Where Do Oil Windfalls Reach?

Local Transmission, State Capacity, and Ethnic Exclusion

Albert Duodu

Aalto University, Department of Economics

albert.duodu@aalto.fi

September, 2025

Abstract

Where do the local effects of a resource boom land? Using a global panel of half-degree grid cells across 195 countries from 1989 to 2014, with within-country identification drawn from the 74 countries containing both petroleum and non-petroleum cells, I compare petroleum and non-petroleum cells within the same country and year, with identification validated by high-frequency oil supply news shocks around OPEC announcements. Petroleum cells brighten by roughly eight percent over the 1998–2008 boom. Aggregate organised conflict shows no systematic escalation, but disaggregated UCDP and ACLED outcomes reveal a modest reallocation toward lower-intensity, non-state and civilian-targeted violence rather than a flat null. The development response is moderated on two independent pre-sample institutional margins, state capacity and ethnic political exclusion: cells in high-capacity countries with politically included populations exhibit the full response, cells where either margin is unfavourable show sharp attenuation, and cells where both bind show no detectable response. The positive cross-country correlation between oil exposure and conflict is recovered in this data only when within-country variation is removed; the within-country local response and the cross-country association are distinct empirical objects. The institutional moderation is consistent with mechanisms in which state capacity and political access shape how local revenues reach resident populations; the design recovers reduced-form differentials in local economic activity rather than measuring transfers, public spending, or household welfare directly.

Keywords: OIL PRICE SHOCKS, ARMED CONFLICT, STATE CAPACITY, RESOURCE CURSE, LOCAL DEVELOPMENT, PRIO-GRID

JEL Codes: D74, O13, Q34, H10, F51

1 Introduction

The political economy of oil windfalls turns on a fundamental tension. A positive price shock raises local incomes and expands the fiscal resources available to the state, but it also raises the value of controlling the territory that generates those revenues. These forces run in opposite directions: the first pacifies, through higher opportunity costs of conflict and greater state coercive capacity; the second destabilises, through a larger prize for armed challengers and through distributional contests the state may be unable to manage (Robinson et al., 2006; Besley and Persson, 2010; Dal Bó and Dal Bó, 2011). Which force dominates is an empirical question.

A large empirical literature studies this relationship at the country level. Early contributions find that countries with a higher share of primary commodity exports face a greater risk of civil war (Fearon, 2005; Collier and Hoeffler, 2004). Cotet and Tsui (2013) subsequently show that adding country fixed effects eliminates the oil-reserves correlation with civil-war onset, an early signal that the cross-country pattern reflects between-country differences rather than the within-country effect of resource wealth. Country-level commodity-shock studies have produced mixed findings (Brückner and Ciccone, 2010; Bazzi and Blattman, 2014), and the link between natural resources and conflict remains widely debated (O'Brochta, 2019; Vesco et al., 2020; Bitoto, 2026). A separate sub-national literature uses within-country variation in single-country settings (Dube and Vargas, 2013; Maystadt and Ecker, 2014) and in multi-country regional designs that pool cells across one continent (Berman et al., 2017; Morelli and Rohner, 2015; Hunziker and Cederman, 2017; Adhvaryu et al., 2021), but does not provide a global cell-level test that nests both development and conflict outcomes.

This paper provides a global cell-level test of the local resource-curse hypothesis. It does not adjudicate the country-level resource-curse correlation, which is a different empirical object from the within-country local response; the design absorbs country-level confounds rather than measuring them. What the paper documents is that local petroleum exposure under common national conditions does not behave the way the country-level correlation would suggest: the within-country local response is not conflict-inducing on average. The positive cross-country correlation between oil exposure and conflict reported in earlier work is recovered in this same data only when within-country variation is removed; that pattern is consistent with the cross-country correlation reflecting between-country confounds rather than the local effect of oil exposure, but the cell-level design cannot demonstrate it.

This paper provides a global within-country cell-level joint test of the local ef-

fects of oil-price exposure on economic development and armed conflict. I assemble a panel of 64,818 half-degree grid cells spanning 195 countries from 1989 to 2014, drawn from the Peace Research Institute Oslo (PRIO) Grid (Tollefsen et al., 2012), and combine these with georeferenced conflict events from the Uppsala Conflict Data Program (UCDP GED), satellite night-lights, global oil prices, and country-year measures of institutional quality. Within-country identification is drawn from the 74 countries with both petroleum and non-petroleum cells; the country-count bookkeeping across sample stages is reported in Table 1.

