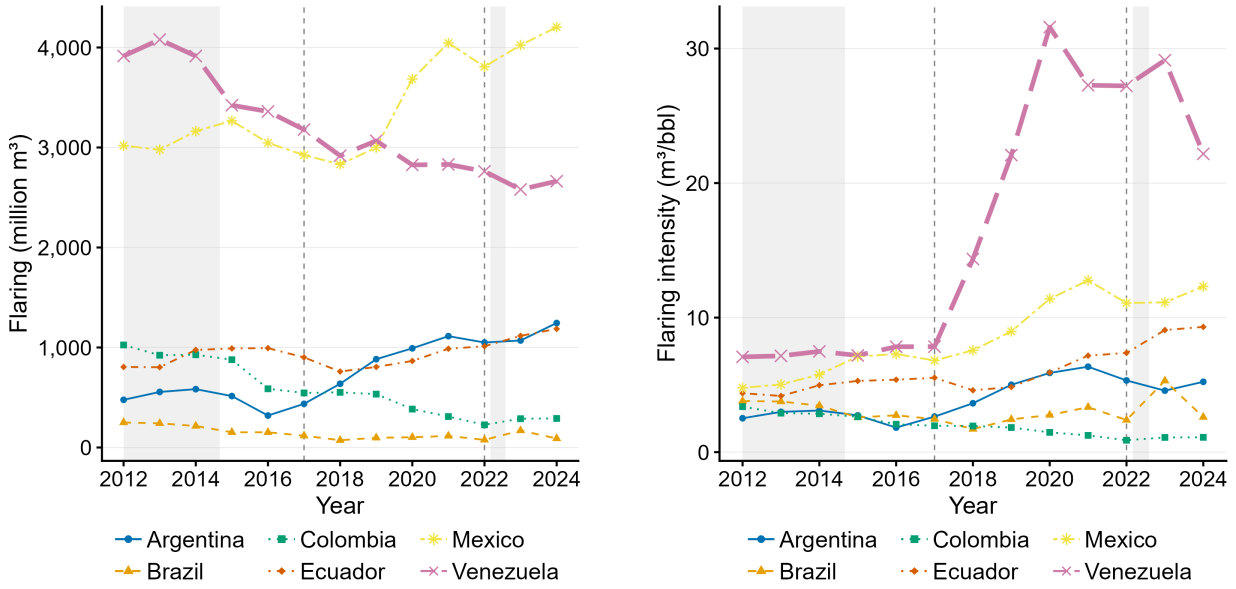
[8] Our six-country LATAM panel is approximately representative of the world's unsanctioned heavy-oil operating environments outside the Middle East: it includes the Mexican Maya basin, the Ecuadorean Oriente, and the Colombian Llanos. It does not cover Canadian oil sands, where the production process is very different, or Middle Eastern conventional crude.

Figure 1: Oil Fields and Flaring Intensity in Latin America



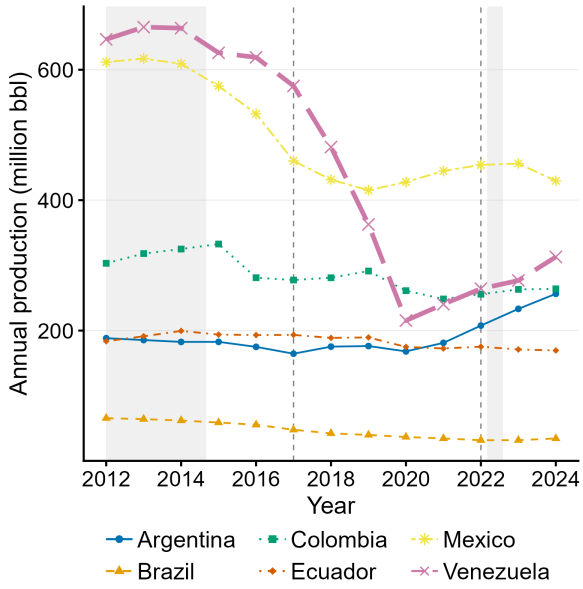
Notes: Each marker represents a Rystad Energy oil field in one of the six sample countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela). Marker color indicates the field's average annual flaring intensity (m^3/bbl), averaged over all field-years in 2012–2024, computed as allocated VIIRS flaring volume divided by annual crude oil production (15 km allocation radius). Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker.

Figure 2: Evolution of Flaring, Production, and Investment by Country

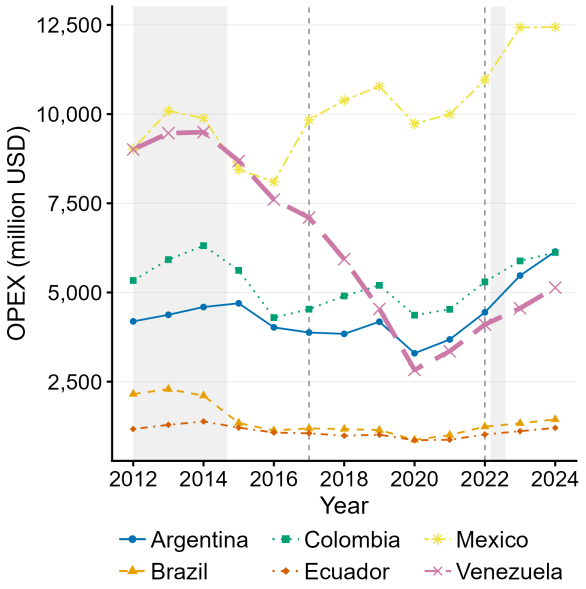


(a) Flaring

(b) Flaring / Production



(c) Production



(d) OPEX

Notes: Raw country-level aggregates (not regression-based) for the six sample countries, 2012–2024. Panel (a): total VIIRS flaring volume allocated to Rystad fields within a 15 km radius (million m^3). Panel (b): production-weighted mean flaring intensity (m^3/bbl, 15 km radius), winsorised at the 99.5th percentile. Panel (c): total annual crude oil production (million bbl). Panel (d): total annual operational expenditure (million USD). Brazil restricted to onshore fields. Vertical dashed lines indicate 2018 (first full year of financial sanctions) and 2022 (OFAC License 41 granted to Chevron). Gray shaded bars indicate periods when the Brent crude oil price exceeds $100/barrel. Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker.

controls. Each margin uses the same Venezuela × Post difference-in-differences design as Table D1 (equations (2) and (3)) but varies the dependent variable to isolate the relevant compositional shift. The three subsections below cover (i) reallocation toward larger fields, (ii) reallocation toward heavier crude, and (iii) reallocation across operators.

5.1 Reallocation toward larger fields

The first composition margin is asset size. When fixed operating costs jump, small fields whose pre-sanctions margins barely covered flat costs exit first, while large fields can spread the same flat-cost burden across more barrels and remain profitable. The clean way to test this within Venezuela is to ask whether, after 2017, the producing-rate of small fields fell by more than the producing-rate of large fields. We estimate

$$\mathbb{1}\{\text{prod}_{it} > 0\} \;=\; \alpha_i + \tau_t + \beta\,(\text{Post}_t \cdot \text{Large}_i^s) + \varepsilon_{it}, \tag{2}$$

where i indexes Venezuelan fields, t calendar years over 2012–2024, $\text{Post}_t = \mathbb{1}\{t \geq 2017\}$, and Large_i^s is a time-invariant indicator that field i's pre-period (2012–2016) value of size measure s exceeds the Venezuelan pre-period median. We try three size measures $s \in \{\text{OOIP, mean pre-period production, Rystad field-size group}\}$. Field fixed effects α_i absorb time-invariant heterogeneity (including Large_i^s itself); year fixed effects τ_t absorb common shocks (including Post_t and the 2014–2016 oil-price collapse). The sample is the ever-producing Venezuelan field-years, so a field that exits contributes its post-period zeros and identifies β through the within-field producing-rate change. Standard errors are clustered at the asset level. The composition prediction is $\beta > 0$: large fields keep producing more often than small fields after sanctions; small fields drop more.

Table 2 reports the estimates. Across all three size measures, β is positive and significant at the 1% level: $+0.21$ for OOIP, $+0.15$ for mean pre-period production, and $+0.24$ for the Rystad size-band cut (300+ MMboe). Read as a composition magnitude: post-2017, large Venezuelan fields are 15–24 percentage points more likely to be producing than small Venezuelan fields, conditional on field and year fixed effects. Equivalently, small fields stopped producing by 15–24 pp *more* than large fields once sanctions hit. The pre-period producing rates were essentially balanced (0.89–0.99 for both groups), so the post-period gap is the entire effect.

The within-VEN comparison absorbs anything that hits all Venezuelan fields uniformly (PDVSA central-procurement frictions, exchange-rate distortions, refinery shutdowns, the 2014–2016 oil-price collapse) into the year fixed effects. What remains is a heterogeneous, size-graded response: when the system was shaken by sanctions, the producing base concentrated on the larger fields. Appendix Figure D1 reports the event-study version with year-by-year Post × Large coefficients normalised to zero in 2016: the gap opens at 2017–2019 and widens through 2024, with no parallel-trends violation in the pre-period.

Table 2: Composition Channel within Venezuela: Producing-rate Change, Large vs. Small Fields

Dependent variable: $\mathbb{1}\{\text{producing}\}$	Large by OOIP (1)	Large by pre-prod (2)	Large by size group (3)
Post × Large (β)	0.2098***	0.1470***	0.2422***
	(0.0413)	(0.0419)	(0.0316)
Pre-period producing rate, large fields	0.955	0.988	0.971
Pre-period producing rate, small fields	0.917	0.887	0.926
Fixed effects			
Asset	✓	✓	✓
Year	✓	✓	✓
Sample	VEN ever-producing	VEN ever-producing	VEN ever-producing
Observations	3,042	3,042	3,042
Fields (large / small)	119 / 115	114 / 120	49 / 185
R^2	0.54132	0.53227	0.53924

Notes: This table tests the composition channel within Venezuela: post-sanctions, do small Venezuelan fields stop producing more than large ones? The specification is $\mathbb{1}\{\text{prod}_{it} > 0\} = \alpha_i + \tau_t + \beta\,(\text{Post}_t \cdot \text{Large}_i^s) + \varepsilon_{it}$ (equation (2)). Sample: ever-producing Venezuelan fields, 2012–2024. $\text{Post}_t = \mathbb{1}\{t \geq 2017\}$. Large_i^s is a time-invariant indicator that field i's pre-period (2012–2016) value of size measure s exceeds the Venezuelan pre-period median: (1) original-oil-in-place (OOIP, MMbbl, fields with positive OOIP); (2) mean pre-period annual crude production (bbl, fields with positive pre-period output); (3) Rystad field-size group, with 'large' defined as the top two bands (*300–1000 MMboe* and *above 1000 MMboe*). Field fixed effects α_i absorb time-invariant heterogeneity (including Large_i^s); year fixed effects τ_t absorb common time shocks (including Post_t and the 2014–2016 oil-price collapse). β is the differential post-sanctions producing-rate change for large vs. small Venezuelan fields. $\beta > 0$ means large fields are β percentage points more likely to be producing post-2017 than small fields, equivalently small fields stopped producing β pp more than large ones—the composition channel. Standard errors clustered at the asset level in parentheses. ***p<0.01, **p<0.05, *p<0.1. Source: Rystad Energy UCube.

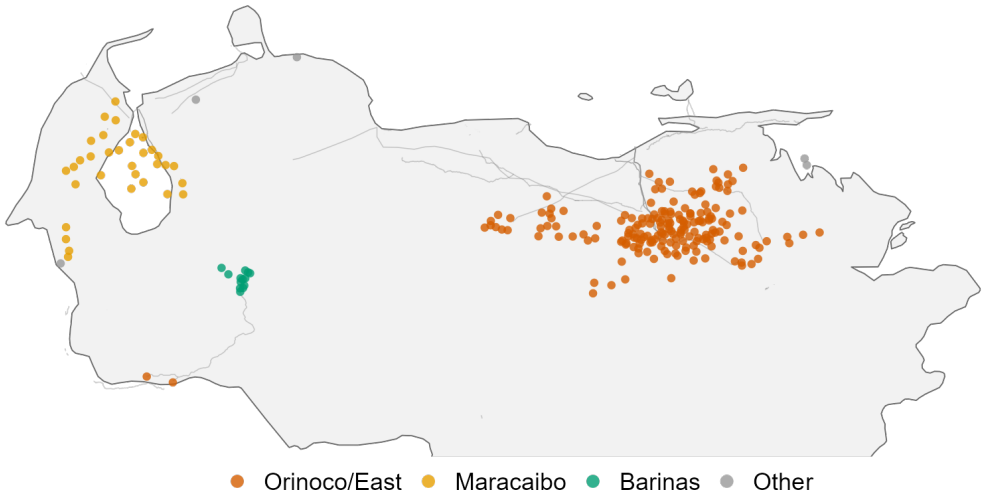
5.2 Reallocation toward heavier crude

The second composition margin is crude heaviness. Venezuelan upstream operations split into two very different gas-handling setups: the Orinoco Belt's 8–12° API extra-heavy crude must be blended with imported naphtha (a light hydrocarbon diluent) to flow through the pipelines and through the equipment that recovers the diluent for re-use, while the Maracaibo basin's medium-API operations need none of that machinery (Section 3). Sanctions cut the naphtha import flow, so the heavy operations should adjust much more sharply than the lighter ones.

In 2016, Venezuelan crude production was approximately 2.28 mbbl/d (Rystad UCube), distributed across the Orinoco extra-heavy belt, the Maracaibo medium-API basin, and a smaller residual of Barinas, eastern, and offshore fields. EIA / Monaldi et al. (2021) attribute roughly 51% of crude output to Orinoco, 30% to Maracaibo, and 19% to the residual at this peak; the Rystad panel we use throughout the regressions allocates production somewhat differently across regions, but the qualitative ordering (Orinoco largest, Maracaibo second, residual smallest) is the same. By 2023, total production had fallen to approximately 0.78 mbbl/d, a $\approx$ 66% decline. Within Venezuela, the post-period reallocation favoured Orinoco: its share of the (smaller) total rose, while Maracaibo's share fell and the residual held roughly flat. The cross-region asymmetry is consistent

with the heaviness-channel prediction: the basin hosting large high-cost fields whose gas-handling equipment cannot easily be turned down retained a disproportionate share, while the basin with smaller conventional fields cut production cleanly. Figure 3 maps the four Venezuelan regions used below.

Figure 3: Venezuelan Oil Fields by Geographic Region



Notes: Each marker represents a Venezuelan oil field with strictly positive crude production in at least one year of the 2012–2024 panel ($n = 234$ fields); fields that did not produce during the sample period are not shown. Fields are classified into four geographic regions based on the Rystad basin assignment: Orinoco/East (East Venezuela Basin and Llanos), Barinas, Maracaibo, and Other. Grey lines show oil and gas pipeline infrastructure from OpenStreetMap and the Global Energy Monitor. Source: Rystad Energy UCube; OpenStreetMap; Global Energy Monitor.

The cross-country regression test interacts Post with indicators for VEN-Heavy and VEN-Non-Heavy fields, where VEN-Heavy is defined as the Orinoco Belt or any Venezuelan field with heavy-crude share $> 50\%$:

$$Y_{ift} = \alpha_i + \tau_t + \beta_h(\text{VEN}_i^{\text{Heavy}} \times \text{Post}_t) + \beta_n(\text{VEN}_i^{\text{Non-Heavy}} \times \text{Post}_t) + \varepsilon_{ift}, \tag{3}$$

with reference category the non-Venezuelan LATAM control fields. The outcome Y_{ift} varies across the eight columns of Table 3: flaring volume at 7.5 km and 15 km, flaring intensity (m^3/bbl) at 7.5 km and 15 km, $\log(1 + \text{production})$, $\log(1 + \text{OPEX})$, $\log(1 + \text{CAPEX})$, and $\log(1 + \text{cumulative CAPEX})$. Field and year fixed effects and asset-clustered standard errors match equation (2); the sample is the six-country LATAM panel restricted to ever-flaring fields, with the per-DV refinement that the flaring-volume and flaring-intensity columns are restricted to field-years with positive flaring at the relevant radius and the production / OPEX / CAPEX columns use the full ever-flaring sample.

Table 3: Sanctions Effects by Crude Oil Heaviness

Dependent Vars:	Flaring Vol. 7.5km (1)	Flaring Vol. 15km (2)	Flaring Int. 7.5km (3)	Flaring Int. 15km (4)	Prod. (log) (5)	OPEX (log) (6)	CAPEX (log) (7)	Cum. CAPEX (log) (8)
Heavy oil (Venezuela) × Post	-5.9228	-2.0061	176.6656***	125.0990***	-0.6578***	-0.6130***	-0.4037	-0.4008***
	(10.9280)	(6.8320)	(59.0852)	(47.0277)	(0.1463)	(0.1482)	(0.3240)	(0.1541)
Non-heavy oil (Venezuela) × Post	-35.3138**	-19.8566	-37.5629	-0.8748	-0.7207	-0.4529	-0.4993	-0.5664***
	(16.8961)	(18.7968)	(46.5616)	(11.9262)	(0.5139)	(0.3739)	(1.0925)	(0.1145)
Fixed effects								
Asset	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓
Mean dep. var.	29.8540	22.0782	76.9685	55.9316	12.0406	14.9673	12.7102	14.8554
Observations	5,023	7,795	5,023	7,795	11,983	11,983	11,983	11,983
Fields	772	1,097	772	1,097	1,097	1,097	1,097	1,097
R^2	0.90678	0.93667	0.68533	0.67789	0.86183	0.83241	0.65119	0.91986

Notes: This table tests heterogeneity in the sanctions effect by crude oil heaviness, interacting Post with indicators for heavy and non-heavy Venezuelan fields (regression (3)). Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Post $= \mathbb{1}\{\text{year} \geq 2017\}$. VEN Heavy: fields in the Orinoco Belt (East Venezuela Basin or Llanos) or with heavy crude share $> 50\%$ (Extra Heavy, Heavy 15–19, Heavy 20–22, Bitumen). VEN Non-Heavy: remaining Venezuelan fields. Reference group: non-Venezuelan LATAM fields. Cols (1)–(2): allocated flaring volume (million m^3). Cols (3)–(4): flaring intensity (m^3/bbl), winsorized at the 99.5th percentile. Col (5): log(1 + crude production in bbl). Col (6): log(1 + OPEX in USD). Cols (7)–(8): log(1 + CAPEX in USD) and log(1 + cumulative CAPEX in USD). Cols (1) and (3) restrict to field-years with positive flaring at 7.5 km; Cols (2) and (4) at 15 km; Cols (5)–(8) use the full sample. All models include field and year fixed effects. Standard errors clustered at the asset level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: Rystad Energy; World Bank VIIRS.

Table 3 reports the estimates. The heavy-vs-non-heavy asymmetry is sharp. At the 15 km flare-allocation radius, $\hat{\beta}_h = +125.1$ m^3/bbl ($p < 0.01$) at heavy Venezuelan fields against a pre-period mean of 56.5 m^3/bbl, roughly a tripling of per-barrel flaring; the non-heavy coefficient is statistically indistinguishable from zero ($\hat{\beta}_n = -0.87$, n.s.). Production and OPEX both fall by $\approx$ 46–48% at heavy fields; CAPEX shows a similar ordering but with wider standard errors. This is the empirical counterpart of the naphtha story laid out in Section 3: only fields whose gas-handling depends on imported naphtha are sharply affected.[9] Appendix Figure D2 reports the flaring-intensity event-study by Venezuelan geographic region together with the unsanctioned controls (Mexican Maya basin, Brazilian onshore): the Orinoco intensity break appears at 2017–2019; the Mexican Maya declines smoothly throughout under Cantarell reservoir maturation, a profile predicted by geology rather than by any policy shock; Brazilian onshore is approximately flat. Neither control replicates the Venezuelan break.

5.3 Reallocation across operators

The third composition margin is across operators. Venezuelan upstream operations are organised through joint ventures between PDVSA and international partners. International partners hold maintenance and procurement networks that PDVSA-only operations lack, and the value of those networks rises sharply when sanctions cut PDVSA's direct procurement access.

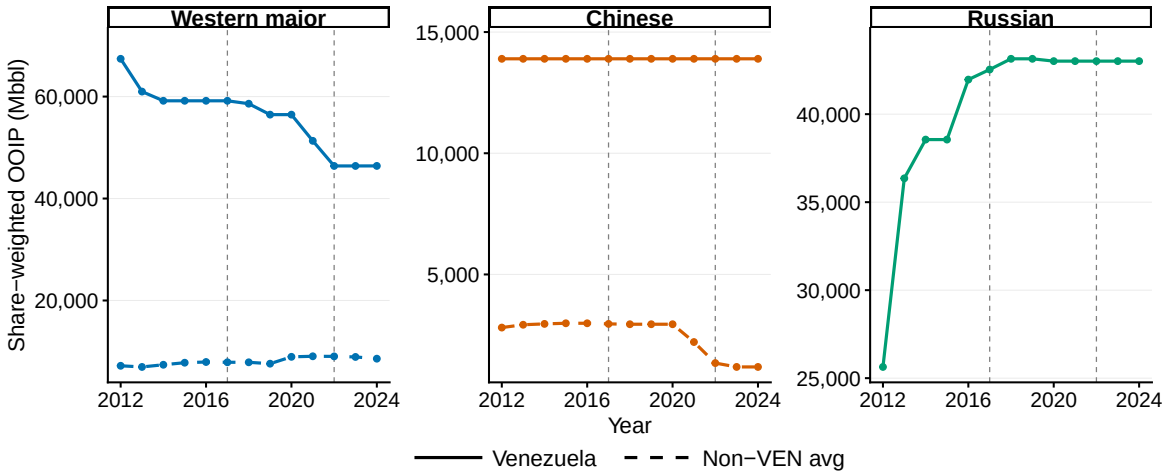
We allocate each field-year's oil-in-place and flaring volume to operators in proportion to their equity share, then sum within international-partner group. So a field that is 40% Rosneft and 60% PDVSA contributes 40% of its OOIP and 40% of its allocated VIIRS flaring to the Russian-partnered series each year. This avoids reclassification artefacts (e.g. a partner exit shifting a field's entire reserve base to a different bin overnight) and produces series that respond *discretely* to the named policy events—August 2017 (E.O. 13808) and November 2022 (OFAC License 41 to Chevron)—not gradually. Both figures restrict to the three international-partner groups (Western including Chevron, Chinese, Russian); the residual is held by PDVSA directly.

Figure 4 traces the share-weighted oil-in-place. Inside Venezuela, the Western-partnered reserve base falls steeply at 2017–2019 (the OFAC Section 5(b) wind-down that removes ConocoPhillips, ExxonMobil, and others) and only partially recovers from late 2022 with OFAC License 41 to Chevron; the Russian-partnered reserve base rises smoothly from 2017 onward; and the Chinese-partnered reserve base is essentially *flat* throughout. The cross-country contrast is what stands out. Chinese-partnered reserves in the non-VEN LATAM controls (dashed line) decline visibly post-2017—Chinese majors trim their LATAM exposure—while their Venezuelan reserves stay put. Russian-partnered fields are essentially absent from the non-VEN controls throughout the panel, so the entire Russian series is a within-Venezuela phenomenon. Together: Chinese partners did not exit Venezuela the way they exited the rest of LATAM, and Russian partners stepped into the

[9] The unconditional Venezuela $\times$ Post DiD on the same outcomes, without the heavy/non-heavy split, is reported in Online Appendix Table D1. We use the heavy/non-heavy specification of equation (3) as the main-text table because the unconditional version pools two production regimes that respond very differently to the same policy shock.

equity gaps left by the Western wind-down.

Figure 4: Oil-in-Place by Ownership Group



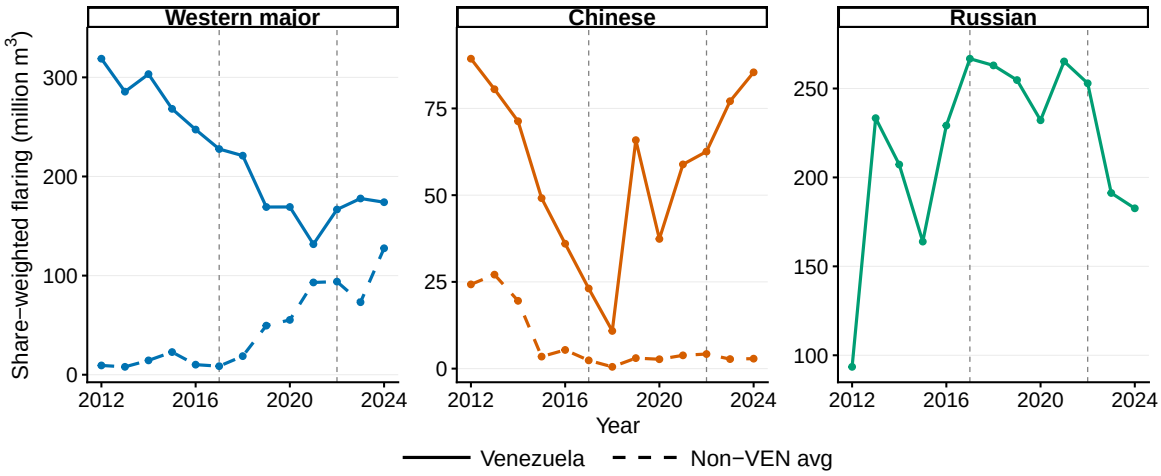
Notes: Share-weighted original oil-in-place (MMboe) attributed to each international-partner group, by year: Western majors (including Chevron, ExxonMobil, Shell, BP, Total, Eni, Equinor, ConocoPhillips, Repsol), Chinese (CNPC, CNOOC, Sinopec, PetroChina), and Russian (Rosneft, Gazprom, Lukoil). Solid lines: Venezuela; dashed lines: average across Argentina, Colombia, and Ecuador. The reserve base is apportioned in proportion to each operator's equity stake in each field-year, avoiding reclassification artifacts from discrete ownership changes. Vertical dashed lines indicate August 2017 (E.O. 13808) and November 2022 (OFAC License 41 to Chevron). Source: Rystad Energy UCube.