Identifying the causal effects of oil windfalls presents two central challenges. First, reverse causality: conflict and political instability may affect oil production and prices rather than the reverse.[1] Second, omitted variables: both oil prices and conflict may respond to global economic conditions or geopolitical shocks. These challenges are particularly acute in cross-country settings, where resource endowments and institutional quality are jointly determined.

To overcome these challenges, I exploit two sources of variation, one spatial and one temporal, both plausibly exogenous to local conditions. The spatial variation arises from geology: only locations with petroleum reserves are directly exposed to oil windfalls. I construct a predetermined measure of local petroleum exposure based on the spatial distribution of oil fields. While not all petroleum cells are continuously active, oil-related activity and revenues are disproportionately concentrated in these locations. The temporal variation comes from global oil price movements. A potential concern is that conflict in oil-producing areas may itself affect oil prices through disruptions to supply. This concern is limited by the geography of the design. World oil prices are determined in integrated global markets and are insensitive to shocks at this spatial scale; no individual cell can meaningfully influence world supply. The asymmetry — global prices affecting local outcomes but not vice versa — is the foundation of the design.

The identifying variation is therefore the differential exposure of petroleum cells relative to non-petroleum cells under common national conditions. This is the analytical move the cross-country resource-curse literature has been missing. Country-level studies cannot separate the local effect of oil from country-level confounds — institutional quality, ethnic composition, regime type, geopolitical position. Comparing

[1] When the United States killed Iranian General Qasem Soleimani in January 2020, crude oil prices jumped by roughly 4 percent within hours (Atlantic Council, 2020; Reuters, 2020). When Houthi forces began attacking commercial shipping in the Red Sea in late 2023—threatening a corridor through which roughly 10 percent of global seaborne oil trade passes—energy markets reacted immediately (U.S. Energy Information Administration, 2024; World Bank, 2024). More recently, escalating conflict involving Iran disrupted flows through the Strait of Hormuz, through which around one-fifth of global oil supply transits, pushing prices sharply upward (Reuters, 2026a,b).

0.5-degree grid cells (roughly 50 km $\times$ 50 km) within the same country and year strips all of these out by construction. By including cell fixed effects, I absorb time-invariant local characteristics; additional country-year fixed effects also absorb every macroeconomic and political shock common to a country in a given year, including national oil revenues, regime changes, and aggregate conflict cycles. I validate the results with high-frequency oil supply news shocks constructed from changes in oil futures prices around pre-scheduled OPEC announcements (Känzig, 2021), which are uncorrelated with local demand by construction. [2] The same comparison reappears as a country-level positive correlation between oil and conflict only when within-country variation is removed, consistent with — though not by itself proof of — the view that the cross-country pattern reflects between-country confounds rather than the local effect of oil exposure identified here.

Three findings emerge. First, oil-price exposure raises local economic activity in petroleum-bearing cells: night-light intensity in petroleum cells rises by roughly eight percent over the 1998–2008 price cycle relative to non-petroleum cells in the same country and year, an estimate stable across alternative controls, weather, lagged price, and identification using OPEC supply news. Second, the same exposure does not raise local conflict in the average petroleum cell on aggregate measures, but it does produce a modest compositional reallocation toward lower-intensity violence. The preferred specification, which excludes ring-adjacent non-petroleum cells that themselves receive conflict-reducing fiscal spillovers, returns small negative point estimates on both events and deaths, and the negative direction does not survive consistently across alternative specifications (supply-shock identification, ACLED, PPML on raw counts). Disaggregating violence by actor type reveals a substantively informative pattern: state-based and large-scale organised violence are unaffected, while non-state violence and violence against civilians show small but statistically detectable positive responses in both UCDP and ACLED. The aggregate null is therefore most defensibly read as no escalation of organised violence accompanied by a small reallocation toward lower-intensity, non-state and civilian-targeted forms, rather than as a flat absence of conflict response or as a quantified peace dividend. Third, the development response is moderated on two independent institutional margins. Both pre-sample state capacity and pre-sample ethnic political exclusion attenuate the local response. In a horse-race specification, each margin pulls the development response down independently, and petroleum cells in low-capacity countries hosting politically excluded groups exhibit essentially none

[2] I also use the Kilian (2008) OPEC production shortfall series as a robustness check. Because these shocks reflect realized disruptions, I treat them as corroborative evidence.

of the response that accrues to high-capacity included cells. The local response is concentrated in cells where neither institutional margin binds.

The pattern is consistent with a fiscal-distribution channel. The local within-country-year response to oil-price exposure is large where the state can convert revenues into public goods (high ICRG capacity) and where the local population has formal political access to those revenues (no EPR-coded ethnic exclusion); the response is statistically indistinguishable from zero where either margin fails. The local conflict null is consistent with offsetting forces operating at the cell scale — a rent-contestation channel that would raise violence in petroleum cells under canonical rent-seeking models, and security and opportunity-cost channels that work in the opposite direction. The data cannot identify the two channels separately at this level of aggregation, and the null should be read as "no systematic conflict escalation" rather than as a quantified peace effect. The within-country design also speaks to the apparent tension with the Colombian estimates of Dube and Vargas (2013): Colombia is one country in the low-capacity, ethnically-divided tail of the institutional distribution, and the global within-country mean reported here is dominated numerically by cells in different institutional environments. A formal model of the dual-margin mechanism that delivers these comparative statics is sketched in Section 2.2.

The paper makes three contributions. *(i)* It provides a global within-country cell-level test of the local resource-curse hypothesis and finds no systematic conflict escalation in the average petroleum cell. The cross-country positive correlation between oil exposure and conflict reported in earlier work (Ross, 2004; Fearon, 2005; Cotet and Tsui, 2013; Lei and Michaels, 2014) is recovered in the same data only when within-country variation is removed; the move from cross-country to within-country identification that I rely on is anticipated by Cotet and Tsui (2013), who show that country fixed effects eliminate the oil-reserves correlation with civil-war onset, and by the commodity-shock literature using country-level price variation (Brückner and Ciccone, 2010; Bazzi and Blattman, 2014). To my knowledge, the cell-level continuous-treatment design that combines a global panel, predetermined geological exposure, and joint estimation of development and conflict responses with country$\times$year fixed effects is new. The contribution is a global within-country local estimand for the joint development-and-conflict response to oil-price exposure — a different empirical object from the country-level association studied in the cross-country literature, and one that is well-identified at this scale of variation. The paper does not adjudicate the cross-country resource-curse correlation; it characterises a different estimand. *(ii)* It documents two independent in-

stitutional gradients in the local development response — state capacity and ethnic political exclusion — in a horse-race specification. The result extends Caselli and Michaels (2013)'s finding that Brazilian municipalities receiving oil windfalls experience large increases in reported public spending but smaller-than-expected improvements in living standards: where Caselli and Michaels show that the spending-to-outcomes mapping is imperfect at the municipal level, this paper shows that the price-shock-to-local-activity response itself varies sharply with administrative capacity and political access at the global cell-year level. The institutional reading also draws on direct oil-and-ethnicity precedents (Morelli and Rohner, 2015; Hunziker and Cederman, 2017), on the ethnic-exclusion-and-rebellion literature initiated by Cederman et al. (2010), on the colonial-partition spatial-discontinuity framework of Michalopoulos and Papaioannou (2016), and on the ethnic-data-linkage contribution of Müller-Crepon et al. (2023). *(iii)* It provides a joint characterisation of development, conflict, and spatial spillovers within a common half-degree grid framework, on the PRIO-GRID platform also used by Girard et al. (2025); Berman et al. (2017) to study artisanal mining.

2 Background

2.1 How Oil Windfalls Reach Local Economies

A world oil price increase does not affect localities uniformly. Its reach is mediated by the geography of petroleum deposits, the structure of extraction contracts, the fiscal institutions that connect the oil sector to local communities, and the political-demographic structure that determines who has access to the state's redistributive machinery. Five channels transmit global price movements to local economies. The first is the *fiscal* channel: production royalties and taxes flow to national governments and are redistributed to subnational units through intergovernmental transfers, public employment, and infrastructure spending in or near production zones (Caselli and Michaels, 2013; Robinson et al., 2006). The second is the *local demand* channel: oil-sector wages and procurement linkages generate multiplier effects for workers and service providers in adjacent communities. The third is the *security* channel: protecting petroleum infrastructure is a state priority, so oil-producing areas receive disproportionate security deployments; in high-capacity states these deployments suppress local disorder and facilitate commerce, but in low-capacity states the same deployments may be absent, ineffective, or themselves a source of predation. The fourth is the *rent-contestation* channel: higher oil prices raise the value of controlling the state apparatus that collects and disburses rents, increasing

the incentive for armed groups to challenge incumbent governments (Robinson et al., 2006; Dal Bó and Dal Bó, 2011). The fifth is the *political-access* channel: in canonical accounts, groups with formal political representation receive a meaningful share of the disbursements that flow through the fiscal channel, while politically excluded groups settled on petroleum territory generate revenues without commensurate local benefits (Cederman et al., 2010; Esteban and Ray, 2008; Michalopoulos and Papaioannou, 2016). The five channels are not mutually exclusive and the cell-level reduced form cannot separately identify them; their joint operation is what the institutional-moderation results below characterise.