Figure 5 repeats the share-weighting exercise for allocated VIIRS flaring volumes. The same equity-share rule applies: a field that is 40% Rosneft contributes 40% of its flaring to the Russian-partnered series each year. The patterns mirror Figure 4 on the equity side and amplify them on the emissions side. Russian-partnered flaring in Venezuela rises from ≈ 94 Mm^3 in 2012 to ≈ 183 Mm^3 by 2024, roughly doubling. Two forces compound to produce that doubling: the Russian-partnered *share* of Venezuelan reserves rises (Figure 4) through the post-2017 equity transfers as Western majors withdraw, and the Russian-partnered fields are heavily concentrated in the Orinoco Belt's extra-heavy operations (Petromonagas in particular), where per-barrel flaring intensity is high in absolute terms and rises sharply post-sanctions (Section 5.2). Russian-partnered flaring therefore rises through both the size and the heaviness margins already documented.[10] Chinese-partnered Venezuelan flaring is roughly flat throughout. Western-partnered Venezuelan flaring falls during the 2017–2019 wind-down and recovers partially from late 2022 with Chevron's licence restoration. Outside Venezuela, the dashed lines tell their own stories that are unrelated to sanctions: Chinese-partnered flaring in the LATAM controls collapses (the same exit pattern visible in the OOIP figure), and Russian-partnered flaring is essentially zero throughout (no Russian joint ventures in the

[10]The geographic-clustering pattern behind these dynamics is sharp: among the 233 Venezuelan fields in the analysis sample, 91% of Russian-partnered fields, 100% of Chinese-partnered, 92% of Other-JV, 67% of Chevron, and 60% of Western-Major fields lie in the Orinoco extra-heavy region, against 95% for PDVSA-only fields. International JVs also concentrate among the largest, most capital-intensive heavy fields. Online Appendix Figure G1 maps each Venezuelan analysis-sample field by ownership bin and heavy-crude status. The collinearity between ownership, location, size, and heaviness motivates the size-by-ownership decomposition in Online Appendix G, which holds size fixed and asks whether joint-venture giants behave differently from PDVSA-sole giants.

LATAM control countries).

Figure 5: Allocated Flaring Volume by Ownership Group



Notes: Share-weighted allocated VIIRS flaring volume (million m^3) attributed to each international-partner group, by year. The flaring volume of each field-year is allocated to operators in proportion to their equity stake (the same rule as Figure 4 applied to flaring instead of oil-in-place). Field-level flaring is the World Bank VIIRS allocation at the 15 km radius, restricted to Onshore fields and to the six LATAM countries. Solid lines: Venezuela; dashed lines: average across Argentina, Colombia, and Ecuador. Vertical dashed lines: August 2017 (E.O. 13808) and November 2022 (OFAC License 41 to Chevron). Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker. The corresponding flaring-intensity-by-owner panel (which divides through by share-weighted production) is reported in Online Appendix D.

The cross-country regression test interacts Post with each of four ownership-group indicators applied to Venezuelan fields:

$$Y_{ift} = \alpha_i + \tau_t + \sum_g \beta_g (\text{VEN}_i^g \times \text{Post}_t) + \varepsilon_{ift}, \tag{4}$$

with $g \in \{$PDVSA-only, Russian, Chinese, Chevron, Western Majors, Other (Venezuela)$\}$, ownership bins assigned using 2012–2016 baseline participation shares with the precedence rule of Table 4, and reference category the non-VEN LATAM fields. Field and year fixed effects and asset-clustered standard errors match equation (2); the sample is the ever-flaring six-country panel of equation (3) intersected with the geo-classification sample used to define the ownership bins.

Table 4 reports the estimates. The flaring-intensity rise and production decline concentrate at PDVSA-only fields and at Western-partnered fields during the 2017–2019 wind-down; Russian and Chinese partnered fields show smaller production declines, consistent with their international partners maintaining procurement access. The Chevron license restoration in November 2022 reverses the Western-partnered decline at the four Chevron joint ventures (PetroPiar, PetroBoscan, PetroIndependiente, PetroIndependencia). We exploit this November 2022 shock as a within-Venezuelan, within-firm quasi-experiment in Section 7. Appendix Figure D3 reports the event-study dynamics by ownership group, showing the same ordering and confirming flat pre-trends.

Table 4: Sanctions Effects by Ownership Group

Dependent Vars:	Flaring Vol. 7.5km (1)	Flaring Vol. 15km (2)	Flaring Int. 7.5km (3)	Flaring Int. 15km (4)	Prod. (log) (5)	OPEX (log) (6)	CAPEX (log) (7)	Cum. CAPEX (log) (8)
PDVSA × Post	-5.1397	8.3450	133.9061**	89.4193	0.0023	0.1308	0.4462	-0.5996***
	(17.0493)	(9.6604)	(67.8889)	(56.0294)	(0.0748)	(0.0987)	(0.3084)	(0.0971)
Russia × Post	2.5803	-5.9854	248.0255**	265.8159**	-0.8476**	-1.3402***	0.8826***	-0.8511***
	(21.9990)	(13.2023)	(116.6912)	(128.9822)	(0.3644)	(0.3659)	(0.2981)	(0.1626)
China × Post	-6.0258	-6.0490	1.8926	-13.7759	-2.5627**	-2.3789**	-3.0399	0.7171
	(6.8535)	(3.7854)	(25.1289)	(9.7026)	(1.1600)	(1.0289)	(2.0415)	(0.8311)
Chevron × Post	-0.2433	-76.1697	-21.4096**	-22.5433***	0.2548	0.0594	0.1413	-0.3162
	(2.6263)	(55.2103)	(9.0177)	(6.7081)	(0.4019)	(0.4070)	(0.2855)	(0.2019)
Western Majors × Post	-13.3526	-11.8381	350.2150	290.0557	-0.4529	-0.5995*	-3.4294**	-0.4695***
	(30.0468)	(24.3866)	(304.5394)	(261.8658)	(0.4289)	(0.3210)	(1.6041)	(0.0841)
Other (Venezuela) × Post	-30.0564**	-17.2802**	152.9799	130.4584	-1.9304***	-1.8506***	-1.8807**	-0.1400
	(11.8059)	(6.8386)	(148.7432)	(107.2288)	(0.3620)	(0.3302)	(0.7387)	(0.4645)
Fixed effects								
Asset	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓
Mean dep. var.	29.8540	22.0782	76.9685	55.9316	12.0406	14.9673	12.7102	14.8554
Observations	5,023	7,795	5,023	7,795	11,983	11,983	11,983	11,983
Fields	772	1,097	772	1,097	1,097	1,097	1,097	1,097
R^2	0.90686	0.93713	0.68611	0.67933	0.86573	0.83661	0.65263	0.92003

Notes: This table tests heterogeneity in the sanctions effect by ownership group, interacting Post with indicators for six Venezuelan ownership bins (regression (3)). Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Post $= \mathbb{1}\{\text{year} \geq 2017\}$. Ownership bins are assigned using pre-2017 baseline participation shares with the following precedence: PDVSA ($\geq$ 99% PDVSA participation), Russia, China, Chevron, Western Majors, Other (any firm from each group $>$ 0%). Reference group: non-Venezuelan LATAM fields. Cols (1)–(2): allocated flaring volume (million m^3). Cols (3)–(4): flaring intensity (m^3/bbl), winsorized at the 99.5th percentile. Col (5): log(1 + crude production in bbl). Col (6): log(1 + OPEX in USD). Cols (7)–(8): log(1 + CAPEX in USD) and log(1 + cumulative CAPEX in USD). Cols (1) and (3) restrict to field-years with positive flaring at 7.5 km; Cols (2) and (4) at 15 km; Cols (5)–(8) use the full sample. All models include field and year fixed effects. Standard errors clustered at the asset level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: Rystad Energy; World Bank VIIRS.

6 The strain channel: methane and biodiversity

The strain channel predicts that surviving fields, pushed closer to the maximum throughput their gas-handling equipment was designed for, vent and leak more gas per barrel. The diagnostic is *atmospheric methane*: when a heavy-oil gathering system overloads, gas that the system can no longer route to a sales pipeline, reinjection compressor, or upgrader feed escapes through flange leaks, vapour-recovery bypasses, and emergency vents (Marks, 2022). Methane is the right outcome for two reasons. First, venting produces no thermal anomaly, so VIIRS flare-detection (the central outcome of Section 5) *misses* it; only the column-averaged atmospheric concentration shows it. Second, methane is the proximate non-CO_2 greenhouse gas in the social cost of carbon literature (International Energy Agency, 2024; Plant et al., 2022), so the magnitude maps directly into externality-equivalent dollars.

We measure methane with TROPOMI, using the same 10×10 km analysis grid and pipeline-distance classification defined in Section 4.3. Restricted to Venezuela, the regression sample comprises 25,520 cell-months spanning 1,280 analysis cells over November 2018 through December 2024 (Veefkind et al., 2012).[11]

Figure 6 shows the geography on which the analysis runs. The black lines are oil and gas pipelines, the yellow-to-red circles are mean cell-level XCH_4 over 2018–2024 (brighter shades = higher concentrations), and the purple circles are NASA FIRMS VIIRS active-fire detections (the flare sites). The pipeline corridors trace the path between Maracaibo oilfields, the Orinoco Belt heavy-oil operations, refineries, and export terminals; methane hotspots co-locate with the corridors and especially with wellhead clusters in the Orinoco Belt and around Lake Maracaibo. The biodiversity sample we use later in this section (Venezuelan eBird checklists; sample statistics in Online Appendix H, Table H1) is concentrated in the populated north and along the same pipeline corridors, so cell-level pollution exposure varies sharply across cells while ecological observations are dense enough to support cell-month identification.

6.1 Exploiting high-frequency response to naphtha-price shocks

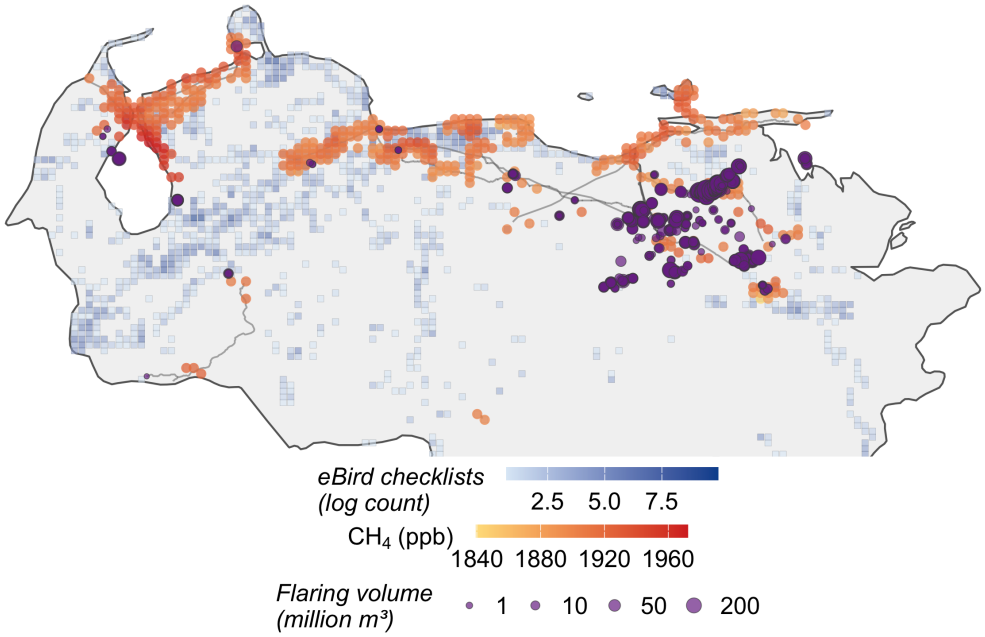
To identify the diluent channel cleanly, we use the spot price of naphtha on the US Gulf Coast (Venezuela's pre-sanctions diluent supplier) as a high-frequency, exogenous proxy for diluent scarcity. If diluent dependence is the channel through which sanctions raise methane, Orinoco pipeline-vicinity methane should respond month-to-month to naphtha-price shocks, and Maracaibo methane should not.

We estimate, on the Venezuelan Orinoco cell-month sample, the contemporaneous-elasticity specification

$$\log \mathrm{XCH}_{4ct} \;=\; \alpha_c + \tau_t + \theta_0 \, \log \mathrm{NaphthaPrice}_t \cdot \mathrm{NearPipe}_c + \varepsilon_{ct}, \tag{5}$$

[11] Online Appendix E documents the L2-versus-L3 choice, the cell-grid construction, our use of the bias-corrected product, and a side-by-side comparison with the atmospheric-inversion methodology of Nathan et al. (2024).

Figure 6: Pipelines, methane, and flare detections in Venezuela



Notes: Black lines: oil and gas pipelines from the combined OSM+GEM network. Yellow-to-red circles: cell-level mean TROPOMI XCH_4 (column-averaged dry-air methane) over 2018–2024, brighter shades indicating higher concentrations. Purple circles: NASA FIRMS VIIRS S-NPP active-fire detections retained as flare sites (types 2 and 3). Source: ESA Sentinel-5P TROPOMI; NASA FIRMS; OpenStreetMap; Global Energy Monitor.

with cell and year-month fixed effects, which gives $\hat{\theta}_0 = +3.05$ ($p = 0.001$): a one-log-point rise in the spot price of naphtha is associated with a 3-log-point rise in Orinoco pipeline-vicinity methane in the same month. A lag-extended version that re-estimates the regression at each of the 1-, 2-, and 3-month lags of log naphtha price (reported in Online Appendix Table E2) yields positive and statistically significant coefficients at every horizon (between $+2.1$ and $+3.1$, all $p < 0.05$), so the contemporaneous elasticity is not driven by single-month timing.

In Maracaibo, where conventional crude does not require naphtha blending, the same interaction is negative at short radii and unstable across distance bands. We interpret this placebo pattern as evidence that the Orinoco response is not a generic hydrocarbon-price effect: conditional on common monthly shocks, higher naphtha prices raise methane only where naphtha is a binding physical input. In conventional Maracaibo operations, higher naphtha prices instead appear to coincide with lower relative pipeline activity, consistent with production being easier to scale down without overloading the gas-handling chain. The full basin-by-distance results are in Table 5.

The naphtha-price test sharpens the trend regression in two ways. First, naphtha prices vary at monthly frequency for reasons orthogonal to Venezuelan oil policy (US Gulf refining, global naphtha demand), so the elasticity is identified off variation that is plausibly exogenous to Venezuela. Second, the prediction is sign-asymmetric across regions in exactly the way the diluent-dependence physics implies: Orinoco responds, Maracaibo does not. Cumulating the elasticity over the post-2018 sanctions window implies a $\approx 0.8\%$ uplift in atmospheric methane near Venezuelan pipelines

Table 5: TROPOMI CH_4: naphtha price × pipeline proximity

	(1) $d = 10$ km	(2) $d = 20$ km	(3) $d = 50$ km	(4) $d = 100$ km
Panel A: Naphtha price × pipeline proximity (all Venezuela)				
log(Naphtha) × NearPipe	0.5427	0.5974	0.2182	-0.4255
	(0.7463)	(0.6140)	(0.5159)	(0.6446)
Observations	25,520	25,520	25,520	25,520
Panel B: Naphtha price × pipeline proximity (Orinoco only)				
log(Naphtha) × NearPipe	2.6774**	3.0533***	2.5254***	0.0002
	(1.1093)	(0.9072)	(0.8361)	(0.9557)
Observations	9,124	9,124	9,124	9,124
Panel C: Naphtha price × pipeline proximity (Maracaibo only)				
log(Naphtha) × NearPipe	-2.5638*	-3.6105**	0.0145	-1.3035
	(1.4016)	(1.5284)	(1.3425)	(1.5904)
Observations	4,218	4,218	4,218	4,218
Cell FE	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓

Notes: This table tests whether global naphtha prices differentially affect methane concentrations near pipeline infrastructure. Sample: Venezuelan 10 km grid cells, 2018–2024, cell-month level. The dependent variable is $\log(\mathrm{XCH_4})$, where $\mathrm{XCH_4}$ is the column-averaged dry-air methane mole fraction (ppb) from TROPOMI; all estimates and standard errors are multiplied by 1,000 for readability. Each column corresponds to a different pipeline proximity threshold d. Panel A estimates the pooled naphtha–proximity interaction across all Venezuelan cells. Panel B restricts the sample to the Orinoco Belt, where extra-heavy crude requires naphtha diluent. Panel C restricts to the Maracaibo Basin, which produces lighter crude that does not require naphtha blending. Naphtha: NAF-C-NWE CIF NW Europe spot ($/tonne, monthly average from Refinitiv). All specifications include cell and year-month fixed effects. Standard errors clustered at the cell level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: ESA Sentinel-5P TROPOMI; Refinitiv; OpenStreetMap.

by 2024 (a year-by-year event-study version of equation (5) yields the same total at calendar-year, semi-annual, and quarterly aggregations). The methane magnitude is reported alongside R in Section 8 as a direction-of-effect confirmation rather than folded into R numerically: the TROPOMI record begins only in November 2018, so there is no clean pre-sanctions methane baseline against which to scale the post-sanctions level, and R is consequently quantified on flaring intensity alone.

6.2 Biodiversity outcomes

To check whether the methane signal we just documented translates into ecological damage at the spatial scale at which households live, we use citizen-science bird observations from eBird, the Cornell Lab of Ornithology platform whose Venezuelan coverage runs to ≈ 134,000 checklists

across 1,508 unique 10 km grid cells over 2012–2024 (Sullivan et al., 2014).[12] For each checklist we record *species richness* (count of unique species) and *log abundance* (log 1 + total individuals).

The natural test runs a horse race between the two pollutants that the conceptual framework predicts should drive ecological damage. Composition (Section 5) operates through field-level flaring, which is observable from VIIRS as fire radiative power. Strain (this section) operates through methane and the volatile organics that co-emit with it, observable from TROPOMI as XCH_4. The two pollutants are imperfectly correlated within Venezuela because they map to different stages of gas-handling failure: VIIRS picks up flare-stack combustion; TROPOMI integrates flare incompleteness, venting, and pipeline leakage. Running both on bird outcomes simultaneously isolates which margin matters at which spatial scale.

We estimate, on Venezuelan eBird checklists 2018–2024,

$$Y_{ict} = \alpha_i + \tau_t + \mu_{b(i),t} + \beta_{\mathrm{CH_4}} \log(1 + \mathrm{CH_4})_{i,R} + \beta_{\mathrm{F}} \log(1 + \mathrm{Flare})_{i,R} + \mathbf{X}'_{ict}\boldsymbol{\delta} + \varepsilon_{ict}, \tag{6}$$

where Y_{ict} is log(1 + abundance) (cols 1–3) or species richness (cols 4–6); $\log(1 + \mathrm{CH_4})$ and $\log(1 + \mathrm{Flare})$ are z-standardised within sample; $\mathrm{CH_4}$ and Flare are aggregated within radius $R \in \{25, 50, 75\}$ km of cell i; α_i is a cell fixed effect, τ_t a year-month fixed effect, and $\mu_{b(i),t}$ a basin-by-year-month fixed effect that absorbs basin-specific seasonal and trend variation. The controls $\mathbf{X}_{ict}$ collect checklist-level effort variables (log duration, log distance, number of observers, traveling indicator, hour and hour^2) together with the cell's proximity to Lake Maracaibo, $-\log(1 +$ distance to Lake Maracaibo$)$, interacted with year and month indicators (a wetland control). Standard errors clustered at the cell level. A basin-interacted version of the same specification reports implied basin-specific slopes (Other / Orinoco / Maracaibo).

Table 6 reports the estimates. The two pollutants act on different spatial scales. VIIRS flaring matters most at $R = 25$ km, the close radius where direct flare-stack disturbance localises: a one-standard-deviation rise in log VIIRS flare is associated with 0.05 log units lower abundance ($p < 0.05$) and roughly three-quarters of a species fewer ($p < 0.01$). At $R = 50$ km the VIIRS signal attenuates and the methane signal takes over: a one-standard-deviation rise in log $\mathrm{CH_4}$ is associated with 0.030 log units lower abundance ($p < 0.01$) and roughly a third of a species fewer ($p < 0.01$); at $R = 75$ km the $\mathrm{CH_4}$ abundance effect attenuates further (-0.021, $p < 0.10$) while the richness effect remains (-0.334, $p < 0.05$). Both effects concentrate in the Orinoco basin (Panel B): Orinoco-specific slopes are uniformly negative across radii, with the $\mathrm{CH_4}$ coefficients on richness and abundance significant at $p < 0.05$ at every R and the VIIRS coefficients significant at the close $R = 25$ radius

[12] Online Appendix H (Table H1) reports sample statistics: mean richness 13.7 species per checklist, median 11; mean log abundance 3.4. Checklist density is highest in the populated north and along the pipeline corridors mapped in Figure 6, providing the spatial overlap with TROPOMI and VIIRS that the regression below exploits. A natural concern with citizen-science data is that observer effort might covary with pollution exposure: if birders avoid (or seek) methane- or flare-heavy cells, the estimated pollution effect on biodiversity would be a selection artefact. We rule this out directly: regressing six standardised effort and coverage outcomes on $\log(1 + \mathrm{CH_4})$ and $\log(1 +$ VIIRS flare$)$ at the $R = 50$ km buffer used below, with cell and year-month fixed effects, all twelve coefficients are individually insignificant (largest $|t| = 1.22$). Pollution exposure does not predict observer effort, so the richness and abundance results below are not selection-driven (Online Appendix Table H5).

Table 6: Atmospheric Pollution, Bird Abundance, and Species Richness

	log(1 + abundance)			Species richness		
	$R = 25$ km	$R = 50$ km	$R = 75$ km	$R = 25$ km	$R = 50$ km	$R = 75$ km
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Pooled specification						
$\log(1 + CH_4)$ (std)	-0.026**	-0.030***	-0.021*	-0.109	-0.329***	-0.334**
	(0.010)	(0.011)	(0.011)	(0.123)	(0.109)	(0.129)
$\log(1 +$ VIIRS flare) (std)	-0.051**	0.001	0.005	-0.748***	-0.050	-0.249
	(0.023)	(0.019)	(0.022)	(0.211)	(0.194)	(0.259)
Panel B: Implied basin-specific slopes						
Other region						
CH_4 (std)	0.027	-0.021	0.001	0.436	0.149	0.335
	(0.027)	(0.025)	(0.025)	(0.326)	(0.301)	(0.284)
VIIRS flare (std)	-0.022	0.028	0.013	-0.749**	0.104	-0.105
	(0.040)	(0.032)	(0.043)	(0.371)	(0.355)	(0.400)
Orinoco basin						
CH_4 (std)	-0.038***	-0.040***	-0.031**	-0.273**	-0.553***	-0.545***
	(0.013)	(0.013)	(0.016)	(0.114)	(0.147)	(0.184)
VIIRS flare (std)	-0.050*	-0.031	-0.009	-0.651**	-0.415	-0.432
	(0.030)	(0.026)	(0.025)	(0.282)	(0.281)	(0.323)
Maracaibo basin						
CH_4 (std)	-0.028	0.002	0.008	0.124	0.088	-0.010
	(0.042)	(0.030)	(0.046)	(0.631)	(0.231)	(0.375)
VIIRS flare (std)	-0.104**	0.102	0.126	-1.029	1.614*	1.773*
	(0.046)	(0.072)	(0.087)	(0.710)	(0.925)	(0.981)
Effort controls	✓	✓	✓	✓	✓	✓
Cell FE	✓	✓	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓	✓	✓
Basin × month FE	✓	✓	✓	✓	✓	✓
Lake prox. × year, month	✓	✓	✓	✓	✓	✓
Observations	49,851	80,094	93,884	49,851	80,094	93,884
Mean dep. var.	3.54	3.49	3.47	14.92	14.44	14.34

Notes: Joint test of pollution exposure on bird outcomes in the post-sanctions period. Outcomes are $\log(1 + \text{abundance})$ in cols (1)–(3) and species richness (number of unique species per checklist) in cols (4)–(6). Sample: Venezuelan eBird checklists 2018–2024 with non-missing biodiversity outcomes. Both exposure variables are standardised (z-scored) within each radius-specific regression sample. CH_4 is the mean TROPOMI column-averaged dry-air methane mixing ratio (ppb) across 10 km cells within radius R; VIIRS flare is the sum of fire radiative power (MW) from NASA FIRMS active-fire detections within radius R. Panel A reports the pooled specification. Panel B reports implied basin-specific slopes from an interaction specification where each exposure is interacted with basin indicators (Other = baseline; Orinoco and Maracaibo = baseline + interaction). All regressions include cell, year-month, and basin×month fixed effects, log-distance to Lake Maracaibo interacted with year and month dummies, and checklist-level effort controls (log duration, log distance, number of observers, traveling indicator, hour-of-day quadratic). Standard errors clustered at the cell level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: eBird; Sentinel-5P TROPOMI; NASA FIRMS VIIRS.

where flare-stack disturbance localises. The horse race therefore reads as complementary rather than substitute: the composition channel of Section 5 appears in the close-radius VIIRS effect, and the strain channel of this section appears in the larger-radius CH_4 effect, and both bite hardest at the heavy operations in Orinoco where the mechanisms come together.

Bird-relocation test. A second concern is whether the pollution-biodiversity association reflects birds physically relocating away from polluted cells rather than dying or failing to recruit. To distinguish the two we estimate a triple-difference of cell-year mean richness on POST $\times$ VEN $\times$ FAR using Colombia as the placebo country.[13] Under the relocation hypothesis the triple should be *positive and large*: displaced birds in Venezuela should crowd into far-from-oil cells, raising richness there relative to the corresponding far-from-oil cells in Colombia. We find $\hat{\beta}_3 = +0.39$ (SE $= 2.14$, $p = 0.85$), statistically indistinguishable from zero. The data therefore reject pure spatial relocation: there is no measurable net flow of birds from polluted to unpolluted Venezuelan cells over the post-sanctions period. The biodiversity reduction we estimate is consistent with population loss, not with birds moving next door.[14]

Strain at large fields: a complementary OOIP gradient. A third probe asks whether bird abundance is depressed near the largest surviving fields, where post-sanctions production concentrates and gas-handling runs closest to capacity. Regressing checklist-level log abundance on $\log(\text{OOIP})$ within 50 km, with observer fixed effects and the same effort controls as Table 6, the Orinoco-OOIP interaction is -0.054 ($p < 0.01$): a one-standard-deviation increase in $\log(\text{OOIP})$ corresponds to roughly 5% fewer birds. The Maracaibo gradient is similarly negative (-0.033, $p < 0.05$) but if anything attenuates over time as production at that basin collapses. Within-post-period trend tests are imprecise (pre-2020 eBird coverage is thin), so we treat this gradient as suggestive support for the strain channel rather than identifying evidence; full estimates appear in Online Appendix Table H2.