The common thread is institutional design. Channels one through three turn on whether the state can deliver — whether it has the administrative capacity to convert revenue into local public goods. Channels four and five turn on who has access to the proceeds — whether organised armed groups can extract them by force, and whether resident populations can extract them through political representation. The local economic impact of an oil windfall therefore depends jointly on the state's ability to deliver public goods and the local population's political access to the rents. This is the dual-margin foundation of the paper's institutional findings.

2.2 A Dual-Margin Model of Local Rent Transmission

The same price shock that raises local incomes also raises the value of controlling the territory that generates them. This duality is the core logic of the resource-curse literature (Sachs and Warner, 1995; Ross, 2004; Mehlum et al., 2006). Its sub-national implications, however, are sharper and more conditional than cross-country regressions can detect. This subsection sets up a compact model that formalises one channel through which oil-price movements map into local economic activity. The model yields comparative statics on institutional moderation of the local response, which the empirical sections test directly. The data identify the reduced-form coefficient β_{gc} and how it varies across cells and institutional contexts; the structural primitives (τ, κ_c, π_g) and the underlying flow of transfers and public spending are not separately observed.

Setup. A country c contains a continuum of cells g. Cell g is petroleum-bearing if $Oil_g = 1$. A petroleum-bearing cell produces a quantity q of oil sold at the world price P_t, generating a pre-tax cell-level rent

$$R_{gt} \;=\; Oil_g \cdot q \cdot P_t.$$

The state collects a fraction $\tau \in (0,1)$ of this rent as tax revenue; the residual $(1-\tau)$ accrues directly to the cell as worker wages and supply-chain spending, of which a share $\rho \in (0,1]$ translates into measurable local economic activity. Of the tax revenue τR_{gt}, a fraction $\kappa_c \in [0,1]$ is converted into local public goods — κ_c indexes the country's *administrative capacity* to deliver public goods, as proxied empirically by the pre-sample ICRG quality-of-government index. Of the public-goods budget, a fraction $\pi_g \in [0,1]$ is allocated to residents of cell g — π_g indexes the local population's *political access*, equal to one for a fully included group and approaching zero for an ethnically excluded settlement, as proxied empirically by the pre-sample EPR political-status indicator.

Total local economic activity in cell g is therefore the sum of a direct-rent component and a public-goods component:

$$Y_{gt} \;=\; \underbrace{(1-\tau)\,\rho\,R_{gt}}_{\text{direct local share}} \;+\; \underbrace{\kappa_c\,\pi_g\,\tau\,R_{gt}}_{\text{state-mediated share}} \quad . \tag{1}$$

Substituting $R_{gt} = Oil_g \cdot q \cdot P_t$ and taking logs gives the empirical reduced form

$$\ln Y_{gt} \;=\; \alpha_g + \lambda_{ct} + \beta_{gc}\left(Oil_g \times \ln P_t\right) + \varepsilon_{gct}, \quad \text{with} \quad \beta_{gc} \;=\; (1-\tau)\,\rho + \kappa_c\,\pi_g\,\tau, \tag{2}$$

where α_g absorbs the cell-level constant ($\ln q$, the direct-share parameters) and λ_{ct} absorbs national time-varying shocks. Equation (2) is exactly the empirical specification used in Section 6; β_{gc} is the cell-specific coefficient that the regression recovers.

Comparative statics. The price-cycle response coefficient β_{gc} is increasing in both institutional moderators:

$$\frac{\partial \beta_{gc}}{\partial \kappa_c} \;=\; \pi_g\,\tau \;\geq\; 0, \qquad \frac{\partial \beta_{gc}}{\partial \pi_g} \;=\; \kappa_c\,\tau \;\geq\; 0, \qquad \frac{\partial^2 \beta_{gc}}{\partial \kappa_c\,\partial \pi_g} \;=\; \tau \;>\; 0.$$

A high-capacity inclusive cell receives the largest response $(1-\tau)\rho + \tau$. A cell in which both margins fail receives only the direct-share floor $(1-\tau)\rho \geq 0$. The positive cross-partial implies that the two moderators are mutual complements: each margin independently attenuates the state-mediated channel, but a cell receives the full response only when both moderators are favourable, and the response collapses toward the floor only when both fail. Empirically, this means cells where only one margin binds should exhibit partial attenuation rather than the full floor — a graded pattern that the regression specification with separate triple interactions on $LowCap_c$ and $Excl_g$ is designed to recover.

Conflict block. Let local conflict C_{gt} respond to two opposing forces, both increasing in the rent R_{gt}: a rent-mobilisation force $\mu(R_{gt})$ raising violence through the prize value of controlling the cell, and a pacification force $\sigma(R_{gt})$ lowering violence through security deployments protecting production sites and through the opportunity-cost effect of higher local wages. The unconditional cell-level response is

$$\frac{\partial C_{gt}}{\partial R_{gt}} = \mu'(R_{gt}) - \sigma'(R_{gt}).$$

The sign of this expression is theoretically ambiguous; whether $\mu' > \sigma'$ or vice versa is empirical. Heterogeneity of C_{gt} by κ_c or π_g is expected only if μ' and σ' respond differentially to the moderators — which the data may not support if both forces scale similarly with R_{gt}.

Testable predictions. The model generates four predictions that the empirical sections test directly.

Prediction 1. Oil-price exposure raises local economic activity on average. Since $\beta_{gc} \geq (1-\tau)\rho \geq 0$ for all (κ_c, π_g), the population-average coefficient $\bar{\beta} = \mathbb{E}[\beta_{gc}]$ is strictly positive provided either the direct-share or the state-mediated channel is active. Tested in Section 6, Table 3.

Prediction 2. The development response is attenuated by low state capacity. Substituting a low-capacity indicator $LowCap_c = \mathbf{1}[\kappa_c < \tilde{\kappa}]$ into equation (2) implies a negative triple-interaction coefficient $\theta = \mathbb{E}[\pi_g]\,\tau\,(\bar{\kappa}_{\text{high}} - \bar{\kappa}_{\text{low}}) < 0$ on $Oil_g \times \ln P_t \times LowCap_c$. Tested in Section 7.1, Table 5.

Prediction 3. The development response is attenuated by ethnic political exclusion. Symmetrically, with $Excl_g = \mathbf{1}[\pi_g < \tilde{\pi}]$, the triple-interaction coefficient γ on $Oil_g \times \ln P_t \times Excl_g$ is negative. The two moderators enter additively in a horse-race specification: each margin pulls the response down independently, and the implied combined response in a cell where both margins bind approaches the floor $(1-\tau)\rho$. Cells where only one margin binds exhibit partial attenuation. Tested in Section 7.2, Table 6.

Prediction 4. The average local conflict response is unsigned and does not exhibit systematic moderation by κ_c or π_g. Since the cell-level response is $\mu' - \sigma'$ and both forces scale with R_{gt}, the sign is theoretically ambiguous and need not be moderated by institutional variables that operate on the rent-distribution side. Tested in

Sections 6 and 7; institutional heterogeneity tested in Table 5 and Table 6.

The model is a qualitative organising device, not a structural-estimation target. It delivers the sign and direction of the comparative statics that the empirical sections then test — both moderators reduce the response on their own and the response approaches zero when both bind — but the precise magnitudes of the off-diagonal cells of the (capacity, inclusion) grid are an empirical question not pinned down by the binary moderators. The strict multiplicative form $\kappa_c \pi_g$ would predict an identical floor whenever either margin fails; empirically, the off-diagonal cells exhibit partial rather than full attenuation, reflecting the fact that $LowCap_c$ and $Excl_g$ are median splits of continuous structural parameters rather than $\{0, 1\}$ switches. The empirical horse-race in equation (3) recovers this richer additive structure as separate $\hat{\theta}$ and $\hat{\gamma}$ coefficients on the two interactions.[3]

What the design identifies and what it does not. The empirical exercise that follows recovers a single reduced-form object: the differential response of cell-level economic activity (night lights) and cell-level conflict (UCDP events and deaths) to global oil-price movements, comparing petroleum cells with non-petroleum cells inside the same country and year. That object is well-identified, and its institutional moderation by state capacity and political access is the central finding of the paper. The design does not observe transfers, public spending, procurement, intergovernmental fiscal flows, household income, or local public-goods provision. The model above sketches one mapping from the data to a rent-distribution mechanism: high $\kappa_c \pi_g$ cells receive a larger share of the state-mediated component, low $\kappa_c \pi_g$ cells receive less. The reduced-form pattern is consistent with that mapping, but the same pattern is also consistent with alternative channels — differential labour-market demand, differential security deployment, differential infrastructure investment, differential local-government tax-base effects — that operate without an explicit transfer of rents from the central state. The moderation results are evidence of *institutionally conditioned local transmission* of oil-price shocks; expressions such as "rents reach the cell" refer to that reduced-form differential rather than to a measured fiscal flow. Disentangling the channels would require direct me-