[13] The estimating equation is

$$\overline{\text{richness}}_{cyt} = \alpha_c + \tau_y + \beta_1 (\text{POST}_y \cdot \text{VEN}_c) + \beta_2 (\text{POST}_y \cdot \text{FAR}_c) + \beta_3 (\text{POST}_y \cdot \text{VEN}_c \cdot \text{FAR}_c) + \varepsilon_{cy},$$

with $\overline{\text{richness}}_{cyt}$ = mean checklist species richness in cell c, year y; cell (α_c) and year (τ_y) fixed effects; $\text{POST}_y = \mathbb{1}\{y \geq 2017\}$; $\text{VEN}_c = \mathbb{1}\{\text{cell in Venezuela}\}$ (with Colombian cells as the placebo control); $\text{FAR}_c = \mathbb{1}\{\text{cell more than 10 km from any Rystad oil field}\}$. The coefficient of interest is the triple β_3. Standard errors clustered at the cell level. Under the relocation hypothesis $\beta_3 > 0$: displaced Venezuelan birds crowd into Venezuelan FAR cells, raising richness there relative to Colombian FAR cells.

[14] The triple-difference panel is restricted to Venezuela and Colombia, so it does not directly test long-haul overwater emigration to the Caribbean islands, Trinidad and Tobago, Guyana, or Brazil. We do not regard this as a serious alternative for our setting because the guilds that drive the Chevron-licence abundance gain (Section 7.2) are forest insectivores, granivores, and frugivores, mostly sedentary tropical residents whose flight stamina and behavioural repertoire do not include sustained overwater crossings. The one guild that does cross open water seasonally, waterbirds, shows small and statistically insignificant Chevron-licence coefficients. A complementary diagnostic using vegetation and land-cover controls (Hansen tree-cover loss; GHSL built-up indicators), reported in Online Appendix H, shows that the pollution-biodiversity association is not driven by differential habitat conversion. Together these diagnostics rule out the leading non-pollution explanations: local relocation, long-haul overwater emigration for the relevant species, and habitat conversion.

7 The maintenance channel: the Chevron quasi-experiment

The maintenance channel predicts that sanctions raise per-barrel emissions at fields whose total production is unchanged or modestly changed, because operating expenditure on leak prevention, gas reinjection, and equipment turnover falls. We exploit one sharply dated relaxation of the maintenance constraint to identify this channel: OFAC General Licence 41, issued on 26 November 2022, which authorised Chevron specifically to resume operations and maintenance at four Venezuelan joint-ventures (PetroBoscan, PetroIndependencia, PetroIndependiente, and PetroPiar).[15] The licence reset Chevron's per-barrel OPEX to roughly its pre-sanctions level while sanctions remained in place at all other Venezuelan fields, generating quasi-experimental variation in maintenance spending that is not also variation in production volume.

We use *per-barrel* OPEX, $\log(1 + \text{OPEX}_{ft}/\text{Production}_{ft})$, as the maintenance regressor, scaled this way for two reasons. First, dollar OPEX correlates mechanically with field size; scaling by production isolates the per-barrel maintenance *intensity* that the channel is supposed to operate through. Second, because flaring and bird outcomes are also per-barrel objects, the elasticity η on $\log(\text{OPEX}/\text{prod})$ has a direct maintenance-channel interpretation. We saturate the specification with indicators for the two policy shocks, $\text{Post 2017}_t = \mathbb{1}\{t \geq 2017\}$ (broad sanctions) and $\text{Post 2022}_t = \mathbb{1}\{t \geq 2023\}$ (Chevron licence; we date Post 2022 to 2023 because the licence took effect in late November of that year), and a Chevron joint-venture indicator:

$$\begin{aligned} Y_{it} \;&= \alpha_i + \tau_t + \eta_0\, X_{it} + \eta_1\, X_{it} \cdot \text{Post 2017}_t \\ &\quad + \eta_2\, X_{it} \cdot \text{Post 2022}_t + \eta_3\, X_{it} \cdot \text{Chevron}_i + \eta_4\, \text{Post 2022}_t \cdot \text{Chevron}_i \\ &\quad + \eta_5\, X_{it} \cdot \text{Post 2022}_t \cdot \text{Chevron}_i + \varepsilon_{it}, \end{aligned} \tag{7}$$

where $X = \log(1 + \text{OPEX}/\text{prod})$, α_i is a field fixed effect (column 1) or cell fixed effect (columns 2–4), τ_t is a year (column 1) or year-month (columns 2–4) fixed effect, and Chevron $= 1$ at the four covered fields (column 1) or at any cell within 30 km of one (columns 2–4). The maintenance coefficient is η_0 (the average within-field elasticity over 2012–2024); η_5 is the Chevron-licence-specific differential after November 2022.

7.1 OPEX intensity, flaring, and bird abundance

Table 7 reports the saturated specification on flaring (column 1, field-year level) and on bird abundance (columns 2–4, cell-checklist level, three increasingly demanding fixed-effect structures). Three results stand out.

[15] The body specification identifies η_5 off three of the four GL-41 joint-ventures, all operating heavy crude with vulnerable gas-handling chains: PetroPiar (VEF0028, Hamaca, Orinoco extra-heavy), PetroIndependencia (represented by VEF0466, Carabobo 5 Early, Orinoco extra-heavy), and PetroBoscan (VEF0082, Boscan, Maracaibo heavy with API around 10°). The fourth GL-41 partner, PetroIndependiente, appears in Rystad as the conventional medium-API Lake Maracaibo asset LL-652, classified as offshore; the diluent-strain mechanism the licence is supposed to relax does not operate at conventional fields (Section 6). Pooling PetroIndependiente into the treatment group nonetheless delivers a similar but smaller elasticity, as expected: $\eta_5 = -0.30$ instead of -0.42 (both $p < 0.001$).

Table 7: Maintenance Channel: Per-Barrel OPEX, Flaring, and Bird Abundance

	Flaring	Bird abundance		
	log(1 + vol.)	log(1 + abund.)		
	(1)	(2)	(3)	(4)
log(1 + OPEX/prod) (η_0)	-0.0462*	-0.1588	-0.4389	+0.1279
	(0.0239)	(0.3731)	(0.3657)	(0.2286)
× Post2017 (η_1)	+0.0577**	+0.2261	+0.5052	-0.0665
	(0.0255)	(0.3688)	(0.3629)	(0.2212)
× Post2022 (η_2)	-0.0202	-0.0675	-0.0218	-0.1471
	(0.0154)	(0.0471)	(0.0477)	(0.1224)
× Chevron (η_3)	-0.8184	-2.2792***	-2.1356***	-1.4534***
	(2.2590)	(0.3506)	(0.3680)	(0.3180)
× Post2022 × Chevron (η_5)	-0.4170***	+1.6407**	+1.4437**	+1.6382*
	(0.0869)	(0.7845)	(0.5947)	(0.8477)
Post2022 × Chevron (eta_4)	+1.5960***	-2.3427	-1.8200	-2.7721
	(0.2696)	(1.9741)	(1.4607)	(2.1272)
Field FE	✓	—	—	—
Cell FE	—	✓	✓	✓
Year (FE for col 1) / Year-Month (cols 2–4)	✓	✓	✓	✓
Observer FE	—	No	✓	—
Observer × Year FE	—	No	No	✓
Observations	2,344	7,565	7,522	7,457

Notes: Direct test of the maintenance channel. The dependent variable is log(1 + flaring volume in million m^3) at the field-year level in column (1), and log(1 + bird abundance per checklist) at the cell-checklist level in columns (2)–(4). The regressor is per-barrel operating expenditure, $\log(1 + \text{OPEX}_{ft}/\text{Production}_{ft})$ in USD bbl^{-1}, interacted with two policy-period indicators ($\text{Post2017}_t = \mathbb{1}\{t \geq 2017\}$, $\text{Post2022}_t = \mathbb{1}\{t \geq 2023\}$) and a Chevron joint-venture indicator. In column (1), Chevron = 1 if the field is one of the four OFAC GL-41 Chevron JVs (PetroBoscan, PetroIndependencia, PetroIndependiente, PetroPiar). In columns (2)–(4), Chevron = 1 if any such field lies within 30 km of the cell, and OPEX/production is the cell-year mean of log(1 + OPEX/prod) over fields within 30 km of the cell. Sample in (1): Venezuelan onshore field-years with positive production, 2012–2024. Sample in (2)–(4): Venezuelan eBird checklists in cells with at least one Venezuelan onshore field within 30 km, 2012–2024. Columns (2)–(4) include checklist-level effort controls (log duration, log distance travelled, number of observers, hour started, traveling indicator). Standard errors clustered at the field level (col 1) or cell level (cols 2–4). ***p<0.01, **p<0.05, *p<0.1. Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker; eBird.

First, the within-field maintenance elasticity on flaring volume ($\eta_0 = -0.046^*$, column 1) is negative and marginally significant: a one-log-point increase in per-barrel OPEX at the average within-Venezuela field is associated with a 4–5% ($1 - \exp(\eta_0)$) reduction in flaring volume at the same field. The Chevron-licence-specific differential $\eta_5 = -0.417^{***}$ is much larger and highly significant: at the covered fields, a one-log-point rise in per-barrel OPEX in the post-licence period reduces flaring volume by an additional $1 - \exp(\eta_5) \approx 34\%$, and a doubling of OPEX ($\Delta \log = \log 2$) cuts flaring by approximately 25%. The licence sharpens the maintenance-channel response at the very fields that received the maintenance-spending shock.

Second, the bird-abundance counterpart η_5 has the same sign, flipped because more maintenance restores ecosystem outcomes: +1.64 (Cell+YM), +1.44 (Cell+YM+Obs), +1.64 (Cell+YM+Obs×Year). The point estimate is stable across all three fixed-effect specifications,

including observer×year fixed effects which absorb within-observer effort drift, and remains statistically significant ($p \leq 0.06$) under each. Read together with the flaring column, these results trace one mechanism: the licence restored Chevron's per-barrel OPEX, which (a) reduced flaring at the covered fields and (b) raised bird abundance in the cells exposed to those fields.

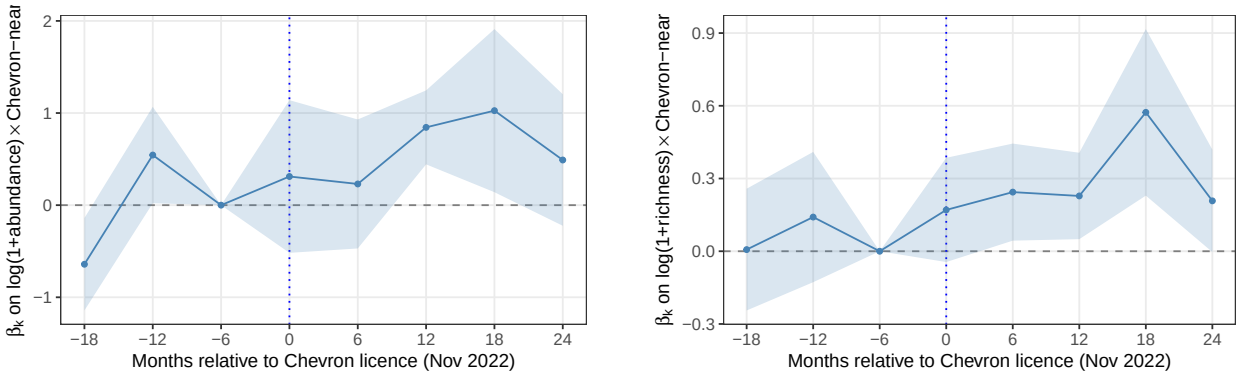
Third, the implied bird-abundance elasticity at Chevron-affected cells is large in absolute terms, roughly an order of magnitude above the US estimates of Liang et al. (2020) for the NO_x Budget Trading Program and Liang et al. (2025) for economic activity. Three features of our setting account for the gap: baseline pollution at sanctioned Venezuelan operations is much higher than at US sources subject to EPA enforcement, multiple co-pollutants (flaring particulates, methane leakage, vented BTEX, diluent residues) move together with maintenance restoration rather than in isolation, and tropical forest insectivores, granivores, and waterbirds are more sensitive to acute disturbance than the US ozone-exposed species in those papers. The order-of-magnitude comparison is consistent with the maintenance channel operating through several pollutants jointly at degraded baseline.

Dynamics: abundance and species richness around the licence. Figure 7 reports the dynamic counterpart of Table 7: per-bin coefficients on NearChevron × bin, in 6-month windows ±24 months around the licence date, for log abundance (panel (a)) and log richness (panel (b)). The reference bin is the six months immediately pre-licence ($k = -6$ m), so each $\hat{\beta}_k$ traces the differential in Y at Chevron-near cells relative to that pre-period baseline and relative to other Venezuelan cells in the same year-month.

The two panels share the same regression structure and identifying variation, but differ in magnitude and timing. *Abundance* (panel (a)) is roughly balanced pre-licence (-0.64 at -18 m and $+0.54$ at -12 m, both wide CIs that span zero), then builds gradually post-licence: $+0.84$ at $+12$ m ($p < 0.001$), peaking at $+1.03$ at $+18$ m ($p = 0.025$), and settling at $+0.49$ by $+24$ m. *Richness* (panel (b)) is roughly half the magnitude at every post-licence horizon, peaking at $+0.57$ at $+18$ m ($p < 0.01$). Pre-period coefficients in panel (b) are near zero and statistically insignificant, which is reassuring for parallel-trends.

Three takeaways. *First*, the lag profile in panel (a) is consistent with the mechanics of clearing a maintenance backlog: parts ordered, equipment swapped, leaks sealed, gas-reinjection capacity restored over months, with ecological response trailing several quarters further. *Second*, the abundance–richness wedge — abundance responds roughly twice as strongly as richness — aligns with how bird communities recover: individuals can return immediately when habitat quality improves, but new *species* require time to disperse in, establish breeding populations, and persist long enough to show up on checklists. *Third*, the saturated specification in Table 7 averages over this dynamic; the event study shows that the static coefficient understates the effect at the peak horizon and overstates it at shorter horizons. Counterpart richness regressions in pooled form are reported in Online Appendix Table H3.

Figure 7: Bird Abundance and Species Richness at Chevron Cells around the Licence Date



(a) $\log(1 + \text{abundance})$ (b) $\log(1 + \text{richness})$

Notes: Per-bin coefficients β_k from $Y_{ict} = \alpha_i + \tau_t + \sum_{k \neq -6\text{m}} \beta_k \mathbb{1}\{\text{bin}_t = k\} \cdot \text{NearChev}_i + \mathbf{X}'\delta + \varepsilon$, with $Y = \log(1 + \text{abundance})$ in panel (a) and $Y = \log(1 + \text{richness})$ in panel (b). Estimated on Venezuelan eBird checklists in cells with at least one field within 30 km, ±24 months around 26 November 2022, with 6-month bins. Cell, year-month, and observer fixed effects; effort controls. Reference bin: −6 to 0 months pre-licence. Vertical blue line at the licence date. Shaded band: 95% confidence interval. Source: eBird; Rystad Energy UCube; World Bank VIIRS.

7.2 Guild-level decomposition

To probe which trophic guilds drive the abundance recovery, we estimate a static Chevron-vs-Western-JV difference-in-differences on each guild's log abundance,

$$\log(1 + Y_{g,ict}) \;=\; \alpha_i + \tau_t + \beta_g \left(\text{NearChev}_i \cdot \text{Post}_t\right) + \mathbf{X}'_{ict}\delta + \varepsilon_{ict}, \tag{8}$$

on the ±24-month window around the licence date, restricted to cells with at least one Chevron or Western JV field within 30 km. Treatment cells are those near a Chevron joint-venture; controls are cells near other Western JVs (which lacked the licence). Effort controls $\mathbf{X}_{ict}$ are checklist-level (duration, distance, observers, hour, traveling).

Online Appendix Table H4 reports estimates across three fixed-effect structures. Forest insectivores show the clearest post-licence gain ($+0.60^{***}$ under Cell+YM, $+0.51^{***}$ under Cell+YM+Obs), consistent with their canopy and understory foraging niches being sensitive to local air-quality and gas-flaring disturbance. Frugivores and granivores show smaller, directionally positive coefficients but never reach conventional significance. Waterbirds show small negative coefficients in the first two specifications and zero under the most demanding observer×year FE; their pre-trends are also the cleanest of the four guilds, so the licence shock plausibly does not move their abundance much. The pattern is consistent with the eBird-based evidence in the existing literature (Liang et al., 2020): pollution remediation moves abundance most for guilds whose foraging niches overlap most directly with the local pollutant.[16]

[16] A species-level decomposition addresses the concern that the guild-level effects might be driven by a small number of species. Of 265 species with at least five checklist detections in the sample, 172 have positive (Chevron × Post) coefficients and 92 have negative; the binomial sign test rejects equal shares at $p < 10^{-6}$, and the sum of point estimates is +4.53 species, in line with the net guild-level gain. Individual-species significance is low (17% of positive and 14% of

8 Quantification, generalisation, and policy implications

The framework in Section 2 predicts that sanctions-induced environmental damage exceeds what aggregate-output reduction alone would imply, with the gap given by the sum of three channel-specific contributions. This section quantifies that gap for the Venezuelan case from 2017 to 2024, extends the qualitative pattern to Russia after 2022 and Iran after 2018, and draws design implications for sanctions architecture.

8.1 Constructing the counterfactual

Let D_{pre} denote total environmental damage in the pre-sanctions baseline year (2016), and D_{post} the damage observed in the post-sanctions period (we use the 2022–2024 average to smooth short-run volatility). Aggregate output fell from $Q_{\text{pre}} \approx 2.28$ mbbl/d to $Q_{\text{post}} \approx 0.78$ mbbl/d over this window, a reduction of $\rho \equiv 1 - Q_{\text{post}}/Q_{\text{pre}} \approx 66\%$.

Under the proportional-decline benchmark of Section 2,

$$D^{\text{ctf}}_{\text{post}} \;\equiv\; (1-\rho)\, D_{\text{pre}} \;=\; \tfrac{Q_{\text{post}}}{Q_{\text{pre}}}\, D_{\text{pre}}, \tag{9}$$

holding per-barrel intensity, the set of operating fields, and proximity to capacity all fixed at their pre-sanctions values. The empirical question is whether $D_{\text{post}} > D^{\text{ctf}}_{\text{post}}$, and by how much.

We construct the ratio

$$R \;\equiv\; \frac{D_{\text{post}}}{D^{\text{ctf}}_{\text{post}}} \;=\; \frac{D_{\text{post}}}{D_{\text{pre}}} \cdot \frac{Q_{\text{pre}}}{Q_{\text{post}}} \tag{10}$$

where $R = 1$ means proportional decline holds and $R > 1$ means damage outruns the output reduction.

8.2 Channel-by-channel decomposition

The outcome we quantify is gas-flaring intensity — the total m^3 flared per barrel of crude lifted. Production-weighted mean intensity in Venezuela is $\bar{\alpha}_{\text{pre}} = 12.6$ m^3/bbl in the pre-period (2012–2016) and $\bar{\alpha}_{\text{post}} = 38.0$ m^3/bbl in the post-period (2022–2024), so directly from the data

$$R \;=\; \frac{\bar{\alpha}_{\text{post}}}{\bar{\alpha}_{\text{pre}}} \;=\; \frac{38.0}{12.6} \;=\; 3.0. \tag{11}$$

Equivalently, $R = F_{\text{post}}/F^{\text{ctf}}_{\text{post}}$, where $F^{\text{ctf}}_{\text{post}} \equiv F_{\text{pre}}\,(Q_{\text{post}}/Q_{\text{pre}})$ is the proportional-decline counterfactual on total flaring volume; simple cancellations turn the volume ratio in equation (10) into the intensity ratio in equation (11). Empirically, total Venezuelan flaring volume is roughly constant between the pre- and post-periods while production fell by 66% — which is the same fact as a

negative species at the 5% level), reflecting limited power at the species level given the 2,600-checklist sample, so the guild-level result is best read as a direction-of-effect finding.

tripling of per-barrel intensity. We work with the per-barrel ratio, as defined in equation (10), because it isolates the degradation of the production *process* from the output decline from sanctions.

Venezuelan flaring damage in 2022–2024 is three times what the proportional-decline benchmark would predict. We now decompose the change in flaring $\frac{\bar{\alpha}_{\text{post}}-\bar{\alpha}_{\text{pre}}}{\bar{\alpha}_{\text{pre}}} = R - 1$ into the three channels of Section 2, with each component anchored to a specific coefficient already estimated in Sections 5–7. Methane (Section 6) and biodiversity (Sections 6.2, 7) confirm the channel-specific signs but are not folded into R numerically since we lack their pre-sanction values; we report them as direction-of-effect confirmations and quantify R on flaring intensity alone.[17]

The standard Oaxaca decomposition of the change in production-weighted mean intensity, with H for heavy fields (Orinoco East or heavy-crude share above 50%) and NH for non-heavy, is

$$\underbrace{\bar{\alpha}_{\text{post}} - \bar{\alpha}_{\text{pre}}}_{\text{numerator of } R-1} = \underbrace{\Delta s_H \left(\bar{\alpha}_{H,\text{pre}} - \bar{\alpha}_{NH,\text{pre}}\right)}_{\text{composition}} + \underbrace{s_{H,\text{post}}\, \Delta\bar{\alpha}_H + \left(1 - s_{H,\text{post}}\right) \Delta\bar{\alpha}_{NH}}_{\text{within-field intensity change}}, \tag{12}$$

where s_H is the heavy share of Venezuelan production. The four production-weighted intensities and the heavy share are reported in Table 8.

Table 8: Inputs to the R decomposition.

	Pre (2012–2016)	Post (2022–2024)	Δ	Source
Panel A: Production-weighted aggregates from raw data				
Heavy intensity $\bar{\alpha}_H$ (m^3/bbl)	12.9	39.6	+26.7	—
Non-heavy intensity $\bar{\alpha}_{NH}$ (m^3/bbl)	8.1	6.7	-1.4	—
Overall intensity $\bar{\alpha}$ (m^3/bbl)	12.6	38.0	+25.4	—
Heavy share of production s_H	0.94	0.95	+0.01	—
Panel B: Structural coefficients from earlier regressions				
	Value			Source
Chevron OPEX–flaring elasticity η_5	−0.36			Table 7
Heavy-VEN log-OPEX shock β_h^{OPEX}	−0.61			Table 3

Notes: Panel A reports production-weighted means computed as $\sum_i F_{it} / \sum_i Q_{it}$ over Venezuelan field-years with positive flaring and positive production. Heavy = Orinoco East geographic region or heavy-crude share above 50%. Sample windows: 2012–2016 (pre) and 2022–2024 (post). F in cubic metres, Q in barrels of crude. Source: VIIRS Global Gas Flaring Tracker (allocated within 15 km of each field); Rystad Energy UCube. Panel B reports the two structural coefficients used to split the within-heavy intensity rise into maintenance and strain contributions: η_5 is the Post-2022 × Chevron × $\log(1 + \text{OPEX/prod})$ triple interaction on $\log(1 + \text{flaring volume})$ (column 1 of Table 7); β_h^{OPEX} is the Heavy × Post coefficient on $\log(1 + \text{OPEX})$ (Table 3).

Two implications follow immediately:

Composition. Most Venezuelan production was already at heavy fields before sanctions (94% of barrels), so the within-VEN heavy-share shift is small: $\Delta s_H = +0.01$. The within-Venezuela survival result, that small fields exit and large fields keep producing, is identified by the Post × Large

[17] Ideally R would combine flaring CO_2 with vented CH_4 on a common global warming potential scale, but the TROPOMI methane record begins only in November 2018.

coefficient in Table 2 (+0.15 to +0.24, all $p < 0.001$), but because the surviving large fields are essentially all heavy already, this within-VEN selection does not move the heavy share much. The resulting contribution to $R - 1$ is

$$r_{\text{comp}} = \frac{\Delta s_H \left(\bar{\alpha}_{H,\text{pre}} - \bar{\alpha}_{NH,\text{pre}}\right)}{\bar{\alpha}_{\text{pre}}} = \frac{0.01 \cdot (12.9 - 8.1)}{12.6} \approx 0.004,$$

essentially zero. The composition channel does very little numerical work in this case because there is no clean-versus-dirty margin to shift along inside Venezuela: the country's pre-sanctions production base was already concentrated on heavy crude.

Maintenance and strain (joint). The remaining $R - 1 \approx 2.0$ is the within-heavy intensity rise. Empirically, $\bar{\alpha}_H$ tripled from 12.9 to 39.6 m^3/bbl ($\Delta\bar{\alpha}_H \approx +26.7$ m^3/bbl). This is consistent with the differential Heavy × Post coefficient on flaring intensity reported in Table 3: $\hat{\beta}_h = +125.1^{***}$ m^3/bbl at the 15 km radius; the production-weighted analogue recovers the same direction at smaller magnitude because most heavy production sits at large fields with lower per-barrel intensity. Weighted by post-period heavy share, the within-heavy intensity rise contributes

$$r_{\text{maint+strain}} = \frac{s_{H,\text{post}}\,\Delta\bar{\alpha}_H + (1 - s_{H,\text{post}})\,\Delta\bar{\alpha}_{NH}}{\bar{\alpha}_{\text{pre}}} = \frac{0.95 \cdot 26.7 + 0.05 \cdot (-1.4)}{12.6} \approx 2.01,$$

which matches $R - 1 = 2.01$ to two decimal places. The within-heavy intensity rise is the entirety of the damage multiplier; non-heavy intensity is essentially flat, consistent with the regression coefficient $\beta_n \approx 0$ on flaring intensity in column (4) of Table 3.

Splitting maintenance from strain. The heavy-field intensity rise bundles two sub-channels in equilibrium because sanctions cut both maintenance budgets and diluent imports simultaneously at the same fields. The coefficients from Sections 6.1 and 7 let us identify each sub-channel.