[3] Equation (2) also assumes τ and ρ are common across countries. In practice both vary: tax shares range from below 30% (US Gulf operators) to above 80% (Saudi Arabia, Norway sovereign funds), and direct-share parameters vary with the labour-intensity of extraction and the depth of local supply-chain linkages. Allowing (τ_c, ρ_c) to vary by country adds a country-level intercept to β_{gc} that is absorbed by the country×year fixed effects in equation (3). The moderator coefficients $\hat{\theta}$ and $\hat{\gamma}$ therefore identify the average institutional moderation across the cross-country distribution of (τ_c, ρ_c), not a single structural parameter — consistent with the reduced-form rather than structural interpretation of the empirical estimates.

diator data (cell-level transfers, public-goods access, or welfare outcomes such as DHS wealth or child-mortality indices).

3 Data and Variable Construction

3.1 Data Sources

The analysis combines geocoded data on economic activity, conflict, and natural-resource endowments at the grid-cell level. The unit of observation is a $0.5° \times 0.5°$ cell-year over 1989–2014. This window is set by data coverage and identification relevance: UCDP GED begins in 1989, DMSP-OLS night-lights begin in 1992, the Känzig (2021) supply-news series covers the full window without truncation, and 2014 is a conservative endpoint before US shale production materially weakened OPEC's pricing power in subsequent years. The window also brackets the largest peacetime commodity price cycle in modern history (Figure 1), which provides the within-country-year variation the regression exploits.

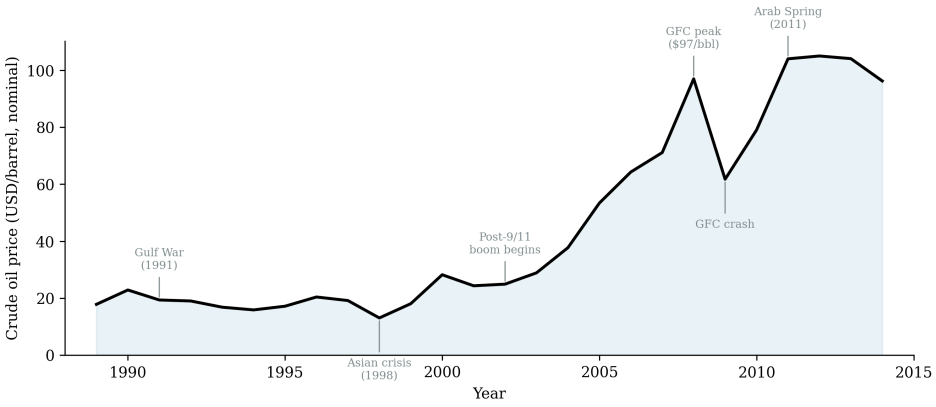

Figure 1: World crude oil price, 1989–2014. Source: World Bank Pink Sheet (nominal USD per barrel). Key events annotated. The shaded band marks the 2002–2008 price boom, the primary source of within-country-year variation exploited in the reduced-form estimates.

The spatial backbone is PRIO-GRID (Tollefsen et al., 2012), which provides a globally consistent grid of cells with harmonized economic, demographic, and geographic variables. From PRIO-GRID, I obtain the treatment indicator for petroleum exposure, local economic outcomes (including night lights), and baseline controls such as population, climate, and geographic characteristics. The grid structure allows comparisons between petroleum and non-petroleum cells within the same country and year.