The within-Chevron OPEX elasticity of flaring volume from Table 7 is $\eta_5 = -0.42^{***}$ (column 1): a one-log-point rise in OPEX per barrel reduces flaring volume by about 42%. The Heavy × Post coefficient on log(OPEX) in Table 3 (column 6) is -0.61^{***}, a 46% post-period drop in heavy-field OPEX. Multiplying gives the maintenance-induced rise in heavy flaring volume: $\eta_5 \cdot \beta_h^{\text{OPEX}} = (-0.42) \cdot (-0.61) = +0.2562$ log units, or about $\exp(0.2562) - 1 \approx +29\%$ on flaring volume from maintenance alone. This calculation assumes the Chevron-licence elasticity η_5, identified at the four GL-41 fields, is the relevant maintenance-elasticity for all heavy Venezuelan fields when maintenance is suppressed by sanctions.

The strain component is then obtained as the residual: the within-heavy rise that remains after the maintenance contribution has been subtracted. Combining, maintenance accounts for roughly 29% of the heavy-field flaring rise and strain for the remaining 71%. Both channels are sanctions-specific: the cross-country comparison (Online Appendix Table D2) shows Mexican Maya and Colombian Llanos wellhead-pipeline trends *fall* at −2.51 and −0.81 per year respectively; only

the Venezuelan Orinoco shows the trend rise.

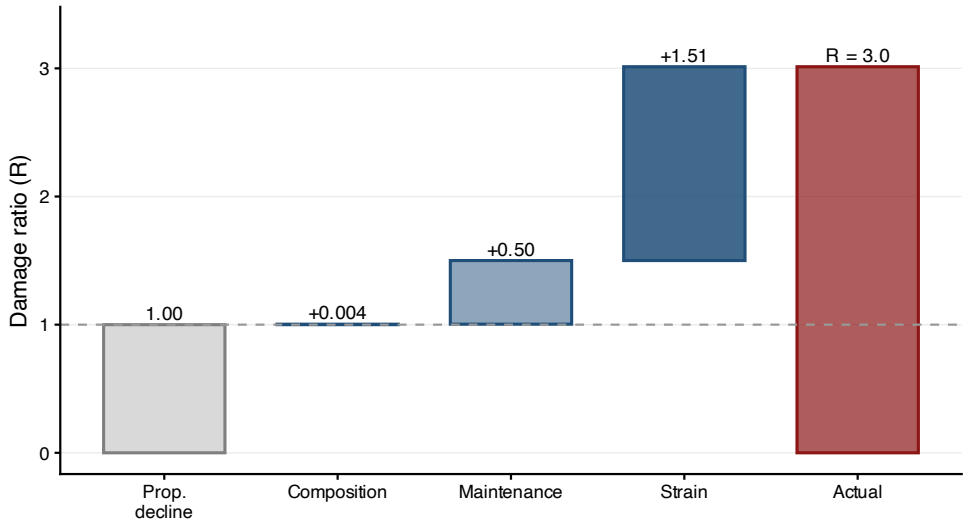
The damage multiplier. Equation (12) therefore gives

$$R \;=\; 1 \;+\; \underbrace{0.004}_{r_{\text{comp}}} \;+\; \underbrace{0.59}_{r_{\text{maint}}} \;+\; \underbrace{1.42}_{r_{\text{strain}}} \;\approx\; 3.0, \tag{13}$$

with sensitivity bounds across the 7.5 km versus 15 km flare-allocation radii (which produce a Heavy $\times$ Post coefficient of $+176.7$ versus $+125.1$, Table 3 cols (3)–(4)) and across the post-period definition (2022 alone vs. 2022–2024 average) giving $R \in [2.5, 3.5]$. Composition does very little work; maintenance and strain together do almost everything. Figure 8 visualises the decomposition.

Population-weighted multiplier. Weighting each field's damage contribution by the population within a 30 km buffer (GHS-POP 2020) raises the multiplier from a field-equal-weighted analogue of $R = 2.72$ to $R^{\text{pop}} = 2.93$, with $R^{\text{pop}} \in [2.86, 3.25]$ across 10–50 km buffers. The surviving heavy operations sit near populated cities at the northern margin of the Orinoco belt and around Lake Maracaibo, so the exposure-weighted welfare loss exceeds the field-equal aggregate.

Figure 8: Decomposition of the damage ratio R into three sanctions-driven channels



Notes: The grey bar marks the counterfactual baseline $R = 1$ that would prevail if per-barrel pollution had simply scaled with the proportional output decline. The three blue bars are the additive contributions to R from equation (13): composition ($+0.004$, the share-weighted reallocation toward heavy fields), maintenance ($+0.59$, from the OPEX-restoration elasticity $\eta_5 = -0.42$ and the 0.61 log-unit drop in per-barrel OPEX at heavy fields), and strain ($+1.42$, the residual within-heavy intensity rise after subtracting maintenance, anchored by the naphtha-price elasticity $\theta_0 = +3.05$). The red bar at the right is the empirically observed ratio $R \approx 3$. Composition is environmentally negligible because 94% of pre-sanctions production was already heavy; the post-sanctions environmental cost is built almost entirely from maintenance and strain.

8.3 Cross-country generalisation

The same three-channel structure should hold under any large sanctions-driven reallocation of oil-and-gas production. Russia post-2014 and Iran post-2018 are the natural external-validity checks. We do not run our own field-level decomposition for either case; the multiple overlapping sanctions waves in Russia and the opacity of Iranian post-2018 production data both require treatment that goes beyond the scope of this paper. External evidence on production composition and methane emissions is consistent with the qualitative pattern we document for Venezuela.

Russia. Russia has been subject to three successive sanctions waves: 2014 (Crimea/Donbass financing and technology restrictions), 2017–2018 (CAATSA), and the most severe in February 2022 (G7+EU oil price cap, EU embargo on seaborne crude, and broadened US restrictions on Russian financial institutions and oil-trade intermediaries). Aggregate Russian crude production fell modestly, from 10.7 to 9.5 million barrels per day between 2021 and 2023 (-11%). The composition signature parallels Venezuela: surviving output has concentrated on the Western Siberian supergiants of the Yamal Peninsula, where state-owned Gazprom and Rosneft retain the technical capacity that European-financed mid-tier fields could not sustain after the 2022 import controls. Russia is already the world's largest oil-and-gas methane emitter at roughly 10–14 Mt per year (International Energy Agency, 2024), and post-2022 evidence points to deterioration: documented flaring spikes accompanied the suspension of Nord Stream exports in mid-2022 (notably at the Portovaya loading station near the Finnish border), and sanctions on imported methane-detection sensors, drones, and gas-handling equipment have constrained the routine repair work that contained Russian leaks pre-invasion (Center on Global Energy Policy, Columbia SIPA, 2023).

Iran. The May 2018 US withdrawal from the Joint Comprehensive Plan of Action and the subsequent maximum-pressure campaign reduced Iranian production from 3.8 to 1.9 million barrels per day at the 2020 trough; partial recovery to 3.2 by 2023–24 has been routed through dark-shipping networks (Fernández-Villaverde et al., 2025). The composition signature parallels Venezuela: surviving production has concentrated at the South Pars condensate field (the world's largest gas-condensate field, shared with Qatar) and at the heavy-oil clusters of Ahvaz and Marun in Khuzestan province. Iranian oil-and-gas methane emissions rose 10% to 15% between 2018 and 2024 even as production fell (International Energy Agency, 2024).

Cross-country contrast: sanctioned vs. operating heavy-oil regions. A natural concern is that what we attribute to sanctions might be a generic feature of heavy-oil extraction. To separate the two, we estimate the same pipeline-vicinity methane trend in unsanctioned heavy-oil regions of Latin America and compare to the Venezuelan estimate. The differential per-year trend in pipeline methane (relative to the Latin American mean) is $+1.30^{**}$ in the Venezuelan Orinoco Belt; -2.51^{***} in the Mexican Maya basin; $+1.52^{***}$ in the Ecuadorean Oriente; and -0.81^{***} in the Colombian Llanos (Online Appendix Table D2 reports all four coefficients under three fixed-effect

structures). The two cleanly comparable unsanctioned regions, Mexico and Colombia, decline; only the sanctioned region, Venezuela, shows the near-pipeline methane rise that the diluent-dependence physics predicts. The Ecuadorean Oriente also trends upward but follows a different profile (steady operational expansion, not a sanctions shock) and shows clean pre-trends, so it does not act as a heavy-oil-extraction baseline that rationalises the Venezuelan rise. This cross-country difference is the sharpest evidence that the within-Venezuela strain pattern is sanctions-driven rather than crude-type-driven.

8.4 Policy implications

Each of the three channels has its own policy lever: field-level revenue targeting for composition, operating-licence carve-outs for maintenance, and supply-input controls for strain. Because the levers are distinct, sanctions can hit the regime's revenue base while keeping the maintenance and strain channels quiet. The G7 price cap on Russia targets revenue from the dirtiest fields; the Chevron carve-out in Venezuela licenses maintenance inputs at those same fields. A package that combines both is coercive-equivalent to a uniform package but environmentally far cheaper; uniform packages forgo this opportunity and let the composition shift toward the dirtiest fields run its course.

Counterfactual exercises. The decomposition in equation (13) lets us trace each policy lever to a direct effect on R. We consider three.

Diluent licensing is the single most consequential lever in our setting. It eliminates the strain channel by restoring naphtha imports at the wellhead. Assigning $r_{\text{strain}} = 0$ in equation (13) delivers $R \approx 1.6$, roughly a halving of the damage multiplier. The naphtha-price elasticity $\theta_0 = +3.05$ (Table 5) is the structural anchor for this number: the channel responds sharply to diluent supply because heavy crude cannot move through pipelines without it.

Operating-licence carve-outs at all heavy fields, modelled on the Chevron General Licence, eliminate the maintenance channel. The OPEX-restoration elasticity from the Chevron quasi-experiment, $\eta_5 = -0.42$ (Table 7, column 1), implies that reversing the -0.61 log-unit OPEX shock at heavy fields (Table 3, column 6) removes the $r_{\text{maint}} \approx 0.59$ contribution. Combined with diluent licensing, R falls to roughly the proportional-decline benchmark of 1.0.

Composition steering through price caps, methane-intensity import standards, or field-named sanctions tilts surviving production toward cleaner fields. The lever does almost no environmental work in Venezuela: 94% of pre-sanctions production was already at heavy fields, the within-heavy intensity gap is small in production-weighted terms (12.9 vs. 8.0 m^3/bbl), and the surviving large heavy fields are mildly cleaner per barrel than the small fields that exited. There is no large clean tier inside Venezuela to preserve.

The lever matters elsewhere. The G7 price cap on Russia (Crozet and Hinz, 2020) is one operational implementation: by capping the per-barrel sale price (initially \$60/bbl in December 2022, lowered to \$47.60 in September 2025), it makes high-marginal-cost legacy fields, which are also

dirtier per barrel, unprofitable first, concentrating surviving production on the cleaner Western-Siberian super-giants. The EU's 2024 Methane Regulation, which sets methane-intensity thresholds for imported natural gas, is the analogous instrument for the oil side under regulatory development. Field-named sanctions are a third implementation, practical when the dirty tier is concentrated in a small number of named assets and impractical when it is hundreds of legacy fields spread across a basin.

The first two levers are environmentally additive: together they bring damage back to $R = 1$. They are also coercive-equivalent, but the equivalence requires a specific design feature. Licensing maintenance inputs at heavy fields does not relax the revenue-side bite of sanctions only if the carve-out severs the link between operations and state revenue. The November 2022 OFAC General Licence 41 illustrates how this is done. The licence permits Chevron to import operational inputs — spare parts, foreign expertise, banking services — at the four covered joint-ventures, but it explicitly prohibits payment of taxes, royalties, or dividends to the Venezuelan state, and routes the proceeds of Chevron's oil sales toward settlement of pre-existing PDVSA debt to Chevron rather than to the Venezuelan treasury. Even when output at the covered fields rises modestly post-licence, the regime sees no new revenue from those additional barrels. A general-purpose carve-out at all heavy fields would have to inherit this same revenue-routing structure: license inputs, not outputs, and keep export and financial sanctions in place. Under that design the combined package is a Pareto improvement over the unconditional sanctions baseline.

The decomposition also implies a measurement template. Aggregate flaring volumes are useful but insufficient: methane concentrations capture leakage that flaring data miss, and biodiversity outcomes integrate flaring, methane, and the co-emitted local pollutants at the spatial scale at which households actually live. The Chevron event study and the methane analysis in this paper together provide a template that can be applied to other sanctioned jurisdictions where production-side data are sparse and satellite remote sensing is the only credible source of variation. Without the three-channel framework, R is just a number; with it, R is the sum of three margins each of which can be targeted by a different policy lever.

9 Conclusion

Sanctions on resource-rich autocracies are widely understood to reduce aggregate output, but less commonly understood to reshape *how* the surviving output is produced. We document three reinforcing channels in the Venezuelan case, composition, maintenance, and strain, and identify each through a distinct research design, with the November 2022 Chevron licence carve-out providing the within-country, within-firm variation that recovers the structural responsiveness of pollution per barrel to the maintenance budget. Combining the channel-specific estimates, total environmental damage from Venezuelan oil extraction in 2022 to 2024 is approximately three times what the proportional decline of output alone would predict, with the maintenance and strain channels at heavy-oil fields carrying essentially all of the gap and the composition

channel determining where that damage lands. The same qualitative pattern, falling or stagnating output, surviving production concentrating on the largest dirtiest fields, rising per-barrel emissions, appears in Russia after 2014, 2017, and 2022 and in Iran after 2018, though only Venezuela offers the within-firm licence shock that makes the structural identification possible.

The decomposition has direct implications for sanctions design. It implies that sanctions design has two distinct objectives, not one: maximising coercive impact on the regime, and minimising the environmental damage that the production reshuffling sanctions induce. The three channels operate through distinct policy levers, so these objectives are not always in conflict. A package that targets revenue from the dirtiest large fields while licensing maintenance and parts at those same fields hits the regime's revenue base hard while leaving the maintenance and strain channels quiet. The Chevron carve-out shows the design feature that makes this work in practice: the licence covers operational inputs but explicitly bars payment of taxes, royalties, or dividends to the Venezuelan state, routing oil-sale proceeds toward settlement of pre-existing PDVSA debt to Chevron, so any incremental production at the covered fields generates no new state revenue. The G7 oil price cap on Russia and this OFAC General Licence in Venezuela have begun to implement the broader logic in different ways. The Chevron quasi-experiment identifies the maintenance channel responding quickly to budget restoration, with biodiversity recovery of a magnitude consistent with the predicted pollution decline. The strain channel runs through physical bottlenecks like diluent supply for Venezuelan extra-heavy crude, so sanctions designs that interfere with these bottlenecks activate strain damage directly while designs that leave them alone do not. The wider takeaway is that the environmental costs of sanctions are first-order, predictable from the structure of extraction technology, and amenable to design through policy levers already in the sanctions toolkit.

References

AFESORGBOR, S. K. (2019): "The Impact of Economic Sanctions on International Trade: How Do Threatened Sanctions Compare with Imposed Sanctions?" *European Journal of Political Economy*, 56, 11–26.

AFESORGBOR, S. K. AND S. STEINBACH (2024): "The Collateral Damage of Economic Sanctions on the Environment," *Journal of Environmental Economics and Management*, 128, 103031.

AGERTON, M., B. GILBERT, AND G. B. UPTON, JR. (2023): "The Economics of Natural Gas Flaring and Methane Emissions in US Shale: An Agenda for Research and Policy," *Review of Environmental Economics and Policy*, 17, 251–273.

AHLVIK, L., J. J. ANDERSEN, J. H. HAMANG, AND T. HARDING (2026): "Quantifying Supply-Side Climate Policies," *Review of Economic Studies*, 93.

ASKER, J., A. COLLARD-WEXLER, C. DE CANNIERE, J. DE LOECKER, AND C. R. KNITTEL (2024):

"Two Wrongs Can Sometimes Make a Right: The Environmental Benefits of Market Power in Oil," Working Paper 33115, National Bureau of Economic Research.

ASKER, J., A. COLLARD-WEXLER, AND J. DE LOECKER (2014): "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 122, 1013–1063.

BAHAR, D., S. BUSTOS, J. R. MORALES, AND M. A. SANTOS (2019): "Impact of the 2017 Sanctions on Venezuela: Revisiting the Evidence," Tech. rep., Brookings Institution.

BARTELSMAN, E., J. HALTIWANGER, AND S. SCARPETTA (2013): "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, 103, 305–334.

BLUNDELL, W. AND A. KOKOZA (2022): "Natural Gas Flaring, Respiratory Health, and Distributional Effects," *Journal of Public Economics*, 208, 104601.

CENTER ON GLOBAL ENERGY POLICY, COLUMBIA SIPA (2023): "Russia's Methane Emissions and the War in Ukraine," Tech. rep., Columbia University SIPA.

CONGRESSIONAL RESEARCH SERVICE (2024): "Venezuela: Overview of U.S. Sanctions," Tech. Rep. IF10715, CRS.

COULOMB, R., F. HENRIET, AND L. REITZMANN (2026): "'Bad' Oil, 'Worse' Oil and Carbon Misallocation," *Review of Economic Studies*, 93, 404–437.

COVERT, T. R. AND R. KELLOGG (2025): "Environmental Consequences of Hydrocarbon Infrastructure Policy," Working Paper 23855, National Bureau of Economic Research, conditionally accepted, *Journal of Political Economy*.

CROZET, M. AND J. HINZ (2020): "Friendly Fire: The Trade Impact of the Russia Sanctions and Counter-Sanctions," *Economic Policy*, 35, 97–146.

EGLINGTON, S. M., D. G. NOBLE, AND R. J. FULLER (2012): "A Meta-Analysis of Spatial Relationships in Species Richness across Taxa: Birds as Indicators of Wider Biodiversity in Temperate Regions," *Journal for Nature Conservation*, 20, 301–309.

EGOROV, K., V. KOROVKIN, A. MAKARIN, AND D. NIGMATULINA (2025): "Trade sanctions against Russia: Stylized facts," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 115, 578–582.

FERNÁNDEZ-VILLAVERDE, J., Y. LI, L. XU, AND F. ZANETTI (2025): "Charting the Uncharted: The (Un)Intended Consequences of Oil Sanctions and Dark Shipping," Working Paper 33486, National Bureau of Economic Research.

FIORETTI, M., J. HE, AND J. TAMAYO (2025): "Prices and concentration: a u-shape? Theory and evidence from renewables," *arXiv preprint arXiv:2407.03504*.

FIORETTI, M., A. IARIA, A. JANSSEN, C. MAZET-SONILHAC, AND R. K. PERRONS (2026a): "Sovereign Hold-Up and Technology Adoption: Evidence from the North Sea," Working Paper.

FIORETTI, M., V. SAINT-JEAN, AND S. C. SMITH (2026b): "The shared costs of pursuing shareholder values," *arXiv preprint arXiv:2103.12138*.

FOWLIE, M., S. P. HOLLAND, AND E. T. MANSUR (2012): "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program," *American Economic Review*, 102, 965–993.

FU, Q., Y.-E. CHEN, C.-L. JANG, AND C.-P. CHANG (2020): "The Impact of International Sanctions on Environmental Performance," *Science of the Total Environment*, 745, 141007.

GRAFF, T. (2026): "Depreciation and Growth: Evidence from Machine Repair in Uganda," Job Market Paper, Harvard University.

GREGORY, R. D. AND A. VAN STRIEN (2010): "Wild Bird Indicators: Using Composite Population Trends of Birds as Measures of Environmental Health," *Ornithological Science*, 9, 3–22.

GUTMANN, J., M. NEUENKIRCH, AND F. NEUMEIER (2023): "The Economic Effects of International Sanctions: An Event Study," *Journal of Comparative Economics*, 51, 1214–1231.

HANCOCK, S. E., J. R. WORDEN, S. SAATCHI, ET AL. (2025): "Methane Emissions from South America: Insights from TROPOMI and GOSAT," *Atmospheric Chemistry and Physics*, 25, 797–816.

HARSTAD, B. (2012): "Buy Coal! A Case for Supply-Side Environmental Policy," *Journal of Political Economy*, 120, 77–115.

HART, O., D. THESMAR, AND L. ZINGALES (2024): "Private sanctions," *Economic Policy*, 39, 203–268.

HAUSMANN, R. AND F. MUCI (2019): "Don't Blame Washington for Venezuela's Oil Woes: A Rebuttal to Weisbrot and Sachs," Tech. rep., Americas Quarterly.

HOGAN, D., E. G. FRANK, J. GANTOIS, AND A. MISSIRIAN (2025): "The Effectiveness of Local Conservation Ballots," in *AEA Papers and Proceedings*, American Economic Association, vol. 115, 403–408.

HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.

INTERNATIONAL ENERGY AGENCY (2024): "Global Methane Tracker 2024," IEA Report.

ITSKHOKI, O. AND E. RIBAKOVA (2024): "The economics of sanctions: From theory into practice," *Brookings Papers on Economic Activity*, 2024, 425–470.

KATOVICH, E. (2023): "Quantifying the Effects of Energy Infrastructure on Bird Populations and Biodiversity," *Environmental Science & Technology*, 58, 323–332.

KELLOGG, R. (2024): "The End of Oil," Working Paper 33207, National Bureau of Economic Research, revise and resubmit, *Review of Economic Studies*.

KHALID, U., M. T. ALI, L. OKAFOR, AND O. I. SANUSI (2024): "Do Sanctions Affect the Environment? The Role of Trade Integration," *Research in Globalization*, 8, 100191.

LADE, G. E. AND I. RUDIK (2020): "Costs of Inefficient Regulation: Evidence from the Bakken," *Journal of Environmental Economics and Management*, 102, 102336.

LANTERI, A. AND A. A. RAMPINI (2025): "Financing the Adoption of Clean Technology," NBER Working Paper No. 33545.

LIANG, Y., I. RUDIK, AND E. Y. ZOU (2025): "Economic Activity and Biodiversity in the United States," *Journal of Public Economics*, 243, 105325.

LIANG, Y., I. RUDIK, E. Y. ZOU, A. JOHNSTON, A. D. RODEWALD, AND C. L. KLING (2020): "Conservation Cobenefits from Air Pollution Regulation: Evidence from Birds," *Proceedings of the National Academy of Sciences*, 117, 30900–30906.

LORENTE, A., T. BORSDORFF, A. BUTZ, O. HASEKAMP, J. DE BRUGH, A. SCHNEIDER, L. WU, F. HASE, R. KIVI, D. WUNCH, D. F. POLLARD, K. SHIOMI, N. M. DEUTSCHER, V. A. VELAZCO, C. M. ROEHL, P. O. WENNBERG, T. WARNEKE, AND J. LANDGRAF (2021): "Methane Retrieved from TROPOMI: Improvement of the Data Product and Validation of the First 2 Years of Measurements," *Atmospheric Measurement Techniques*, 14, 665–684.

LYON, D. R., B. HMIEL, R. GAUTAM, M. OMARA, K. A. ROBERTS, Z. R. BARKLEY, K. J. DAVIS, N. L. MILES, V. C. MONTEIRO, S. J. RICHARDSON, S. CONLEY, M. L. SMITH, D. J. JACOB, L. SHEN, D. J. VARON, A. DENG, X. RUDELIS, N. SHARMA, K. T. STORY, A. R. BRANDT, M. KANG, E. A. KORT, A. J. MARCHESE, AND S. P. HAMBURG (2021): "Concurrent variation in oil and gas methane emissions and oil price during the COVID-19 pandemic," *Atmospheric Chemistry and Physics*, 21, 6605–6626.

MADANI, K. (2020): "How International Economic Sanctions Harm the Environment," *Earth's Future*, 8, e2020EF001829.

——— (2021): "Have International Sanctions Impacted Iran's Environment?" *World*, 2, 231–252.

MADHOK, R. (2025): "Infrastructure, Institutions, and the Conservation of Biodiversity in India," *Journal of the Association of Environmental and Resource Economists*, 12, 1705–1745.

MARKS, L. (2022): "The Abatement Cost of Methane Emissions from Natural Gas Production," *Journal of the Association of Environmental and Resource Economists*, 9, 165–198.

MASNADI, M. S., H. M. EL-HOUJEIRI, D. SCHUNACK, Y. LI, J. G. ENGLANDER, A. BADAHDAH, J.-C. MONFORT, J. E. ANDERSON, T. J. WALLINGTON, J. A. BERGERSON, ET AL. (2018): "Global Carbon Intensity of Crude Oil Production," *Science*, 361, 851–853.

MONALDI, F., I. HERNÁNDEZ, AND J. LA ROSA REYES (2021): "The Collapse of the Venezuelan Oil Industry: The Role of Above-Ground Risks Limiting Foreign Investment," *Resources Policy*, 72, 102116.

NATHAN, B. J., T. LAUVAUX, W. THIÉRY, R. WETS, AND N. P. M. VAN LIPZIG (2024): "Assessing Methane Emissions from Collapsing Venezuelan Oil Production Using TROPOMI," *Atmospheric Chemistry and Physics*, 24, 6845–6863.

OMARA, M., D. ZAVALA-ARAIZA, D. R. LYON, B. HMIEL, K. A. ROBERTS, AND S. P. HAMBURG (2023): "Developing a Spatially Explicit Global Oil and Gas Infrastructure Database for Characterizing Methane Emission Sources at High Resolution," *Earth System Science Data*, 15, 3761–3790.

PLANT, G., E. A. KORT, L. T. MURRAY, J. D. MAASAKKERS, AND I. ABEN (2022): "Evaluating urban methane emissions from space-based observations," *Remote Sensing of Environment*, 268, 112756.

RESTUCCIA, D. AND R. ROGERSON (2008): "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics*, 11, 707–720.

RODRÍGUEZ, F. R. (2022a): "Sanctions and Oil Production in Venezuela," *Latin American Economic Review*, 31.

——— (2022b): "Sanctions, Economic Statecraft, and Venezuela's Crisis," Tech. rep., Sanctions & Security Research Project / Fourth Freedom Forum.

ŞEKERCIOĞLU, Ç. H. (2006): "Increasing Awareness of Avian Ecological Function," *Trends in Ecology & Evolution*, 21, 464–471.

SHAPIRO, J. S. AND R. WALKER (2018): "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade," *American Economic Review*, 108, 3814–3854.

SINGER, G. (2026): "Complementary Inputs and Industrial Development: Can Lower Electricity Prices Improve Energy Efficiency?" London School of Economics, Working Paper.

SULLIVAN, B. L., J. L. AYCRIGG, J. H. BARRY, R. E. BONNEY, N. BRUNS, C. B. COOPER, T. DAMOULAS, A. A. DHONDT, T. DIETTERICH, A. FARNSWORTH, D. FINK, J. W. FITZPATRICK, T. FREDERICKS, J. GERBRACHT, C. GOMES, W. M. HOCHACHKA, M. J. ILIFF, C. LAGOZE, F. A. LA SORTE, M. MERRIFIELD, W. MORRIS, T. B. PHILLIPS, M. REYNOLDS, A. D. RODEWALD, K. V. ROSENBERG, N. M. TRAUTMANN, A. WIGGINS, D. W. WINKLER, W.-K. WONG, C. L. WOOD, J. YU, AND S. KELLING (2014): "The eBird Enterprise: An Integrated Approach to Development and Application of Citizen Science," *Biological Conservation*, 169, 31–40.

SUN, C. AND S. WANG (2025): "Collateral Damage: The Environmental Consequences of U.S. Sanctions," *Journal of Environmental Economics and Management*, 130, 103103.

TOMBE, T. AND J. WINTER (2015): "Environmental Policy and Misallocation: The Productivity Effect of Intensity Standards," *Journal of Environmental Economics and Management*, 72, 137–163.

U.S. ENERGY INFORMATION ADMINISTRATION (2024): "Country Analysis Executive Summary: Venezuela," Tech. rep., EIA.

VEEFKIND, J. P., I. ABEN, K. MCMULLAN, H. FÖRSTER, J. DE VRIES, G. OTTER, J. CLAAS, H. J. ESKES, J. F. DE HAAN, Q. KLEIPOOL, M. VAN WEELE, O. HASEKAMP, R. HOOGEVEEN, J. LANDGRAF, R. SNEL, P. TOL, P. INGMANN, R. VOORS, B. KRUIZINGA, R. VINK, H. VISSER, AND P. F. LEVELT (2012): "TROPOMI on the ESA Sentinel-5 Precursor: A GMES Mission for Global Observations of the Atmospheric Composition for Climate, Air Quality and Ozone Layer Applications," *Remote Sensing of Environment*, 120, 70–83.

WEISBROT, M. AND J. D. SACHS (2019): "Punishing Civilians: U.S. Sanctions on Venezuela," *Challenge*, 62, 299–321.

WORLD BANK (2024): "Global Gas Flaring Tracker Report," Tech. rep., World Bank Group, Global Gas Flaring Reduction Partnership.

——— (2025): "Global Gas Flaring Tracker Report," Tech. rep., World Bank Group, Global Gas Flaring Reduction Partnership.

XU, Q. AND T. KIM (2022): "Financial Constraints and Corporate Environmental Policies," *Review of Financial Studies*, 35, 576–635.

Online Appendix

The Environmental Costs of Sanctions: Flaring, Venting, and Biodiversity in Venezuela

A A stylised two-field model of reallocation under sanctions

This appendix formalises the three-channel framework introduced in Section 2. The setup is the simplest two-field model in which a sanctions-induced rise in per-field fixed costs endogenously concentrates surviving production on the larger field; on top of the resulting composition wedge we layer a maintenance channel and a convex-damage strain channel, giving the damage decomposition $D = \sum_i \alpha_i(m_i)\, q_i + (\beta_i/2)\, q_i^2 + F\,\mathbb{1}\{q_i > 0\}$ used in the body.

A.1 Setup

A regime controls two heterogeneous fields $i \in \{S, L\}$: field S (small, low-cost, clean) has capacity Q_S, marginal cost c_S, and per-barrel external damage α_S; field L (large, high-cost, dirty) has capacity Q_L, marginal cost $c_L > c_S$, and damage $\alpha_L > \alpha_S$. Each operating field carries a per-period fixed cost F (overhead, sanctions-evasion infrastructure, foreign expertise). The regime takes the world price p as given, must meet a revenue target G (so total production $q_S + q_L \geq q^* \equiv G/p$), ignores externalities, and chooses production to minimise operating cost:

$$\min_{q_S, q_L \in [0, Q_i]} \quad c_S q_S + c_L q_L + F \cdot \mathbb{1}\{q_S > 0\} + F \cdot \mathbb{1}\{q_L > 0\} \quad \text{s.t.} \quad q_S + q_L \geq q^*. \tag{A1}$$

We assume $Q_S < q^* < Q_L$, so the small field alone cannot meet the target while the large field alone can.

A.2 The reallocation threshold

Two solutions satisfy the constraint: Option A operates both fields ($q_S = Q_S$, $q_L = q^* - Q_S$) at total cost $c_S Q_S + c_L(q^* - Q_S) + 2F$; Option B operates only the large field ($q_S = 0$, $q_L = q^*$) at cost $c_L q^* + F$. Option A is cheaper iff

$$F \;<\; (c_L - c_S)\, Q_S, \tag{A2}$$

i.e., when the per-field fixed cost is below the variable-cost penalty of routing Q_S barrels through the expensive field.

We model sanctions as a multiplicative shock to the per-field fixed cost, $F \mapsto \lambda F$ with $\lambda > 1$ (extra overhead from sanctions-evasion banking, parts substitution, security, foreign- expertise sourcing). After the shock, Option B beats Option A iff $\lambda F > (c_L - c_S) Q_S$. Combining,

$$F \;<\; (c_L - c_S)\, Q_S \;<\; \lambda F \tag{A3}$$

gives an interior range of parameters in which the regime *endogenously reallocates*: both fields operate pre-sanctions, only the large field operates post-sanctions. Equivalently, the sanctions multiplier must exceed $\lambda^* \equiv (c_L - c_S)Q_S/F$ for the small field to exit. λ^* falls when fixed costs are large relative to the marginal-cost gap, when the small field is small, or when the two fields have similar marginal costs; in the Venezuelan context, mature small fields with declining yields exhibit large F/Q_S ratios, so λ^* is modest.

A.3 Three reinforcing channels of damage

When the reallocation condition (A3) binds, total external damage rises through three reinforcing channels.

Composition. Reassigning Q_S barrels from S to L under constant per-barrel intensities raises damage by $(\alpha_L - \alpha_S)Q_S$.

Maintenance. Per-barrel intensity at field i depends on maintenance spending m_i: $\alpha_i(m_i) = \alpha_i^0\, g(m_i)$ with $g(0) = 1$, $g'(m_i) < 0$, $g''(m_i) > 0$. Sanctions block the imported parts and foreign expertise that maintenance requires, forcing m_i down at every operating field; the channel applies to both surviving fields, not only to those receiving reallocated volume.

Strain (convex damage). Replacing the linear L-segment $\alpha_L q_L$ with $\alpha_L^0 q_L + (\beta_L/2)\, q_L^2$ for $\beta_L > 0$ captures convex damage from gas-handling capacity binding at high throughput. Pushing q_L from $q^* - Q_S$ to q^* raises damage by $(\beta_L/2)Q_S\,(q_L^{\text{post}} + q_L^{\text{pre}})$.

Combining,

$$\Delta D \;=\; \underbrace{(\alpha_L - \alpha_S)\, Q_S}_{\text{composition}} \;+\; \underbrace{\alpha_L^0 \left[g(m_L^{\text{post}}) - g(m_L^{\text{pre}})\right] q^*}_{\text{maintenance}} \;+\; \underbrace{\tfrac{\beta_L}{2}\, Q_S\, (q_L^{\text{post}} + q_L^{\text{pre}})}_{\text{strain}}. \tag{A4}$$

Each term is positive under sanctions; an aggregate-output analysis that reports only "production fell by X%" captures none of them. The biodiversity outcomes used in Section 6.2 integrate the three channels at the local ecological scale.

A.4 Illustration

Figure A1 plots the marginal external damage curve against quantity. The vertical line at q^* is the revenue-target "demand". In each panel, the area under the marginal damage curve from 0 to q^* is the total external damage.

- **Panel (a)** (pre-sanctions). Both fields operate. The green rectangle is the S-side damage, $\alpha_S Q_S$. The red rectangle is the L-side damage, $\alpha_L(q^* - Q_S)$. The marginal damage curve (dashed purple) jumps from α_S to α_L at Q_S.

- **Panel (b)** (post-sanctions). The same Q_S barrels of production are now sourced from L. The marginal damage curve is flat at α_L. The green rectangle (dim) marks the baseline damage that *would* have come from S producing those barrels; the **gold** rectangle, sitting on top, is the additional damage from the reallocation, exactly $\Delta = (\alpha_L - \alpha_S)Q_S$. The red L-side rectangle is unchanged.

The two panels use identical parameters ($Q_S = 3$, $q^* = 5$, $\alpha_S = 8$, $\alpha_L = 35$, $c_S = 25$, $c_L = 35$, $F = 20$, $\lambda = 4$, $p = 60$). With these values:

- Pre: $F = 20 < (c_L - c_S)Q_S = 30$, so Option A (cost = 185) beats Option B (cost = 195). Both fields operate; $D_{\text{pre}} = 94$.
- Post: $\lambda F = 80 > (c_L - c_S)Q_S = 30$, so Option B (cost = 255) beats Option A (cost = 305). Only L operates; $D_{\text{post}} = 175$.
- $\Delta D^{\text{comp}} = (35 - 8) \cdot 3 = 81$, an 86% increase relative to D_{pre} from the composition channel alone. Adding plausible maintenance ($g(m_L^{\text{post}})/g(m_L^{\text{pre}}) \approx 1.5$, i.e. a 50% rise in per-barrel intensity at the unmaintained large field) and strain ($\beta_L = 2$) channels roughly doubles ΔD. The figure visualises only the composition component to keep the geometry clean; the maintenance and strain channels would each scale up the gold rectangle further.

B Heavy-oil gas-handling mechanism

This appendix expands the five-step physical chain summarised in Section 3.2 of the body, with the petroleum-engineering references that establish each step.

Gas-oil ratio at the reservoir. The gas-oil ratio at the Orinoco reservoir is determined by reservoir pressure, temperature, and composition; published industry-standard values for extra-heavy Orinoco crude are on the order of 30–80 standard cubic feet of associated gas per barrel of crude lifted (Monaldi et al., 2021). The operator does not choose this ratio; it is a thermodynamic co-product of the lift process.

Diluent dependence. At surface conditions, Orinoco crude has a viscosity on the order of 10^4 centipoise at 30°C, functionally a solid for pipeline purposes. To be transportable, it must be blended with 20–40% by volume of light hydrocarbons (naphtha, condensate, or upgraded synthetic crude) at the wellhead or at a local gathering centre (U.S. Energy Information Administration, 2024). Naphtha imports were the largest single category of US-sourced petroleum exports to Venezuela before the January 2019 PDVSA designation; the designation severed that supply channel and forced PDVSA to source smaller volumes through intermediary networks at premium prices (Fernández-Villaverde et al., 2025).

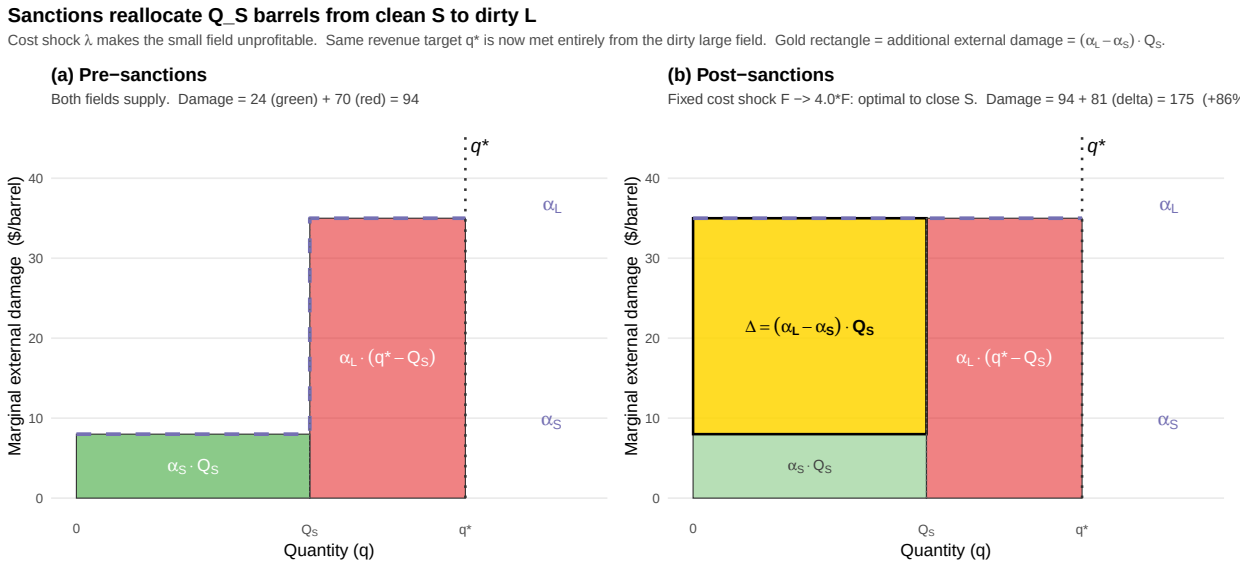


Figure A1: **Reallocation under sanctions.** The two panels share identical model parameters, differing only in the fixed-cost level (F in Panel a, λF in Panel b with $\lambda = 4$). Dashed purple step function: marginal external damage $\alpha(q)$ along the aggregate supply curve. Dotted vertical line: revenue target $q^* = G/p$. Green: damage from the small field, $\alpha_S Q_S$. Red: damage from the large field over $[Q_S, q^*]$, identical pre and post. Gold (Panel b only): the additional external damage from reallocation, $\Delta = (\alpha_L - \alpha_S)Q_S$. Parameters used: $Q_S = 3$, $q^* = 5$, $\alpha_S = 8$, $\alpha_L = 35$, $c_S = 25$, $c_L = 35$, $F = 20$, $\lambda = 4$, $p = 60$. Total damage rises from 94 to 175 (+86%).

Throttle constraints on heavy-oil wells. Orinoco production uses cyclic steam stimulation, steam-assisted gravity drainage, or electric submersible pumping. Once these systems are running, they cannot be paused for short periods without risk of permanent damage: wax and asphaltene precipitation in the wellbore, sand bridging, pump seizure, and ultimately well loss. The operator faces a binary choice when diluent runs short: continue lifting crude that cannot be moved (and stockpiling it), or shut the well permanently. Both choices break the gas-handling chain.

Storage asymmetry. When operators continue producing without sufficient diluent, the heavy crude accumulates in tank farms at the upgrader feed (Petropiar, Petromonagas, Petrocedeño, PetroSan Felix). Crude has a marginal cost of storage near zero per barrel; gas does not. Associated gas exits the separator at low pressure (5–50 psi), is uneconomical to compress for storage, and must be moved or burned within minutes of separation. The gas sinks are: sales-gas pipeline (capped by pipeline capacity and downstream demand); reinjection into the reservoir (requires functioning compressors and dedicated injection wells); feedstock for the upgrader's hydrogen-production unit and on-site power generation (requires the upgrader to be running). When the upgrader is starved of crude because the diluent-deficient gathering system cannot move it, all three sinks collapse simultaneously.

Flare and vent as residual sinks. Once the upstream sinks fail, the only ways to relieve pressure in the gathering system are to flare the gas (the safer choice) or to vent it (the cheaper choice when the flare stack itself has deteriorated). This is the physical reason for the asymmetric flaring-vs-venting shares we observe in the strain-channel results (Section 6): flare events spike during periods of acute diluent stress (high naphtha price), while venting (which generates the methane signal but no thermal anomaly visible to VIIRS) rises more steadily as overall maintenance degrades.

Why this physics is Orinoco-specific. Conventional medium-API Maracaibo crude (11–19° API) flows through pipelines without naphtha blending. Maracaibo gas-handling does not depend on diluent imports; its environmental degradation under sanctions runs through aging-pipeline leakage rather than the flare-substitution episodes documented for Orinoco. The asymmetry between the two basins is the physical counterpart of the asymmetry our model formalises: small clean fields can be cleanly shut down, large dirty ones cannot be.

C Flaring Data and Matching Procedure

This appendix describes the satellite flaring data and the two-stage procedure used to link World Bank VIIRS flare-site detections to Rystad Energy oil fields, producing the field-level flaring variables used throughout the paper.

C.1 Data Source

We obtain individual flare-site observations from the World Bank's Global Gas Flaring Tracker (GGFR), which publishes annual estimates derived from NOAA's Visible Infrared Imaging Radiometer Suite (VIIRS) nightfire algorithm. The input file, `2012-2024 Flare Volume Estimates by Individual Flare Location`, is publicly available from the World Bank's GGFR data portal.[18] Each record contains the geographic coordinates (latitude, longitude) of a detected flare site, its estimated annual flaring volume (in millions of cubic metres), the associated field name and operator (where available from VIIRS metadata), a location classification (onshore/offshore), and the year of observation. We retain all upstream flare-site-years for the six sample countries (Argentina, Brazil restricted to onshore, Colombia, Ecuador, Mexico, and Venezuela) over 2012–2024, yielding approximately 10,200 flare-site-year observations.

C.2 Two-Stage Matching Algorithm

The central empirical challenge is attributing each satellite-detected flare to a specific Rystad oil field. We implement a two-stage procedure that combines name-based and distance-based matching, with strict no-double-counting guarantees.

[18] World Bank Global Gas Flaring Reduction Partnership: https://www.worldbank.org/en/programs/gasflaringreduction.

Stage 1: Fuzzy Name Matching. Many World Bank flare-site records carry a location name that corresponds, sometimes imprecisely, to a Rystad field name. We exploit this by computing Jaro–Winkler string distances between every (WB flare-site name, Rystad field name) pair within the same country.[19] A match is accepted if and only if three conditions hold simultaneously:

1. *Distance threshold:* The Jaro–Winkler distance between the best-matching pair is ≤ 0.12 (on a 0–1 scale where 0 is an exact match).
2. *Gap requirement:* The best match beats the second-best match by at least 0.03, ensuring that the assignment is unambiguous.
3. *Mutual best match:* The WB field's best Rystad match and the Rystad field's best WB match agree, i.e., the mapping is bijective.

Of 721 candidate pairs evaluated across the six countries, 472 satisfy all three criteria. Name-matched flare-site-years are directly linked to the corresponding Rystad field, consuming approximately 3,975 flare-site-years and accounting for 44.5% of the total World Bank flaring volume in the sample.

Stage 2: Geographic Allocation. Flare-site-years *not* consumed by Stage 1 are allocated to nearby Rystad fields using a production-weighted inverse-distance gravity model. For each remaining flare-site-year, we identify all Rystad fields within a specified radius r whose annual production is strictly positive and compute allocation weights:

$$w_i = \frac{\text{Production}_i}{(d_i + d_0)^2},$$

where d_i is the great-circle (Haversine) distance in metres between the flare site and field i, and $d_0 = 100$ m is a small floor that prevents singularity when satellite and field coordinates coincide. This production-weighted inverse-squared-distance (PW-IDW) scheme assigns most of the flaring volume to the nearest and most productive fields while allowing smaller neighbours within the radius to receive a share. When all candidate fields within the radius have zero or missing production in a given year, the scheme falls back to pure inverse-squared distance ($1/(d_i + d_0)^2$). Weights are normalized within each flare-site-year so that $\sum_i w_i = 1$, ensuring that the allocated volume exactly equals the World Bank observation.

C.3 Mass-Balance Accounting

Every flare-site-year in the World Bank data is assigned one of three mutually exclusive dispositions: (i) allocated by name match, (ii) allocated by geographic proximity, or (iii) unallocated (no Rystad field within the radius). The set of flare-site-years consumed by name matching is excluded from

[19] We normalize names before comparison: lowercasing, stripping accents and generic suffixes (e.g., "field," "area"), and collapsing whitespace. Candidate pairs are restricted to those sharing at least one token (word) to reduce computation.

geographic allocation, preventing any flare-site-year from being counted twice. The total World Bank volume therefore decomposes exactly:

$$V^{\text{WB}} = V^{\text{name}} + V^{\text{geo}} + V^{\text{unalloc}}.$$

At the 15 km radius, the combined name-match and geographic allocation captures 91.5% of total World Bank flaring volume. The unallocated 8.5% consists of flare sites that have no Rystad field with positive production within 15 km, typically isolated flares in remote locations or near fields that Rystad records as inactive. Appendix Table C1 reports allocation rates by country.

C.4 Matching Coverage

Table C1 reports coverage of the satellite-to-field flaring allocation procedure at 7.5 km and 15 km matching radii, evaluated from the perspective of World Bank VIIRS satellite flare detections over 2012–2024. At 15 km, the volume allocation rate exceeds 90% for all countries except Brazil (onshore), where the low match rate reflects the predominantly offshore orientation of Brazil's flaring activity.

Table C1: Satellite Flaring Coverage by Country and Matching Radius

Country	Radius	WB flare sites	Matched sites	Match rate (%)	WB volume (Mm^3)	Allocated volume	Vol. alloc. rate (%)
Argentina	7.5 km	2,499	1,707	68.3	11,194.5	7,823.3	69.9
Argentina	15 km	2,499	2,231	89.3	11,194.5	10,105.2	90.3
Brazil	7.5 km	1,552	234	15.1	15,117.2	1,459.6	9.7
Brazil	15 km	1,552	288	18.6	15,117.2	1,852.1	12.3
Colombia	7.5 km	847	759	89.6	7,374.5	7,186.7	97.5
Colombia	15 km	847	782	92.3	7,374.5	7,213.6	97.8
Ecuador	7.5 km	998	975	97.7	14,383.7	14,338.1	99.7
Ecuador	15 km	998	998	100.0	14,383.7	14,383.7	100.0
Mexico	7.5 km	2,223	1,884	84.8	64,326.3	59,341.0	92.3
Mexico	15 km	2,223	1,979	89.0	64,326.3	61,593.5	95.8
Venezuela	7.5 km	2,057	1,443	70.2	101,701.8	76,523.9	75.2
Venezuela	15 km	2,057	1,895	92.1	101,701.8	100,438.7	98.8
Total	7.5 km	10,176	7,002	68.8	214,098.0	166,672.6	77.8
Total	15 km	10,176	8,173	80.3	214,098.0	195,586.8	91.4

Notes: Coverage is evaluated from the perspective of World Bank VIIRS satellite flare detections. A flare site is matched if it is allocated to at least one Rystad field via the two-step procedure (Jaro–Winkler name match plus production-weighted inverse-distance gravity). Volume allocation rates exceed site-count rates because large, high-volume flare sites near major fields are easier to match; unallocated sites are predominantly small or remote. Brazil includes onshore fields only; Mexico includes all fields.

C.5 Sensitivity of Matching to Radius Choice

A key consideration in the geographic allocation stage is the choice of matching radius. Oil fields can have irregular shapes extending several kilometres from the Rystad-reported centroid, and the

World Bank's satellite-detected flare coordinates need not coincide with the centroid of the field to which the flare belongs. To quantify this discrepancy, we exploit the name-matched pairs from Stage 1 as a validation set: because fuzzy string matching independently confirms that both the World Bank flare site and the Rystad field refer to the same physical asset, the great-circle distance between their respective coordinates provides a direct measure of the geographic offset for known matches.

Table C2 reports the distribution of these offsets for the 51 Venezuelan name-matched pairs. Over 43% of pairs have a flare-to-centroid distance exceeding 5 km, and one quarter exceed 7.5 km. At a 7.5 km matching radius, 13 confirmed matches, including fields such as Yucal Placer (14.8 km), Acema A (14.6 km), and Miga (12.1 km), would fall outside the radius and be excluded from the purely distance-based allocation. The 15 km radius recovers nine of these, reducing the miss rate from 25.5% to 7.8%.

For example, at the Yucal Placer field the World Bank flare site and the Rystad centroid are 14.8 km apart, so a 7.5 km radius would exclude the field while the 15 km radius captures it. This example is representative: oil-field infrastructure is spatially extended, and a radius that is too restrictive systematically excludes valid matches.

Table C2: Geographic Distance Between Name-Matched Flare Sites and Rystad Field Centroids: Venezuela

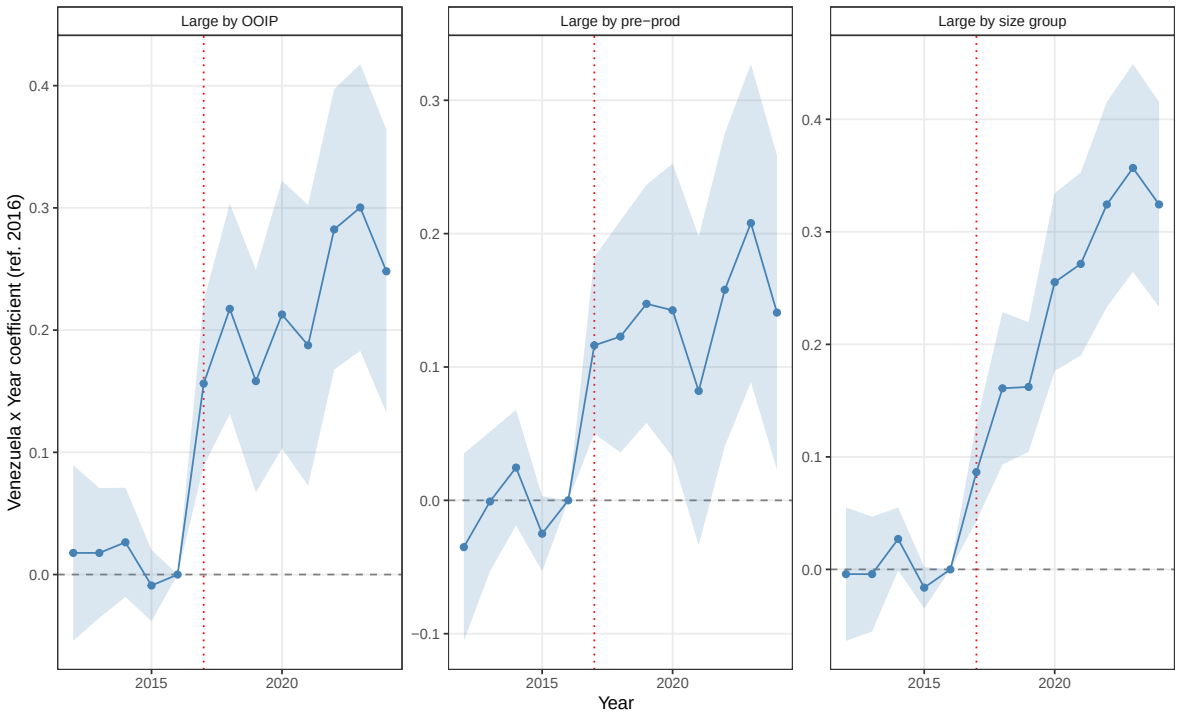
Distance band	Pairs	Share (%)
0–5 km	29	56.9
5–7.5 km	9	17.6
7.5–10 km	6	11.8
10–15 km	3	5.9
$>$15 km	4	7.8
Total	51	100.0

Notes: Distribution of great-circle distances between World Bank VIIRS flare-site coordinates and Rystad field centroids for the 51 accepted name-matched pairs in Venezuela. Name-matched pairs are confirmed by fuzzy string matching (Jaro–Winkler distance $\leq$ 0.12, mutual best match), so both sources identify the same physical field despite the coordinate discrepancy. Pairs in the upper distance bands would be missed by a purely distance-based approach at the corresponding radius threshold.