Conflict outcomes are drawn from the UCDP Georeferenced Event Dataset (GED)

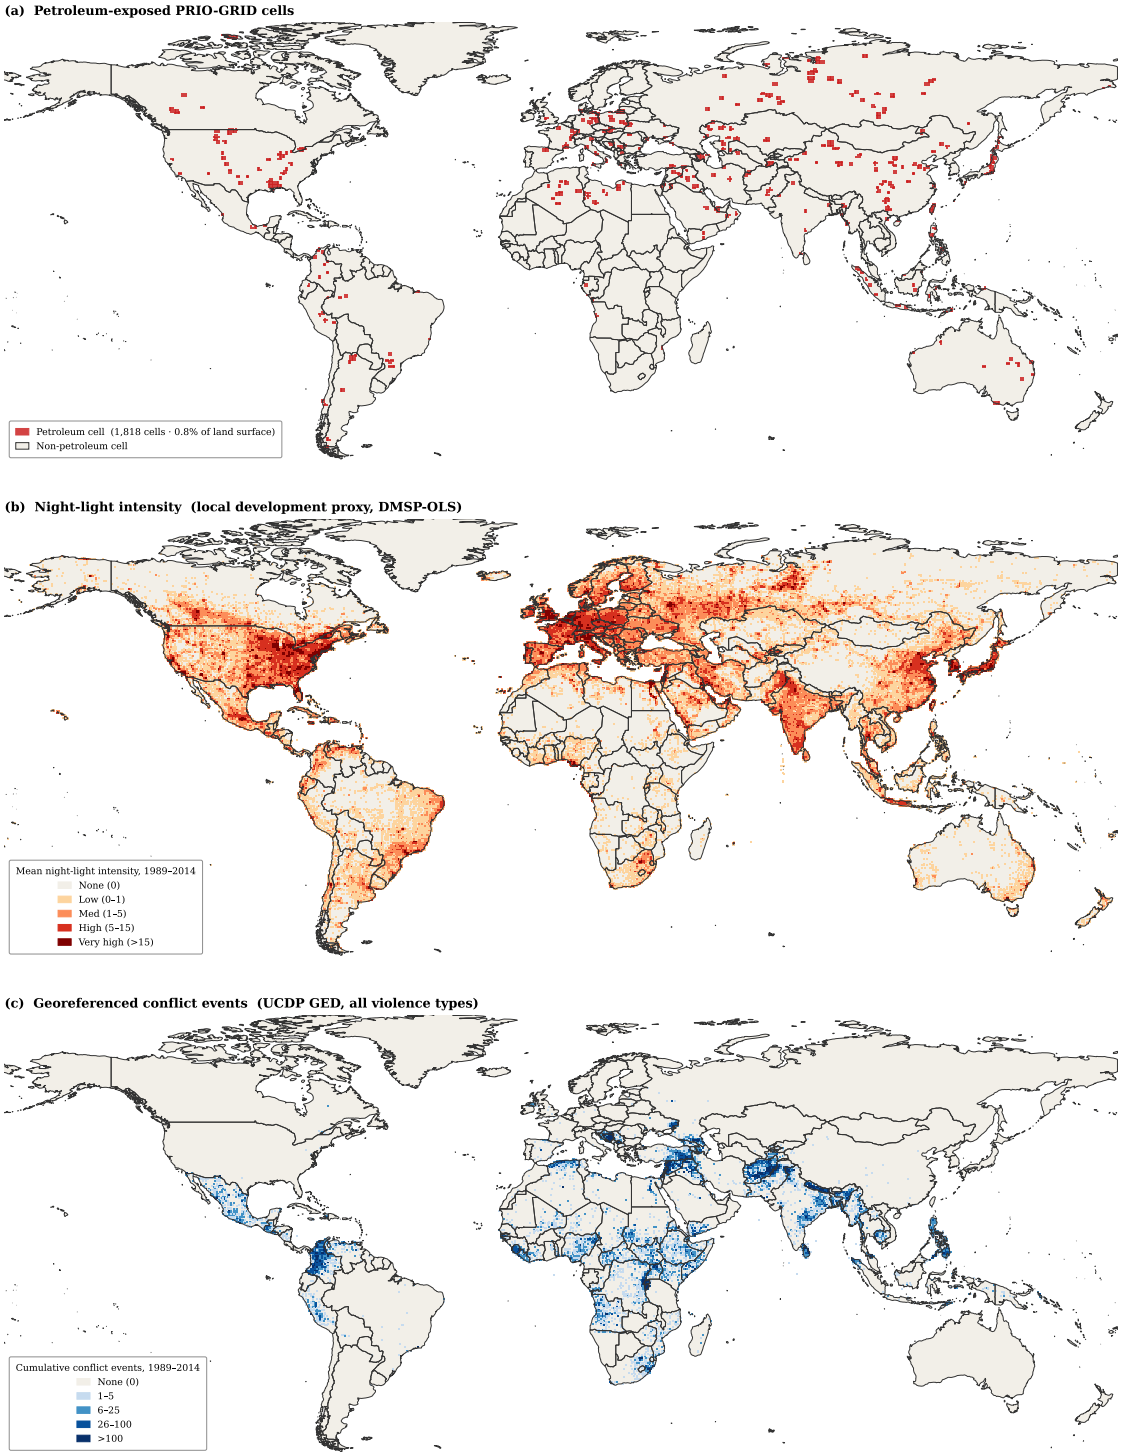

Figure 2: **Petroleum exposure, local development, and conflict across PRIO-GRID cells, 1989–2014.** Panel (a) shows the global distribution of petroleum-exposed 0.5-degree cells (1,818 cells, 2.9% of the land surface) using the Lujala et al. (2007) onshore petroleum deposit dataset embedded in PRIO-GRID. Panel (b) shows mean calibrated night-light intensity (DMSP-OLS) across 1989–2014 as a proxy for local economic activity. Panel (c) shows the cumulative count of georeferenced conflict events (UCDP GED v25.1, all types of violence) over the same period. The spatial unit is the PRIO-GRID cell (Tollefsen et al., 2012).

(Sundberg and Melander, 2013; Davies et al., 2025), which records organized violence events with precise geographic coordinates. Events are assigned to grid cells and aggregated to annual cell-level measures. This yields measures of conflict incidence and intensity that are directly comparable across locations.

Oil prices are measured using the World Bank Pink Sheet. The key explanatory variable combines global price movements with local petroleum exposure, allowing identification from differential responses within countries to common price shocks.

Additional datasets provide supporting variables used in extensions. Institutional characteristics are drawn from the Quality of Government dataset (Teorell et al., 2026) and related sources, while sector-specific governance is captured using the NRGI Resource Governance Index.[4] These variables enter only in heterogeneity analyses and do not affect the baseline identification.

To isolate exogenous variation in oil prices, I use two measures of global supply shocks. The preferred series is the oil supply news measure of Känzig (2021), constructed from high-frequency futures price changes around OPEC announcements.[5] As a complementary source, I use the exogenous OPEC production shortfalls compiled by Kilian (2008).[6] Both series capture shifts in global oil supply that are plausibly orthogonal to local economic conditions. Figure A1 plots the annual Känzig surprise series from 1989 to 2014 alongside the Kilian series over the overlapping 1989–2003 window. The Känzig series spans the full sample and captures major oil market episodes, including the Gulf War, the early-2000s OPEC responses, the 2008 price collapse, and the 2014 supply expansion. This broader coverage is important for identification, as it provides exogenous variation throughout the sample period, including the commodity super-cycle.