D Additional Exhibits

This appendix collects the baseline difference-in-differences table that underlies the central results in Sections 5–6, and the event-study counterparts of the composition, geography, and ownership pooled estimates.

Figure D1: Composition Channel: Event-Study Estimates of the Producing-and-Large Share



Notes: Per-year coefficients from a Venezuela-by-year event study of equation (2), with the omitted year set to 2016. Outcomes match the columns of Table 2: in the first three panels, $Y_{ift} = \mathbb{1}\{\text{prod}_{ift} > 0\} \cdot \text{Large}_i$ where Large is defined respectively by OOIP, pre-period mean production, and Rystad field-size group; the fourth panel uses heavy-oil indicator; the fifth panel is the unconditional $\mathbb{1}\{\text{prod}_{ift} > 0\}$ benchmark. Field and year fixed effects; standard errors clustered at the asset level. Shaded bands are 95% confidence intervals. Vertical dotted line at 2017 marks E.O. 13808. The producing-and-large coefficients are flat through 2016 and break at 2017–2019 in the size measures, consistent with the composition channel.

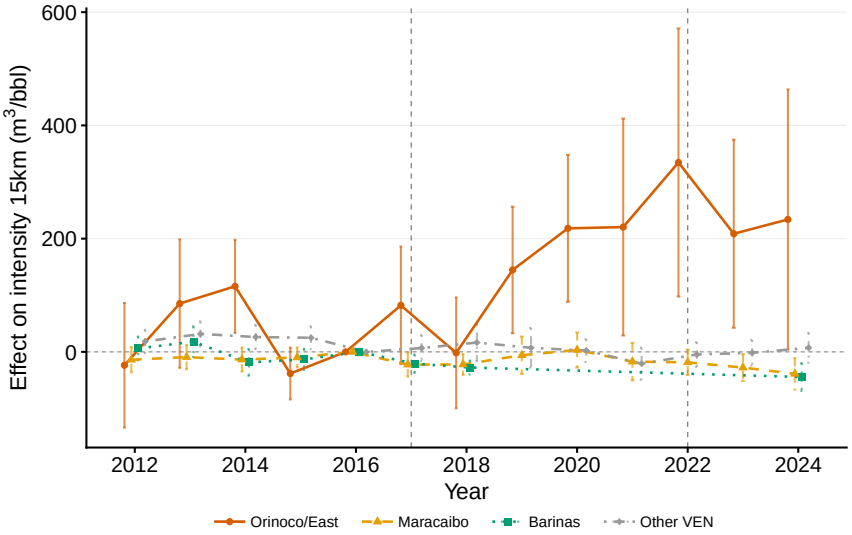
Cross-country wellhead-pipeline trends by basin. Table D2 reports the country–basin breakdown of the per-year wellhead-pipeline methane trend referenced in the body strain section, under three increasingly demanding fixed-effect structures. Only the Venezuelan Orinoco/East row shows a positive differential trend that survives all three; the Mexican Maya and Colombian Llanos rows are negative throughout. This rules out a generic heavy-oil-extraction story for the wellhead-elevation pattern observed in Venezuela.

Table D1: Effect of Sanctions on Flaring, Production, and Investment

Dependent Vars:	Flaring Vol. 7.5km (1)	Flaring Vol. 15km (2)	Flaring Int. 7.5km (3)	Flaring Int. 15km (4)	Prod. (log) (5)	OPEX (log) (6)	CAPEX (log) (7)	Cum. CAPEX (log) (8)
Venezuela × Post	-7.8181	-2.9685	162.8513***	118.3075***	-0.6618***	-0.6029***	-0.4098	-0.4113***
	(10.3728)	(6.5629)	(56.0808)	(44.7596)	(0.1412)	(0.1412)	(0.3145)	(0.1464)
Fixed effects								
Asset	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓
Mean dep. var.	29.8540	22.0782	76.9685	55.9316	12.0406	14.9673	12.7102	14.8554
Observations	5,023	7,795	5,023	7,795	11,983	11,983	11,983	11,983
Fields	772	1,097	772	1,097	1,097	1,097	1,097	1,097
R^2	0.90670	0.93664	0.68477	0.67769	0.86183	0.83240	0.65119	0.91985

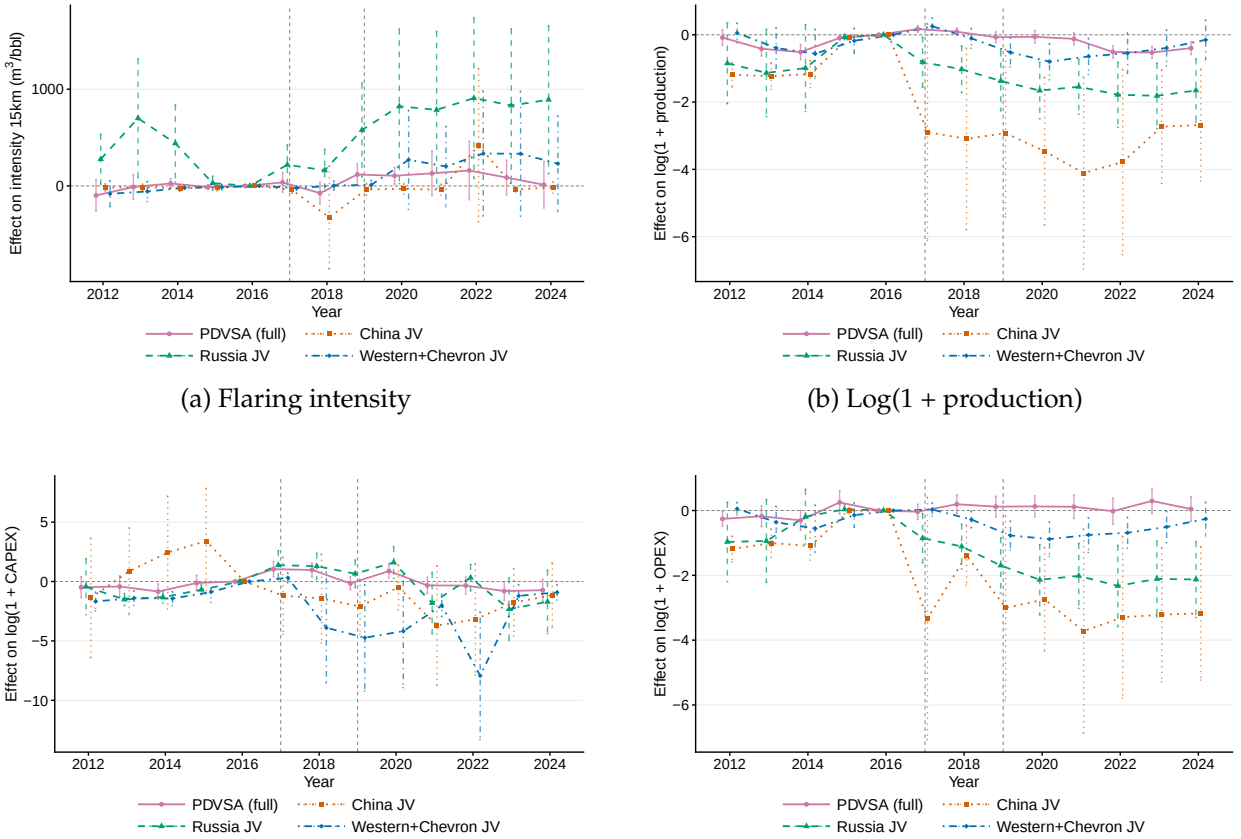
Notes: This table reports difference-in-differences estimates of the effect of US sanctions on Venezuelan oil fields (regression (6)). Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Post $= \mathbb{1}\{\text{year} \geq 2017\}$. Cols (1)–(2): allocated flaring volume (million m^3). Cols (3)–(4): flaring intensity (m^3/bbl), winsorized at the 99.5th percentile. Col (5): log(1 + crude production in bbl). Col (6): log(1 + OPEX in USD). Cols (7)–(8): log(1 + CAPEX in USD) and log(1 + cumulative CAPEX in USD). Cols (1) and (3) restrict to field-years with positive flaring at 7.5 km allocation radius; Cols (2) and (4) at 15 km; Cols (5)–(8) use the full sample. All models include field and year fixed effects. Standard errors clustered at the asset level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: Rystad Energy; World Bank VIIRS.

Figure D2: Flaring Intensity by Geographic Region: Event-Study



Notes: Event-study estimates of equation (3) (heavy/non-heavy DiD) interacted with geographic region indicators, normalised to zero in 2016. Outcome: flaring intensity (m^3/bbl) at the 15 km flare-allocation radius, winsorised at the 99.5th percentile. Geographic regions plotted: Venezuelan Orinoco/East, Venezuelan Maracaibo, the Mexican Maya basin, and Brazilian onshore. Reference: remaining LATAM fields. Field and year fixed effects; standard errors clustered at the asset level. Shaded bands are 95% confidence intervals. The intensity break in Venezuelan Orinoco appears at 2017–2019; Mexican Maya declines smoothly under Cantarell maturation; Brazilian onshore is approximately flat. Source: Rystad Energy UCube; World Bank VIIRS.

Figure D3: Event-Study Estimates by Ownership Group



(a) Flaring intensity

(b) Log(1 + production)

(c) Log(1 + CAPEX)

(d) Log(1 + OPEX)

Notes: Event-study estimates from equation (4) interacted with four ownership-group indicators: PDVSA-only (≥99% PDVSA equity), Russian, Chinese, and Western (including Chevron, collapsed). Ownership bins assigned using 2012–2016 baseline participation shares. Sample: six LATAM countries, 2012–2024, field-year level. Reference: non-Venezuelan LATAM fields, normalised to zero in 2016. Source: Rystad Energy UCube; World Bank VIIRS.

Table D2: Per-year log-XCH_4 trend at wellhead-and-pipeline cells, by country and basin.

Per-year trend ($\times$1,000)	(1)	(2)	(3)
LATAM baseline (pipe + well)	-0.92^{***} (0.03)	-0.54^{***} (0.03)	-0.54^{***} (0.03)
Venezuela — Orinoco/East	$+1.30^{***}$ (0.44)	$+1.51^{***}$ (0.40)	$+1.50^{***}$ (0.36)
Venezuela — Maracaibo	$+0.90^{***}$ (0.27)	$+1.32^{***}$ (0.27)	$+1.19^{***}$ (0.28)
Venezuela — Barinas	$+0.31^{***}$ (0.04)	$+0.82^{***}$ (0.08)	$+0.84^{***}$ (0.10)
Venezuela — Other	$+1.15^{***}$ (0.29)	$+1.44^{***}$ (0.33)	$+1.18^{***}$ (0.34)
Mexico — Maya basin	-2.51^{***} (0.15)	-2.14^{***} (0.12)	-2.13^{***} (0.10)
Ecuador — Oriente	$+1.51^{***}$ (0.17)	$+0.49^{***}$ (0.18)	$+0.67^{***}$ (0.15)
Colombia — Llanos	-0.81^{***} (0.14)	$+0.82^{***}$ (0.13)	$+0.98^{***}$ (0.14)
Cell FE	✓	✓	✓
Year-month FE	✓	✓	
Country $\times$ Year FE		✓	
Country $\times$ Year-month FE			✓
SE clustered at cell	✓	✓	✓
Cell-month observations	2,587,874	2,587,874	2,587,873

Specification. Each column is a separate regression of $\log \mathrm{XCH}_4$ on the time index t (years since 2018) interacted with a wellhead-and-pipeline indicator (cells within 20 km of an OSM+GEM pipeline *and* within 10 km of any oil-or-gas wellhead), with that interaction further interacted with seven country–basin indicators:

$$\log \mathrm{XCH}_{4ict} = \alpha_i + \tau_t + \phi\,(t \cdot \mathrm{PipeWell}_i) + \sum_g \psi_g\,(t \cdot \mathrm{PipeWell}_i \cdot \mathbb{1}\{i \in g\}) + \varepsilon_{ict},$$

where i indexes a 10×10 km TROPOMI cell, c its country, t calendar time in years from 2018, and $g \in \{\text{VEN Orinoco/East, VEN Maracaibo, VEN Barinas, VEN Other, MEX Maya, ECU Oriente, COL Llanos}\}$ are the country–basin indicators (LATAM Other is the omitted reference). Coefficients are reported as the per-year change in $\log \mathrm{XCH}_4$ multiplied by 1,000 to ease reading. The row labelled *LATAM baseline* is ϕ (the trend in the omitted reference); each subsequent row is the corresponding ψ_g (the differential trend in basin g, on top of the baseline). *Fixed effects.* Column (1) is the baseline specification used in the body: cell and year-month fixed effects, with cell-clustered standard errors. Column (2) adds country $\times$ year fixed effects, absorbing country-specific yearly shocks. Column (3) replaces year-month with country $\times$ year-month fixed effects, fully saturating country-by-time variation. *Reading.* Only the Venezuelan Orinoco/East row shows a sustained *positive* differential trend across all three specifications, consistent with the corridor-suppression-and-wellhead-elevation pattern that the diluent-dependence physics predicts. The Mexican Maya and Colombian Llanos rows are robustly *negative*, ruling out a generic heavy-oil-extraction explanation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Source: ESA Sentinel-5P TROPOMI L3; OpenStreetMap + Global Energy Monitor pipelines.

E TROPOMI Methane Data and Pipeline Proximity

This appendix describes the TROPOMI satellite data used in Section 6, explains the distinction between Level 2 and Level 3 data products and why we use the latter, documents the construction of the cell-level methane panel and pipeline proximity measures, and compares our reduced-form approach with the atmospheric inversion methodology of Nathan et al. (2024).

E.1 TROPOMI Instrument

The TROPOspheric Monitoring Instrument (TROPOMI) is aboard ESA's Sentinel-5 Precursor satellite, launched in October 2017, with routine data collection beginning in late 2018. TROPOMI provides near-global daily measurements of column-averaged dry-air methane mole fractions (XCH_4, in parts per billion). Each individual pixel measures approximately 5.5×7 km at nadir (roughly 38.5 km^2), so that a single daily overpass generates approximately 33,000 pixels covering Venezuelan territory.[20]

E.2 Level 2 versus Level 3 Data Products

TROPOMI data are disseminated at two processing levels that differ fundamentally in their granularity:

- *Level 2 (L2):* Individual satellite overpasses, with one observation per pixel per overpass. Each L2 pixel carries rich quality-assurance (QA) metadata (quality-value flags, aerosol optical depth in the near-infrared (NIR) and short-wave infrared (SWIR) bands, SWIR surface albedo, cloud radiance fraction, and retrieval uncertainty), enabling the analyst to apply custom, pixel-level filters. L2 data must be downloaded from Copernicus or SRON archives as individual orbit files.
- *Level 3 (L3):* Pre-aggregated composites (typically monthly) produced by applying the operational processor's standard QA filters and averaging surviving pixels onto a regular grid. L3 products are available via the Google Earth Engine (GEE) catalogue and are convenient for large-scale analysis. However, because pixel-level QA metadata are lost during aggregation, custom post-hoc filters, such as those in Nathan et al. (2024), cannot be applied retroactively.

Scale of the L2 data. With approximately 33,000 pixels per daily overpass covering Venezuela, a six-year panel (2019–2024) would require processing roughly $33{,}000 \times 365 \times 6 \approx 72$ million raw pixel observations, before any quality filtering. Even after applying standard QA thresholds (which eliminate most tropical pixels due to cloud cover), the filtered dataset would contain several million pixel-level records, each requiring individual retrieval-quality decisions. Downloading and processing this volume of L2 data represents a substantial data-engineering task that is beyond the scope of this study. Moreover, applying the full Nathan et al. methodology would additionally

[20] The Level 2 product specifications and validation are documented in Lorente et al. (2021).

require coupling L2 retrievals with an atmospheric transport model (GEOS-Chem or WRF-Chem) to invert observed concentrations into emission fluxes, a specialised atmospheric chemistry exercise rather than a reduced-form econometric analysis.

Why L3 suffices for our purpose. Our empirical design does not require pixel-level QA decisions. We exploit *within-cell temporal variation* via cell and year-month fixed effects, so any systematic filtering difference between L2 and L3 that affects the level of XCH_4 in a cell is absorbed by the cell fixed effect. Any common temporal trend in data quality is absorbed by the year-month fixed effect. The coefficient of interest, the *differential* time trend near pipelines, is identified from variation that is orthogonal to both dimensions. To verify this, we conduct sensitivity analyses (reported below in Section E.9) showing that stricter QA filtering within the L3 product *strengthens* rather than attenuates our coefficients. Moreover, monthly aggregated data fits better our biodiversity analyses as daily methane spikes are less likely to affect birds richnes and abundance metrics.

E.3 Data Acquisition and Grid Construction

Data Acquisition. Level 3 monthly composites of bias-corrected XCH_4 are retrieved via the Google Earth Engine (GEE) Python API. For each 10 km grid cell covering Venezuela, we extract the monthly mean of all quality-filtered TROPOMI pixels whose centroids fall within the cell boundary. Following the quality-assurance approach in Nathan et al. (2024), we retain only pixels with aerosol optical depth < 0.3, a threshold that improves retrieval accuracy in the tropics at the cost of reduced coverage during hazy periods, and discard retrievals flagged as low quality by the operational processor. Values outside the physically plausible range of 1,700–2,100 ppb are excluded as artifacts.

Spatial Grid. The regression grid consists of 10 km $\times$ 10 km cells covering Venezuela, defined on a UTM projection and reprojected to WGS 84 (EPSG:4326) for centroid extraction. The same grid is used for both the methane panel and the biodiversity analysis in Section 6. Cell centroids are computed as the geometric centre of each grid polygon. The final regression sample comprises 25,520 cell-months spanning 1,280 cells over 2018–2024, with a mean XCH_4 of 1,908 ppb and a standard deviation of 28.8 ppb.

Temporal Coverage and Limitations. The GEE Level 3 product provides consistent coverage from November 2018 onward. Because the first US sanctions Executive Order targeting PDVSA (E.O. 13850) was signed in November 2018 and the comprehensive embargo (E.O. 13857) followed in January 2019, the TROPOMI data begin essentially contemporaneously with the sanctions. There is therefore no substantive pre-sanctions baseline for a clean before-and-after comparison, which motivates the time-trend specification used in Section 6.

E.4 Within-Venezuela near-pipeline trend

The body strain regression of equation (5) uses naphtha-price variation to identify the diluent channel at near-pipeline cells. As a complement, a pooled within-Venezuela specification fixes location at the cell level and asks whether methane rose *faster* near pipelines than far from them, without the naphtha interaction:

$$\log \mathrm{XCH}_{4ct} \;=\; \alpha_c + \tau_t + \phi\,(t \cdot \mathrm{NearPipe}_{d,c}) + \varepsilon_{ct}, \tag{E1}$$

where t is a linear time index, $\mathrm{NearPipe}_{d,c}$ is the indicator that cell c falls within distance d of a pipeline, and α_c, τ_t are cell and year-month fixed effects. ϕ is the differential per-year trend in pipeline-vicinity methane. The same specification can be enriched with regional indicators (Other / Maracaibo / Orinoco) to test which Venezuelan region drives the trend.

Table E1 reports the estimates at $d \in \{10, 20, 50, 100\}$ km. Panel A gives the pooled trend; Panel B decomposes by Venezuelan region. Coefficients are multiplied by 1,000 for readability. The pooled trend is positive at every distance band ($+0.95$ at $d = 10$ km, falling monotonically to $+0.57$ at $d = 100$ km, all $p < 0.001$). The regional decomposition shows that Other-region cells (no Orinoco, no Maracaibo) drive most of the level signal, while Maracaibo shows a positive differential and Orinoco shows a near-pipeline trend with a negative differential at moderate distance bands. The cell fixed effect absorbs all time-invariant differences between cells, so the identifying variation is the within-cell rise in methane over the post-2018 window relative to the year-month average.

E.5 Naphtha-price lag structure

Table E2 reports the lag-extended version of the naphtha-price regression in body Table 5, with the contemporaneous coefficient and the four-month-lag structure estimated jointly on the Venezuelan Orinoco cell-month sample.

E.6 Cross-country LATAM extension

The within-Venezuela trend regression of equation (E1) can be extended to a six-country LATAM cell-month panel (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, and Venezuela) to confirm that the Venezuelan near-pipe methane rise documented in the body is sanctions-specific rather than a global near-pipe drift. We use a balanced-cell sub-sample of 3,183 cells observed in both pre-July 2021 and post-July 2021 ($n = 158{,}484$ cell-months), because the LATAM-wide TROPOMI download was first run on a 2,500-centroid sub-sample before being extended to the full $\sim$ 72,000-cell grid; the balanced restriction guarantees that the same set of cells contributes to both pre- and post-period identification of the trend. The pooled regression is

$$\log \mathrm{XCH}_{4ict} \;=\; \alpha_i + \tau_t + \phi_1\,(t \cdot \mathrm{NearPipe}_{d,i} \cdot \mathrm{VEN}_c) + \phi_2\,(t \cdot \mathrm{NearPipe}_{d,i}) + \delta\,(t \cdot \mathrm{VEN}_c) + \varepsilon_{ict}, \tag{E2}$$

Table E1: TROPOMI CH_4: pipeline proximity and time trend

	(1) $d = 10$ km	(2) $d = 20$ km	(3) $d = 50$ km	(4) $d = 100$ km
Panel A: Near pipeline × time trend				
NearPipe × trend	0.9490***	0.7951***	0.6039***	0.5715***
	(0.1500)	(0.1288)	(0.1153)	(0.1269)
Panel B: Regional heterogeneity (Near pipeline × time trend)				
NearPipe × trend (Other)	1.2918***	1.0075***	0.6866***	0.5075***
	(0.2192)	(0.1801)	(0.1361)	(0.1410)
× Maracaibo (diff.)	0.1646	0.1776	0.5016**	0.8479***
	(0.3652)	(0.2859)	(0.1952)	(0.1780)
× Orinoco (diff.)	-0.8526***	-0.5557**	-0.4248**	-0.2237
	(0.2943)	(0.2418)	(0.1695)	(0.1417)
Observations	25,520	25,520	25,520	25,520
Cell FE	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓

Notes: This table estimates the differential methane trend near pipeline infrastructure using TROPOMI satellite data (regression 6 and 5). Sample: 1,280 Venezuelan 10 km grid cells, 2018–2024, cell-month level. The dependent variable is $\log(\mathrm{XCH_4})$, where $\mathrm{XCH_4}$ is the column-averaged dry-air methane mole fraction (ppb) from TROPOMI; all estimates and standard errors are multiplied by 1,000 for readability. Each column corresponds to a different pipeline proximity threshold d. Panel A reports the pooled near-pipeline time-trend coefficient $\hat{\phi}$, where $t = \text{year} + (\text{month} - 1)/12 - 2018$ is a linear time trend. Panel B decomposes the trend by region: Row 1 reports the Other-region baseline ($\hat{\phi}_1$), Rows 2–3 report the Maracaibo ($\hat{\phi}_2$) and Orinoco ($\hat{\phi}_3$) differentials relative to this baseline. All specifications include cell and year-month fixed effects. Standard errors clustered at the cell level in parentheses. ***p<0.01, **p<0.05, *p<0.1. Source: ESA Sentinel-5P TROPOMI; OpenStreetMap.

Table E2: Naphtha-Price Lag Structure in the Orinoco Methane Panel

	(1) Contemp.	(2) Lag-1	(3) Lag-2	(4) Lag-3
$\log(\text{Naphtha}_t) \times \text{NearPipe}_{20}$	+3.053*** (0.907)			
$\log(\text{Naphtha}_{t-1}) \times \text{NearPipe}_{20}$		+2.235** (1.092)		
$\log(\text{Naphtha}_{t-2}) \times \text{NearPipe}_{20}$			+2.136** (1.019)	
$\log(\text{Naphtha}_{t-3}) \times \text{NearPipe}_{20}$				+2.882*** (1.067)
Cell FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Cell-month obs	9,138	9,138	9,138	9,138

Notes: Coefficients are the per-(log naphtha) change in $\log(\mathrm{XCH_4})$ in cells within 20 km of a pipeline, multiplied by 1,000. Sample is restricted to Venezuelan grid cells in the Orinoco Belt. Each column enters a single lag of the naphtha price interacted with NearPipe_{20} into the regression in isolation; the joint specification with all four lags entered together is omitted because the price series is highly autocorrelated, so the joint coefficients are unstable and individually uninterpretable. All specifications include cell and year-month fixed effects; standard errors clustered at the cell level in parentheses. The sample is the same 9,138 cell-months across 589 Orinoco cells in all four columns. *p<0.10, **p<0.05, ***p<0.01.

estimated separately at $d \in \{10, 20, 50, 100\}$ km, with cell α_i and year-month τ_t fixed effects, where $\text{NearPipe}_{d,i} = \mathbb{1}\{\text{cell } i \text{ within } d \text{ km of a pipeline}\}$ and $\text{VEN}_c = \mathbb{1}\{\text{cell } i \text{ in Venezuela}\}$. The three coefficients of interest trace four combinations of (country, pipeline-proximity) cells over time: ϕ_2 is the LATAM-wide pipeline-vicinity time trend (the non-Venezuela baseline near-pipe effect); δ is the Venezuela-wide trend differential vs LATAM (whether average Venezuelan methane rises faster than average LATAM methane, regardless of pipeline proximity); and ϕ_1 is the triple-interaction, the Venezuelan near-pipe trend differential relative to the LATAM-wide near-pipe trend.

Table E3 reports two summaries. Row 1 is the within-Venezuela near-pipe trend $\phi = \phi_1 + \phi_2$ from the Venezuela-only specification, the same object as in Table E1 once the units are aligned: it is positive and statistically significant at every distance band ($+0.00037$ to $+0.00078$ per year, equivalent to $+0.37$–0.78 in the per-mille convention used elsewhere in the appendix). Row 2 is δ, the Venezuela-vs-LATAM *country-wide* trend differential: negative at every distance band and significant at $d \leq 50$ km. The two rows pin down the spatial pattern. Venezuelan near-pipeline cells rise faster than their LATAM analogues, but Venezuelan territory *as a whole* rises more slowly than LATAM. The only configuration consistent with both signs is that Venezuelan cells *far from pipelines* are trending sharply down while *near-pipeline* cells trend up, consistent with the near-pipeline naphtha-price response documented in Section 6.1.

E.7 Comparison with Nathan et al. (2024)

Nathan et al. (2024) study methane emissions from Venezuelan oil infrastructure using a fundamentally different methodology. Where we use TROPOMI Level 3 monthly composites (10 km grid, Nov 2018–2024) and run a reduced-form cell+year-month panel regression to test whether methane concentration trends differ near pipelines after sanctions, Nathan et al. use TROPOMI Level 2 daily retrievals ($\sim$5.5$\times$7 km native pixels, 2018–2020) and atmospheric inversion models (GEOS-Chem, WRF-Chem) to recover emission fluxes (Tg/year) by basin. The two approaches answer different questions: ours, whether methane rose faster near pipelines under sanctions; theirs, how much methane Venezuela's oil sector emits in total. The two approaches are complementary. Nathan et al. quantify Venezuela's total methane budget at 7.5 Tg per year, with the Lake Maracaibo basin alone contributing 1.2 Tg (16% of the national total), and document that methane emission intensity approximately doubled, from 11% to 23%, even as production collapsed. Our reduced-form regressions cannot recover these flux estimates but confirm that the *concentration signature* of infrastructure leakage is detectable in the satellite record and is spatially concentrated near pipelines, with the steepest trends in the Maracaibo Basin.

E.8 Replication of Nathan et al.'s Key Patterns

To validate that our L3 data capture the same empirical patterns documented by Nathan et al. (2024) using L2 inversions, we test three of their central findings using our concentration panel.

Table E3: Methane Trend in Venezuela versus LATAM Controls

	$d = 10$ km (1)	$d = 20$ km (2)	$d = 50$ km (3)
VEN × NearPipe × t $(\hat{\phi}_1)$	+0.00103***	+0.00088***	+0.00072***
	(0.00017)	(0.00014)	(0.00012)
NearPipe × t $(\hat{\phi}_2)$	-0.00032***	-0.00029***	-0.00026***
	(0.00002)	(0.00001)	(0.00001)
VEN × t $(\hat{\delta})$	-0.00081***	-0.00089***	-0.00106***
	(0.00007)	(0.00007)	(0.00009)
Within-VEN near-pipe trend $\hat{\phi} = \hat{\phi}_1 + \hat{\phi}_2$	+0.00078***	+0.00067***	+0.00047***
	(0.00015)	(0.00012)	(0.00011)
Cell-month observations	2,588,253	2,588,253	2,588,253

Notes: Specification:

$$\log \mathrm{XCH}_{4ict} = \alpha_i + \tau_t + \phi_1(t \cdot \mathrm{NearPipe}_{d,i} \cdot \mathrm{VEN}_c) + \phi_2(t \cdot \mathrm{NearPipe}_{d,i}) + \delta(t \cdot \mathrm{VEN}_c) + \varepsilon_{ict},$$

where i indexes a 10×10 km cell, c its country, t calendar time in years from 2018, α_i a cell fixed effect, τ_t a year-month fixed effect, $\mathrm{NearPipe}_{d,i}$ the indicator that cell i lies within d km of an OSM+GEM combined pipeline, and VEN_c the indicator that cell i is in Venezuela. $\hat{\phi}_1$ is the triple-interaction Venezuela-near-pipe trend differential vs LATAM-near-pipe (the headline); $\hat{\phi}_2$ is the LATAM near-pipe-vs-far time-trend differential (non-Venezuela baseline pipeline effect); $\hat{\delta}$ is the Venezuela-overall trend differential vs LATAM (the country-wide shift, not pipeline-specific). The last row reports $\hat{\phi} = \hat{\phi}_1 + \hat{\phi}_2$ from a Venezuela-only specification (the within-Venezuela near-pipe trend, identical in interpretation to equation (E1)); the Venezuela-only value differs slightly from the joint-spec sum because of differences in fixed-effect support and clustering. Sample: cell-month TROPOMI log XCH_4 panel for the six-country LATAM grid (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2018–2024, restricted to the balanced-cell subsample of 3,183 cells (of $\sim$ 71,000) observed both pre- and post-July 2021. Coefficients are reported in their native log-XCH_4-per-year units and are not multiplied by 1,000 (unlike Table E1); to compare with the body, multiply by 1,000. Each column is a separate NearPipe threshold d. Standard errors clustered at the cell level in parentheses. ***p<0.01, **p<0.05, *p<0.1. Source: ESA Sentinel-5P TROPOMI; OpenStreetMap + Global Energy Monitor combined pipelines.

Pattern 1: Maracaibo is the methane hotspot. Nathan et al. find that Lake Maracaibo is Venezuela's dominant methane source, contributing 1.2 Tg/year. In our data, classifying cells by oil basin, mean XCH_4 is highest in the Maracaibo region (1,925 ppb) versus the Orinoco Belt (1,902 ppb) and the rest of Venezuela (1,907 ppb). Cells within 20 km of pipelines in the Maracaibo Basin show a mean of 1,951 ppb, the highest of any subgroup, consistent with localised infrastructure emissions.

Patterns 2 and 3: Methane rose while production collapsed, and intensity increased. Nathan et al. document that Venezuelan methane emissions increased despite severe production declines, with emission intensity approximately doubling (from 11% to 23%). Indexing both near-pipeline mean XCH_4 and national crude oil production to 2018 = 100, our data show that by 2020 the methane index rose to 100.2 while the production index fell to 39.4; by 2023, the methane index reached 102.8 against a production index of 51.9. The divergence between stable (or ris-

ing) concentrations and collapsing production implies a sharp increase in methane intensity. A concentration-based intensity proxy — near-pipeline mean XCH_4 divided by total Venezuelan production, indexed to 2018 = 100 — peaked at 254.6 in 2020, a 2.5-fold increase, before partially recovering as production stabilised. While concentration-to-production ratios are not directly comparable to emission-based intensities, the direction and order of magnitude are consistent with Nathan et al.'s inversion results.

All three patterns documented by Nathan et al. using L2 inversions are clearly visible in our L3 concentration data, suggesting that the L3 product is adequate for detecting the spatial and temporal signatures of infrastructure-related methane emissions.

E.9 Sensitivity to QA Filtering

A potential concern with using L3 data is that the pre-aggregated composites may include lower-quality pixels that would be discarded under the stricter L2 filters applied by Nathan et al. (2024). To address this, we re-estimate the main methane regressions (Tables E1 and 5) under two alternative specifications that approximate stricter QA filtering within the L3 framework:

1. *Low-uncertainty filter (B):* The L3 product includes a retrieval precision band (`XCH4_S5P_L3_CH4_uncertainty`), which reports the posterior uncertainty (in ppb) of each pixel's column-averaged methane estimate from the operational TROPOMI algorithm. We average this band to the cell-month level and drop cell-months with mean retrieval uncertainty above the 75th percentile (3.35 ppb), retaining 19,107 of 25,520 cell-months. This removes observations where the composite is likely dominated by hazy or marginal-quality retrievals.

2. *20-km aggregation (E):* We aggregate the 10 km grid to ∼20 km resolution, closer to the effective resolution of Nathan et al.'s analysis, yielding 14,482 cell-months.

Results. Under the low-uncertainty filter, the pooled time-trend coefficient $\hat{\phi}$ (multiplied by 1,000) *increases* from 0.95 to 1.08 at $d = 10$ km and from 0.80 to 0.94 at $d = 20$ km, with stronger statistical significance throughout. The Orinoco differential remains negative and significant (-0.48 vs. -0.56 at $d = 20$ km). The naphtha–Orinoco interaction remains stable (3.04 vs. 3.05 at $d = 20$ km). This finding, that stricter filtering *strengthens* the time-trend results, is reassuring: it suggests that any noise introduced by marginal pixels in the L3 product attenuates rather than inflates our estimates.

Under 20-km aggregation, the pooled coefficient weakens (0.53 at $d = 10$ km, versus 0.95 at baseline), consistent with spatial dilution: aggregating to coarser cells blends near-pipeline and far-from-pipeline pixels, attenuating the localised signal. Importantly, the regional decomposition retains the same qualitative pattern (the Orinoco differential remains negative and significant at all distances), confirming that the spatial structure of the results is robust.

Table E4: Sensitivity of TROPOMI Results to QA Filtering

Table	Coefficient	Variant		
		(A) Baseline	(B) Low uncertainty	(E) 20-km aggregation
Panel A: Pooled time trend (Table E1, Panel A)				
	$\hat{\phi}$ at $d = 10$ km	0.95^{***}	1.08^{***}	0.53^{***}
		(0.15)	(0.15)	(0.15)
	$\hat{\phi}$ at $d = 20$ km	0.80^{***}	0.94^{***}	0.64^{***}
		(0.13)	(0.13)	(0.17)
	$\hat{\phi}$ at $d = 50$ km	0.60^{***}	0.84^{***}	0.44^{***}
		(0.12)	(0.12)	(0.14)
	$\hat{\phi}$ at $d = 100$ km	0.57^{***}	0.94^{***}	0.55^{***}
		(0.13)	(0.13)	(0.15)
Panel B: Regional decomposition (Table E1, Panel B)				
	Other baseline at $d = 20$ km	1.01^{***}	1.14^{***}	0.93^{***}
		(0.18)	(0.17)	(0.23)
	Maracaibo Δ at $d = 20$ km	0.18	-0.09	0.31
		(0.29)	(0.28)	(0.37)
	Orinoco Δ at $d = 20$ km	-0.56^{**}	-0.48^{**}	-0.74^{***}
		(0.24)	(0.23)	(0.26)
Panel C: Naphtha price $\times$ Orinoco (Table 5, Panel B)				
	$\widehat{\text{Naphtha} \times \text{NearPipe}}$ at $d = 10$ km	2.68^{**}	2.08^{*}	-0.06
		(1.11)	(1.15)	(0.22)
	$\widehat{\text{Naphtha} \times \text{NearPipe}}$ at $d = 20$ km	3.05^{***}	3.04^{***}	-0.31
		(0.91)	(0.94)	(0.34)
Cell-months		25,520	19,107	14,482

Notes: All coefficients multiplied by 1,000. Standard errors (clustered at cell level) in parentheses. Variant (A) reproduces the baseline estimates from Tables E1 and 5. Variant (B) drops cell-months with TROPOMI retrieval uncertainty above the 75th percentile (3.35 ppb). Variant (E) aggregates the 10 km grid to ∼20 km resolution. Pipeline geometries combine OpenStreetMap and Global Energy Monitor (GOIT/GGIT) data. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

E.10 Pipeline Proximity Measures

Pipeline Data Sources. We combine pipeline geometries from two open sources. OpenStreetMap (OSM), obtained via the Overpass API, provides 404 pipeline segments (≈1,440 km) after excluding non-hydrocarbon features; its volunteer-contributed coverage is concentrated along major trunk corridors. We supplement OSM with curated transmission-level routes from the Global Energy Monitor's Global Oil Infrastructure Tracker (GOIT) and Global Gas Infrastructure Tracker (GGIT), which add 23 segments (≈2,670 km) after clipping to Venezuela's national boundary.[21] The combined dataset contains 427 features totalling approximately 4,110 km.

Coverage Benchmarking. The EIA estimates that Venezuela has 25 operational crude-oil pipeline systems totalling 2,139 miles (3,442 km), drawing on Global Energy Monitor data (U.S. Energy Information Administration, 2024). Our combined dataset slightly exceeds this figure because it also

[21]Global Energy Monitor, GOIT/GGIT Pipeline Routes: https://github.com/GlobalEnergyMonitor/GOIT-GGIT-pipeline-routes. License: Creative Commons CC BY 4.0.

includes gas pipelines (≈679 km from five GEM gas projects), refined-product lines, and segments that may no longer be operational. After excluding GEM gas projects, the oil-only component of our dataset (≈3,430 km) is consistent with the EIA benchmark.

Deduplication. Because GEM analysts use OSM as a spatial reference but independently retrace pipeline geometries from aerial imagery, the same physical pipeline can appear in both datasets with slightly offset line strings. A naive union would double-count these segments. We address this by (i) projecting all OSM features to UTM Zone 20N and creating a 1 km buffer around their union, (ii) computing the fraction of each GEM segment's length that falls inside this buffer, and (iii) dropping any GEM segment with more than 80% overlap. At a 1 km buffer, 13 of the 23 GEM segments have zero overlap with OSM, and the maximum overlap fraction is 7.5% (a single liquid-pipeline segment). The two sources thus cover largely distinct portions of the network, and no segment approaches the 80% removal threshold.

Alternative Sources. The Oil and Gas Infrastructure Mapping database (OGIM; Omara et al., 2023) provides an alternative global pipeline dataset. However, its Venezuela coverage is limited: the OGIM v2.7 Data Source References list only two ArcGIS Online maps (from 2017 and 2021) as Venezuela-specific pipeline sources, whereas countries such as Argentina and Brazil draw on eight or more official government datasets. We downloaded OGIM v2.7 and found 108 Venezuelan pipeline segments totalling approximately 16,000 km, of which 6,300 km are classified as oil, nearly twice the EIA's 3,442 km benchmark for operational crude-oil trunk lines. All segments lack facility names, operator information, and operational status, suggesting the underlying sources may include decommissioned, planned, or duplicate lines. By contrast, our OSM + GEM oil-only total (≈3,430 km) closely matches the EIA figure. We therefore retain the OSM + GEM combination as our pipeline source.

Distance Computation. For each 10 km grid cell, we compute the minimum geodesic distance (using the S2 geometry library) from the cell centroid to the nearest pipeline segment. We define a binary proximity indicator:

$$\text{NearPipe}_{d,i} = \mathbb{1}\{\text{distance from cell } i \text{ to nearest pipeline} \leq d \text{ km}\},$$

with distance thresholds $d \in \{10, 20, 50, 100\}$ km. The indicator is time-invariant (pipeline geometry is fixed to the OSM download date) and enters the regression specifications as an interaction with the time trend or post-period indicator. Cell fixed effects absorb the level of the indicator itself, so identification comes from the *differential temporal evolution* of methane concentrations in near-pipeline versus far-from-pipeline cells.

Regional Classification. To decompose the methane signal by oil-producing region, we assign grid cells to the Maracaibo Basin, Orinoco Belt, or Other category based on their geographic overlap

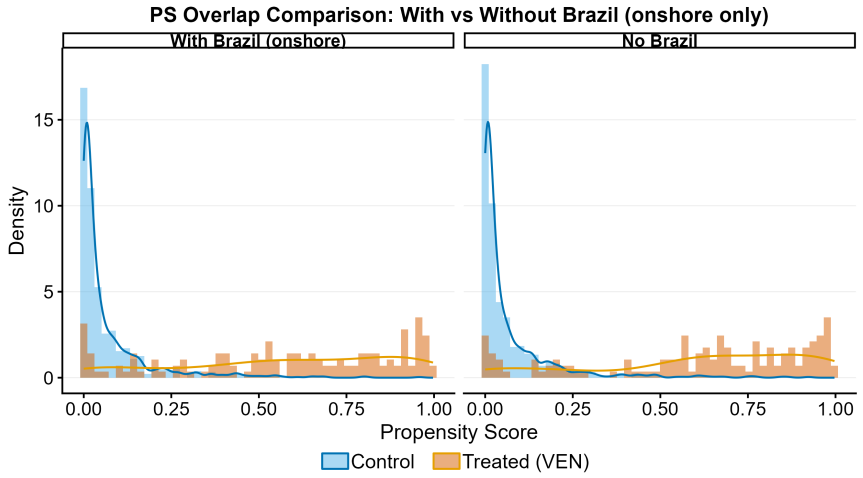
with field-level region assignments from Rystad Energy. The regional interaction terms $\text{NearMara}_{d,i}$ and $\text{NearOrinoco}_{d,i}$ in regression (5) are the product of $\text{NearPipe}_{d,i}$ and the corresponding regional indicator.

F Inverse Probability Weighting

The baseline difference-in-differences in Table D1 relies on the parallel-trends assumption for identification. A potential concern is that Venezuelan fields differ systematically from control-country fields along observable geological dimensions (reservoir depth, porosity, temperature, API gravity) and that these differences could generate divergent trajectories even absent sanctions. To address this, we re-estimate the main specifications using stabilised inverse probability weights (IPW) derived from a logistic propensity-score model. The propensity score is estimated using pre-period (2012–2016) field-level covariates: geological characteristics (water depth, reservoir depth, porosity, temperature), baseline production volumes, and an onshore/offshore indicator. The resulting weights upweight control fields that are observationally similar to Venezuelan fields and downweight those that are dissimilar, producing a pseudo-population in which the covariate distributions are balanced.

Figure F1 displays the propensity-score distributions for Venezuelan and non-Venezuelan fields. There is substantial overlap across most of the support, though the treated distribution has a concentration at higher propensity scores reflecting Venezuela's distinctive geology (Orinoco Belt extra-heavy deposits have few direct analogues elsewhere). We trim the propensity scores at the 2nd and 98th percentiles and cap the resulting weights at the 99th percentile to mitigate the influence of extreme scores.

Figure F1: Propensity-Score Distributions: Venezuela vs. Control Countries



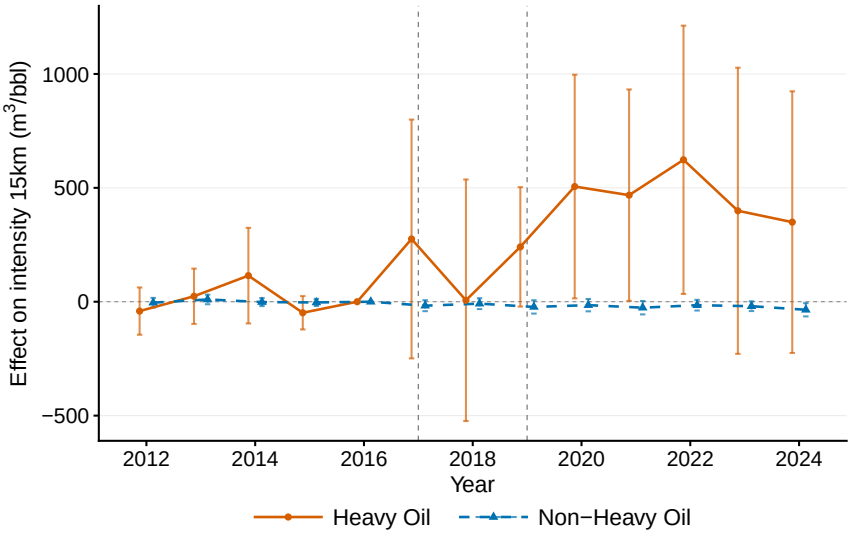
Notes: Kernel density plots of estimated propensity scores for Venezuelan (treated) and non-Venezuelan (control) fields. Propensity scores are estimated via logistic regression of a Venezuela indicator on pre-period (2012–2016) field-level covariates: water depth (m), reservoir depth (m), porosity (%), reservoir temperature (°C), mean baseline crude oil production (bbl/day), and an onshore indicator. The propensity score is used to construct stabilised inverse probability weights (IPW) for the reweighted estimates in Tables F1, F2, and F3. Propensity scores are trimmed at the 2nd and 98th percentiles, and resulting weights are capped at the 99th percentile, to limit the influence of extreme scores. Source: Rystad Energy UCube.

Table F1 reports the corresponding pooled estimates. The production coefficient implies a decline of approximately 42% ($\exp(-0.537) - 1 \approx -0.42$), comparable to the unweighted baseline,

and cumulative CAPEX declines by approximately 44% ($\exp(-0.573) - 1 \approx -0.44$). Flaring intensity at the 15 km radius rises substantially under IPW ($+259.8$ m^3/bbl, $p < 0.10$, larger than the unweighted estimate), consistent with the reweighting placing more emphasis on observationally similar control fields. The key takeaway is that reweighting to balance observable covariates does not materially alter the main findings: the sanctions-induced production collapse, investment decline, and flaring intensity increase are robust to IPW adjustment.

Figure F2 and Table F2 present the IPW-weighted heavy-oil heterogeneity results, confirming that the heaviness channel is robust to propensity-score reweighting.

Figure F2: IPW Event-Study Estimates: Flaring Intensity by Crude Oil Heaviness



Notes: IPW-weighted event-study estimates from regression (8) interacted with a heavy-oil indicator. Heavy: fields in the Orinoco Belt (East Venezuela Basin or Llanos) or with heavy crude share >50% (extra heavy, heavy 15–22° API, bitumen). Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Each coefficient $\hat{\beta}_k$ measures the year-k differential between heavy (or non-heavy) Venezuelan fields and non-Venezuelan LATAM fields, normalised to zero in 2016. Outcome: flaring intensity in m^3/bbl (allocated at 15 km radius, winsorized at the 99.5th percentile, restricted to field-years with positive flaring). Pooled production, CAPEX, and OPEX coefficients are reported in Table F2. Shaded areas show 95% confidence intervals. Standard errors clustered at the asset level. All models include field and year fixed effects. Stabilised inverse probability weights estimated from a logistic propensity-score model on pre-period geological and production covariates, propensity scores trimmed at the 2nd and 98th percentiles; weights capped at the 99th percentile. Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker.

Figure F3 presents the IPW-weighted ownership group event studies, and Table F3 reports the corresponding pooled estimates, confirming the robustness of the six-group heterogeneity results.

Table F1: Effect of Sanctions: Inverse Probability Weighting

Dependent Vars:	Flaring Vol. 7.5km (1)	Flaring Vol. 15km (2)	Flaring Int. 7.5km (3)	Flaring Int. 15km (4)	Prod. (log) (5)	OPEX (log) (6)	CAPEX (log) (7)	Cum. CAPEX (log) (8)
Venezuela × Post	-13.3773	-5.9391	314.9786	259.2161*	-0.5400***	-0.4660**	-0.8036	-0.5830***
	(22.7235)	(13.4226)	(234.9367)	(154.8156)	(0.2090)	(0.2361)	(0.6229)	(0.1098)
Fixed effects								
Asset	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓
Mean dep. var.	29.8540	22.0782	76.9685	55.9316	12.0406	14.9673	12.7102	14.8554
Observations	4,820	7,459	4,820	7,459	11,567	11,567	11,567	11,567
Fields	700	997	700	997	997	997	997	997
R^2	0.85670	0.87990	0.61712	0.64530	0.87689	0.85131	0.65467	0.91666

Notes: This table replicates Table D1 using inverse probability weighting (IPW) to improve covariate balance between Venezuelan and non-Venezuelan fields (regression (6)). Propensity scores are estimated via logit on pre-period geology and production characteristics; stabilised weights are trimmed at the 2nd and 98th percentiles and capped at the 99th percentile. Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Post $= \mathbb{1}\{\text{year} \geq 2017\}$. Cols (1)–(2): allocated flaring volume (million m^3). Cols (3)–(4): flaring intensity (m^3/bbl), winsorized at the 99.5th percentile. Col (5): log(1 + crude production in bbl). Col (6): log(1 + OPEX in USD). Cols (7)–(8): log(1 + CAPEX in USD) and log(1 + cumulative CAPEX in USD). Cols (1) and (3) restrict to field-years with positive flaring at 7.5 km; Cols (2) and (4) at 15 km; Cols (5)–(8) use the full sample. All models include field and year fixed effects. Standard errors clustered at the asset level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: Rystad Energy; World Bank VIIRS.

Table F2: Crude Oil Heaviness: Inverse Probability Weighting

Dependent Vars:	Flaring Vol. 7.5km (1)	Flaring Vol. 15km (2)	Flaring Int. 7.5km (3)	Flaring Int. 15km (4)	Prod. (log) (5)	OPEX (log) (6)	CAPEX (log) (7)	Cum. CAPEX (log) (8)
Heavy oil (Venezuela) × Post	-0.1973	-0.0148	374.1344	334.9423*	-0.5752**	-0.5462**	-1.0594	-0.6114***
	(22.2379)	(13.1766)	(263.5776)	(181.1025)	(0.2500)	(0.2725)	(0.7285)	(0.1153)
Non-heavy oil (Venezuela) × Post	-79.2304***	-27.2070	19.4102	-12.6349	-0.3711**	-0.0819	0.4227	-0.4466**
	(14.3418)	(37.6271)	(24.1966)	(14.2755)	(0.1619)	(0.3242)	(0.2949)	(0.1745)
Fixed effects								
Asset	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓
Mean dep. var.	29.8540	22.0782	76.9685	55.9316	12.0406	14.9673	12.7102	14.8554
Observations	4,820	7,459	4,820	7,459	11,567	11,567	11,567	11,567
Fields	700	997	700	997	997	997	997	997
R^2	0.85851	0.88026	0.62088	0.64973	0.87692	0.85146	0.65493	0.91666

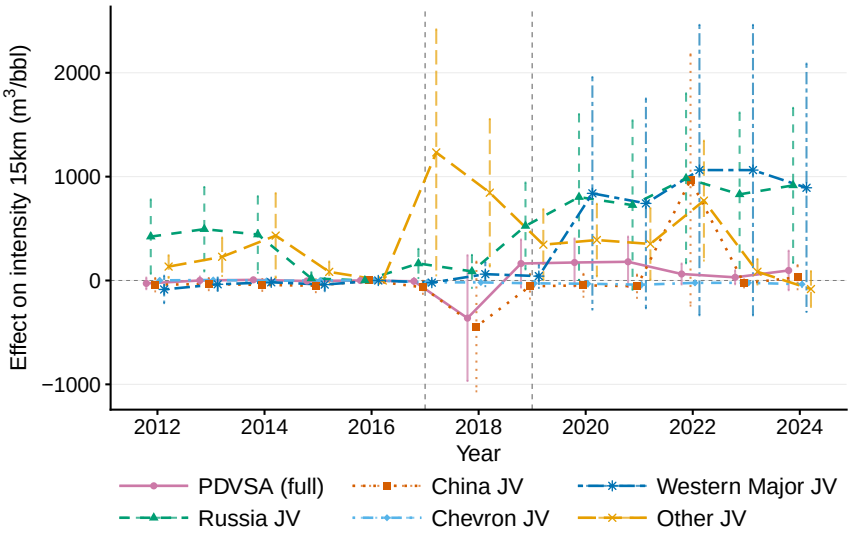
Notes: This table replicates Table 3 using inverse probability weighting (IPW) to improve covariate balance between Venezuelan and non-Venezuelan fields (regression (3)). Propensity scores are estimated via logit on pre-period geology and production characteristics; stabilised weights are trimmed at the 2nd and 98th percentiles and capped at the 99th percentile. Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Post $= \mathbb{1}\{\text{year} \geq 2017\}$. VEN Heavy: fields in the Orinoco Belt (East Venezuela Basin or Llanos) or with heavy crude share $> 50\%$ (Extra Heavy, Heavy 15–19, Heavy 20–22, Bitumen). VEN Non-Heavy: remaining Venezuelan fields. Reference group: non-Venezuelan LATAM fields. Cols (1)–(2): allocated flaring volume (million m^3). Cols (3)–(4): flaring intensity (m^3/bbl), winsorized at the 99.5th percentile. Col (5): log(1 + crude production in bbl). Col (6): log(1 + OPEX in USD). Cols (7)–(8): log(1 + CAPEX in USD) and log(1 + cumulative CAPEX in USD). Cols (1) and (3) restrict to field-years with positive flaring at 7.5 km; Cols (2) and (4) at 15 km; Cols (5)–(8) use the full sample. All models include field and year fixed effects. Standard errors clustered at the asset level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: Rystad Energy; World Bank VIIRS.

Table F3: Ownership Heterogeneity: Inverse Probability Weighting

Dependent Vars:	Flaring Vol. 7.5km (1)	Flaring Vol. 15km (2)	Flaring Int. 7.5km (3)	Flaring Int. 15km (4)	Prod. (log) (5)	OPEX (log) (6)	CAPEX (log) (7)	Cum. CAPEX (log) (8)
PDVSA × Post	-3.5715	6.2079	92.5588*	38.9608	-0.0788	0.0709	0.5049**	-0.5531***
	(9.7142)	(6.1075)	(47.5696)	(30.1079)	(0.0836)	(0.1903)	(0.2563)	(0.1296)
Russia × Post	0.9189	-6.3144	221.3453*	221.5442*	0.2895	-0.2061	0.2249	-0.8987**
	(20.4375)	(11.9722)	(113.6451)	(118.8860)	(0.5572)	(0.3250)	(0.4874)	(0.3485)
China × Post	-4.2421	-5.3910	16.5598	-14.7370	-4.0276***	-3.3403***	-5.1437**	0.1027
	(7.2843)	(4.4765)	(30.1969)	(15.6290)	(1.5007)	(1.1473)	(2.0338)	(0.6196)
Chevron × Post	-0.9653	-17.9461	-25.8149***	-29.2706***	0.7133	0.5168	-0.0492	-0.1663
	(4.1166)	(26.5897)	(9.6118)	(10.0618)	(0.5589)	(0.5635)	(0.2166)	(0.2719)
Western Majors × Post	-16.6493	-18.1082	478.5744	494.3365	-0.8775***	-0.7275**	-2.6241	-0.5274***
	(38.1416)	(39.1401)	(361.6821)	(377.9493)	(0.1863)	(0.2913)	(2.0364)	(0.0859)
Other (Venezuela) × Post	-34.8135***	-13.1930*	45.3420	407.4666**	-1.5718**	-1.4230*	-2.5159	-0.6044***
	(8.3049)	(7.8101)	(65.2520)	(191.0140)	(0.6425)	(0.8450)	(1.6524)	(0.1631)
Fixed effects								
Asset	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓
Mean dep. var.	29.8540	22.0782	76.9685	55.9316	12.0406	14.9673	12.7102	14.8554
Observations	4,820	7,459	4,820	7,459	11,567	11,567	11,567	11,567
Fields	700	997	700	997	997	997	997	997
R^2	0.85693	0.88027	0.62553	0.65519	0.88010	0.85388	0.65675	0.91670

Notes: This table replicates Table 4 using inverse probability weighting (IPW) to improve covariate balance between Venezuelan and non-Venezuelan fields (regression (3)). Propensity scores are estimated via logit on pre-period geology and production characteristics; stabilised weights are trimmed at the 2nd and 98th percentiles and capped at the 99th percentile. Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Post $= \mathbb{1}\{\text{year} \geq 2017\}$. Ownership bins are assigned using pre-2017 baseline participation shares with the following precedence: PDVSA ($\geq$ 99% PDVSA participation), Russia, China, Chevron, Western Majors, Other (any firm from each group $>$ 0%). Reference group: non-Venezuelan LATAM fields. Cols (1)–(2): allocated flaring volume (million m^3). Cols (3)–(4): flaring intensity (m^3/bbl), winsorized at the 99.5th percentile. Col (5): log(1 + crude production in bbl). Col (6): log(1 + OPEX in USD). Cols (7)–(8): log(1 + CAPEX in USD) and log(1 + cumulative CAPEX in USD). Cols (1) and (3) restrict to field-years with positive flaring at 7.5 km; Cols (2) and (4) at 15 km; Cols (5)–(8) use the full sample. All models include field and year fixed effects. Standard errors clustered at the asset level in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: Rystad Energy; World Bank VIIRS.

Figure F3: IPW Event-Study Estimates: Flaring Intensity by Ownership Group



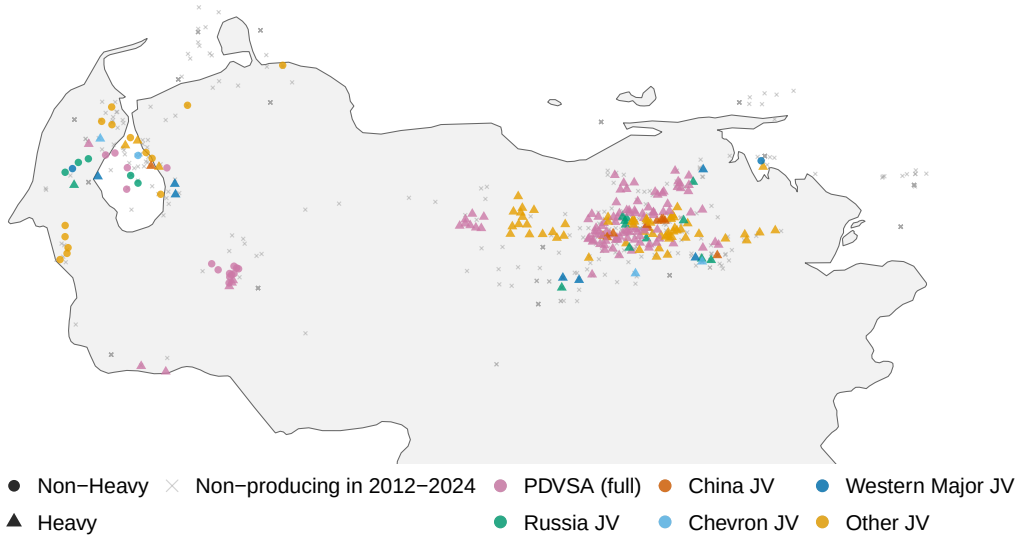
Notes: IPW-weighted event-study estimates from regression (8) interacted with six ownership-group indicators. Ownership bins assigned using 2012–2016 baseline participation shares (see Table 4 notes for precedence rule). Sample: six LATAM countries (Argentina, Brazil onshore, Colombia, Ecuador, Mexico, Venezuela), 2012–2024, field-year level. Each coefficient $\hat{\beta}_k$ measures the year-k differential between a Venezuelan ownership group and non-Venezuelan LATAM fields, normalised to zero in 2016. Outcome: flaring intensity in m^3/bbl (allocated at 15 km radius, winsorized at the 99.5th percentile, restricted to field-years with positive flaring). Pooled production, CAPEX, and OPEX coefficients are reported in Table F3. Shaded areas show 95% confidence intervals. Standard errors clustered at the asset level. All models include field and year fixed effects. Stabilised inverse probability weights estimated from a logistic propensity-score model on pre-period geological and production covariates, propensity scores trimmed at the 2nd and 98th percentiles; weights capped at the 99th percentile. Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker.

G Ownership and Size Heterogeneity

This appendix provides the size-by-ownership decomposition that underlies the within-Venezuela ownership discussion in Section 5.3 and the size-survival discussion in Section 5.1. The central question is whether the post-2017 reallocation toward giant fields reflects firm capacity (international joint ventures sustaining the plants they operate) or pure scale (any owner can keep the largest fields running).

Figure G1 maps Venezuelan fields in the analysis sample by ownership bin and heavy-crude status. Three patterns of geographic clustering stand out. First, the international joint ventures concentrate sharply in the Orinoco Belt extra-heavy region: 8/8 Chinese-partnered fields (100%), 10/11 Russian-partnered (91%), 33/36 Other-JV (92%), 2/3 Chevron, and 3/5 Western-Major fields lie in Orinoco East, versus 74/78 (95%) of PDVSA-only fields. Second, by crude type the partnership pattern flips: Chevron fields are uniformly heavy (3/3), Western Majors mostly heavy (4/5, 80%), but Chinese (1/8, 13%) and PDVSA-only (12/78, 15%) fields are predominantly non-heavy. Third, the geographic concentration in Orinoco is not a coincidence with size: the international JVs are also concentrated among the largest, most capital-intensive fields, which are the ones whose gas-handling chains are most exposed to the diluent shortage. The body-text ownership coefficients in Section 5.3 therefore mix three collinear margins (location, crude type, size); the size-by-ownership breakdown below isolates the size-versus-firm dimension.

Figure G1: Venezuelan Oil Fields by Ownership Group and Heaviness



Notes: Each marker represents a Venezuelan oil field in the analysis sample ($n = 233$ producing fields with valid coordinates and ownership-bin assignment, 2012–2024). Marker colour indicates the field's pre-period ownership bin (Chevron JV, Western Major JV, Chinese JV, Russian JV, PDVSA full ownership, Other JV); marker shape indicates crude heaviness (triangle = heavy, defined as Orinoco location or heavy-crude share above 50%; circle = non-heavy). Grey lines: oil-and-gas pipelines from the OSM + GEM combined network. Source: Rystad Energy UCube; OpenStreetMap; Global Energy Monitor.

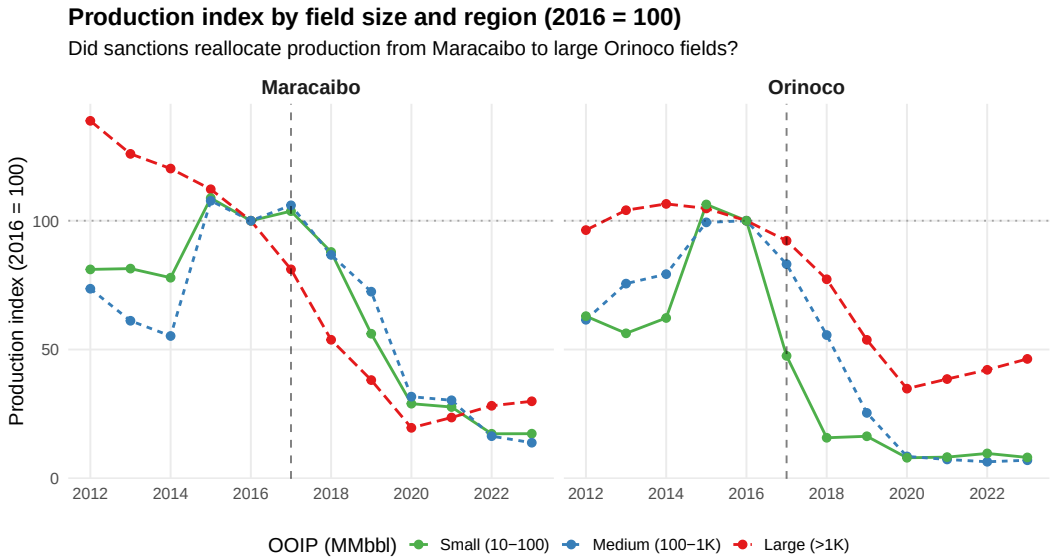
Table G1 breaks each Venezuelan field-size bin into PDVSA sole-ownership (100% equity) versus joint ventures (PDVSA typically holding 60–80% with a minority international partner). Among giant fields, the divergence is pronounced: PDVSA-sole giants *increased* flaring by 7.6% (from 5,192 to 5,584 million m^3/year) while their production fell by 68%, so per-barrel flaring intensity roughly tripled. Joint-venture giants, in contrast, reduced flaring by 23% alongside a 59% production decline, consistent with international partners maintaining some degree of operational discipline. The divergence is concentrated in the giant-field bin: at smaller sizes, PDVSA-sole and joint-venture trajectories track each other closely. The reallocation toward giants is therefore not a pure scale story; it is a scale-and-firm story, in which the surviving giants are disproportionately the ones with foreign technical partners. Figure G2 extends this by region and OOIP tercile and shows that the same composition shift, large fields surviving while small ones exit, holds within Orinoco and Maracaibo separately.

Table G1: Flaring and Production by Field Size and Ownership — Venezuela

		Flaring (M m^3/yr)			Production (M bbl/yr)			
Size	Owner	Pre	Post	Δ%	Pre	Post	Δ%	Fields
Small (<30)	Joint venture	40	9	-78.4	8	2	-77.0	45/33
	PDVSA sole	25	25	+2.2	1	1	-41.9	50/49
Medium (30-300)	Joint venture	920	352	-61.7	89	27	-69.1	46/46
	PDVSA sole	155	79	-49.3	14	6	-57.1	41/41
Large (300-1k)	Joint venture	432	367	-15.1	80	52	-34.7	12/12
	PDVSA sole	483	467	-3.4	33	23	-29.5	13/14
Giant (>1000)	Joint venture	830	638	-23.1	366	152	-58.6	13/13
	PDVSA sole	5,192	5,584	+7.6	358	116	-67.6	11/12

Notes: "PDVSA sole" denotes fields where PDVSA holds 100% of equity; "Joint venture" includes all fields with minority partners (PDVSA typically holds 60–80%) as well as non-PDVSA operators (8 fields). Pre: 2012–2016, post: 2017–2024. Fields column shows pre/post count. Flaring: VIIRS-allocated volume at 15 km. Production from Rystad.

Figure G2: Production Index by Region and Field Size



Notes: Production indexed to 2016 = 100, by Venezuelan oil-producing region and OOIP tercile. Large Orinoco fields retained roughly 46% of 2016 output through 2023; small and medium Orinoco fields collapsed to 7–8%. Maracaibo large fields retained only 30%. Source: Rystad Energy UCube.

H Biodiversity Appendix

eBird Sample. We use complete Venezuelan eBird checklists over 2012–2024 with duration 5–300 minutes, distance $\leq$ 5 km, and at least one species reported. The cleaned panel comprises 133,673 checklists across 1,508 unique 10 km grid cells. Species are matched to trophic guilds using EltonTraits 1.0 with a deterministic priority rule (91% match rate). Table H1 reports summary statistics for the panel; the spatial overlap of checklist locations with pipelines, TROPOMI methane, and VIIRS flare detections is shown in Figure 6 of the body.

Table H1: Summary Statistics: eBird Checklists in Venezuela, 2012–2024

	N	Mean	SD	P10	P25	P50	P75	P90
Panel A: Biodiversity outcomes								
Species richness	133,673	13.66	11.49	3.00	6.00	11.00	18.00	28.00
Shannon entropy	116,552	2.02	0.81	0.95	1.52	2.10	2.58	3.00
Gini–Simpson index	116,552	0.77	0.21	0.50	0.72	0.84	0.90	0.94
Log abundance	116,552	3.40	1.19	1.95	2.71	3.47	4.16	4.79
Pielou's evenness	113,026	0.86	0.13	0.69	0.82	0.91	0.95	0.97
Panel B: Effort variables								
Duration (minutes)	133,673	32.36	43.18	5.00	9.00	15.00	36.00	75.00
Distance (km)	133,673	0.63	1.29	0.00	0.00	0.10	0.69	1.84
Number of observers	133,673	1.59	1.22	1.00	1.00	1.00	2.00	3.00

Notes: This table reports summary statistics for the Venezuelan eBird sample used in the biodiversity analysis. Sample: complete eBird checklists from Venezuela, 2012–2024, assigned to 10 km grid cells, with duration between 5 and 300 minutes and distance $\leq$ 10 km for traveling protocols. Panel A: species richness is the number of unique species per checklist; abundance is the total number of individual birds counted per checklist. Panel B: effortcovariates used as controls in biodiversity regressions; traveling protocol equals 1 if the observer moved during the count, 0 if stationary. Source: eBird Basic Dataset (Cornell Lab of Ornithology); EltonTraits 1.0.

Robustness to Land-Cover and Built-Up Confounds. A natural concern is that the pollution-biodiversity association might be confounded with tree-cover loss or with urban build-up; both would mechanically suppress local bird abundance through habitat-loss channels independent of pollution. We address the concern with two diagnostics. First, augmenting the pollution-richness regression with Hansen tree-cover loss (2001–2023) and its post-period interaction leaves the CH_4 coefficient essentially unchanged while the tree-cover-loss interaction is itself approximately zero ($+0.0018$, $p = 0.95$). The pollution effect is therefore not driven by differential tree-cover dynamics. Second, GHSL built-up data show no asymmetric urbanisation pattern across cell pollution quartiles (built-up share of grid loss roughly equal across quartiles). Both diagnostics support reading the pollution-biodiversity association as a pollution effect, not a habitat-conversion artifact.

Strain Channel: OOIP Gradient. Table H2 reports the cross-sectional gradient of bird abundance and species richness against $\log(\text{OOIP})$ within 50 km of each cell, interacted with basin indicators,

on the post-2017 Venezuelan eBird sample restricted to cells near at least one field with positive OOIP. All specifications include year-month fixed effects and effort controls (log duration, log distance, number of observers, traveling indicator, hour-of-day quadratic); columns (2) and (5) add observer fixed effects, and columns (3) and (6) replace them with observer×year fixed effects to absorb birder learning (Madhok, 2025). Standard errors are clustered at the cell level. The Orinoco log(OOIP) interaction is negative across all columns and significant on log abundance with observer fixed effects (column 5: -0.054, $p < 0.01$); the Maracaibo interaction is similarly negative (-0.033, $p < 0.05$). Species richness loads with the same sign but is imprecise once observer effects are absorbed, consistent with the extensive margin having already adjusted to chronic exposure.

Table H2: Oil Reserves and Bird Abundance

	Total Richness			log(1+Abundance)		
	(1)	(2)	(3)	(4)	(5)	(6)
log(OOIP) × Orinoco	−0.886***	−0.025	−0.125	−0.102***	−0.053***	−0.046**
	(0.201)	(0.131)	(0.142)	(0.025)	(0.018)	(0.021)
log(OOIP) × Maracaibo	−0.526***	0.088	0.028	−0.169***	−0.033**	−0.035**
	(0.178)	(0.119)	(0.112)	(0.043)	(0.013)	(0.014)
Year-month FE	✓	✓	✓	✓	✓	✓
Observer FE	No	✓	No	No	✓	No
Observer × Year FE	No	No	✓	No	No	✓
Effort controls	✓	✓	✓	✓	✓	✓
Observations	40,481	40,418	40,336	40,481	40,418	40,336

Notes: Dependent variables: total species richness (columns 1–3) and log of 1 plus total individual abundance (columns 4–6). Sample: post-2017 Venezuelan eBird checklists within 50 km of an oil field with positive OOIP. log(OOIP) is the log of cumulative Original Oil In Place (million barrels) of fields within 50 km, included as a control in all specifications. Coefficients report the differential association of log(OOIP) with biodiversity in each basin relative to the rest of Venezuela. Effort controls: log duration, log distance, number of observers, protocol indicator, hour-of-day quadratic. Columns (3) and (6) include observer×year fixed effects to account for birder learning (Madhok, 2025). Standard errors clustered at cell level. ***p<0.01, **p<0.05, *p<0.1.

Maintenance Channel: Species Richness. Counterpart to columns (2)–(4) of Table 7 with $\log(1 +$ richness$)$ as the dependent variable.

Maintenance Channel: Bird Abundance Event Study. Dynamic counterpart to the saturated specification in Table 7: an event study in 6-month bins around 26 November 2022, traced for total log abundance at Chevron-near cells.

Maintenance Channel: Guild-Level Decomposition. Counterpart to the saturated specification in Table 7, with $\log(1 +$ abundance$)$ broken out by trophic guild and using the Chevron-vs-Western-JV $\pm$24-month window of equation (8).

Table H3: Maintenance Channel: Per-Barrel OPEX and Bird Species Richness

	log(1 + richness)		
	(1)	(2)	(3)
$\log(1 + \text{OPEX/prod})$ (η_0)	+0.0138	-0.1279	-0.1736
	(0.1291)	(0.1015)	(0.1831)
$\times$ Post2017 (η_1)	+0.0329	+0.1644*	+0.2007
	(0.1217)	(0.0965)	(0.1801)
$\times$ Post2022 (η_2)	-0.0689*	-0.0357	-0.1050
	(0.0366)	(0.0314)	(0.0720)
$\times$ Chevron (η_3)	-0.3118***	-0.2801	+0.0628
	(0.1130)	(0.2082)	(0.1828)
$\times$ Post2022 $\times$ Chevron (η_5)	+0.2037	+0.1022	+0.0070
	(0.2317)	(0.2744)	(0.5012)
Post2022 $\times$ Chevron (δ_2)	-0.2755	+0.1015	+0.1648
	(0.6003)	(0.6600)	(1.2210)
Cell FE	✓	✓	✓
Year-Month FE	✓	✓	✓
Observer FE	No	✓	—
Observer $\times$ Year FE	No	No	✓
Observations	7,565	7,522	7,457

Notes: Counterpart to Table 7 with log(1 + checklist species richness) as the dependent variable. Same right-hand side, sample, and clustering as columns (2)–(4) of the main-text maintenance table. ***p<0.01, **p<0.05, *p<0.1. Source: Rystad Energy UCube; World Bank VIIRS Global Gas Flaring Tracker; eBird.

Pollution–Effort Selection at $R = 50$ km. Table H5 tests whether satellite-measured pollution predicts eBird effort or coverage outcomes at the radius where the main CH_4–richness specification is most significant. Six standardised effort outcomes are regressed on standardised $\log(1 + CH_4)$ and $\log(1 + \text{VIIRS flare})$ at $R = 50$ km, with cell and year-month fixed effects and cell-clustered standard errors. All twelve coefficients are individually insignificant (largest $|t| = 1.22$), confirming that the richness results at this radius are not artefacts of differential observer behaviour.

Table H4: Chevron Licence and Bird Abundance by Trophic Guild

	NearChev × Post: β_g on $\log(1 + \text{abund}_g)$		
	Cell+YM	Cell+YM+Obs	Cell+YM+Obs×Yr
	(1)	(2)	(3)
Forest insectivores	+0.6041***	+0.5085***	+0.0373
	(0.1771)	(0.1720)	(0.3081)
Frugivores	+0.2863	+0.2031	+0.3470
	(0.2422)	(0.2187)	(0.3276)
Granivores	+0.2596	+0.2529	+0.6338
	(0.1951)	(0.2170)	(0.3916)
Waterbirds	-0.2643	-0.1808	+0.0544
	(0.2187)	(0.1524)	(0.3708)
Cell + Year-Month FE	✓	✓	✓
Observer FE	No	✓	—
Observer × Year FE	No	No	✓
Observations	2,584	2,579	2,573

Notes: Guild-level Chevron-license event around 26 November 2022, using the specification of equation (8): $\log(1 + \text{abund}_g)_{ict} = \alpha_i + \tau_t + \beta_g\,(\text{NearChev}_i \cdot \text{Post}_t) + \mathbf{X}'_{ict}\boldsymbol{\delta} + \varepsilon_{ict}$. Sample: cells within 30 km of one of the four OFAC GL-41 Chevron joint-ventures (PetroBoscan, PetroIndependencia, PetroIndependiente, PetroPiar) or another Western JV, observed within ±24 months of the licence date. Treatment is NearChev × Post; the reference group is Western-JV cells without a Chevron field nearby. All columns include cell and year-month fixed effects and checklist-level effort controls; columns (2)–(3) add observer fixed effects; column (3) uses observer × year fixed effects. Standard errors clustered at the cell level. ***p<0.01, **p<0.05, *p<0.1. Source: eBird; EltonTraits; Rystad Energy UCube.

Table H5: Selection Diagnostic: Pollution Exposure and Birding Effort ($R = 50$ km)

	log(1 + CH_4)		log(1 + VIIRS flare)	
Dependent variable	Coef.	(SE)	Coef.	(SE)
log(1 + # checklists)	−0.016	(0.015)	−0.021	(0.031)
Mean log(duration)	−0.007	(0.018)	−0.011	(0.030)
Mean distance (km)	0.018	(0.016)	−0.017	(0.032)
Mean # observers	0.020	(0.017)	0.008	(0.035)
Traveling share	−0.022	(0.018)	−0.045	(0.038)
Mean start hour	0.000	(0.020)	0.017	(0.037)
Cell FE		✓		
Year-month FE		✓		
SE clustering		Cell		
N (cell-months)		7,239		

Notes: This table tests whether pollution exposure predicts birding effort or coverage, as a selection diagnostic. Each row reports a separate regression of a standardised (z-scored) effort or coverage outcome on standardised log(1 + CH_4) and log(1 + VIIRS flare) at $R = 50$ km. Sample: post-period Venezuelan cell-months with non-missing TROPOMI methane data (2018–2024). All specifications include cell and year-month fixed effects. Coefficients are in standard-deviation units: a one-standard-deviation increase in the exposure is associated with a β-standard-deviation change in the outcome. No coefficient is significant at conventional levels, indicating that methane and flaring do not predict birding effort or coverage at this radius. Standard errors clustered at the cell level. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$. Source: eBird; Sentinel-5P TROPOMI; NASA FIRMS VIIRS.