3.2 Data Construction and Variables

The panel is constructed by combining PRIO-GRID cells with conflict events and global price series. The master panel covers 64,818 grid cells over 1989–2014. Conflict events

[4] I proxy state capacity by the ICRG quality-of-government composite `icrg_qog` from the Political Risk Services Group's International Country Risk Guide, averaged over 1984–1988 to obtain a country-level pre-sample score. I also use V-Dem electoral-democracy index `vdem_polyarchy` (Coppedge et al., 2025) to proxy for political openness. Natural Resource Governance Institute (NRGI) Resource Governance Index (Natural Resource Governance Institute, 2013) measures *oil-sector-specific* governance quality across 58 resource-rich countries, scoring four sub-components: institutional and legal setting, reporting practices, safeguards and quality controls, and enabling environment

[5] Oil supply news shocks (December 2017 vintage) are available at https://raw.githubusercontent.com/dkaenzig/oilsupplynews/master/oilSupplyNewsShocks_2017M12.xlsx.

[6] I use the Kilian (2008) exogenous OPEC oil production shock series, available at https://sites.google.com/site/lkilian2019/research/data-sets

are assigned to grid cells using geographic coordinates and aggregated to annual cell-level outcomes. Oil prices and global supply shocks are merged by year.

Sample flow. Analytic samples differ across tables for two reasons: outcome-specific coverage gaps (DMSP-OLS night lights begin in 1992 and end in 2013) and design-specific data requirements (ring exclusion, pre-period moderator merges). Table 1 reports the country count at each stage.

Table 1: Country-count bookkeeping across sample stages

Filter	Cells	Countries
Master PRIO-GRID panel, 1989–2014 (0.5° cells)	64,818	196
With valid Gleditsch–Ward country code	64,818	195
of which: countries with at least one petroleum *and* at least one non-petroleum cell (within-country identifying universe)	54,097	74
8-neighbor ring-excluded sample (preferred OLS, Table 3)	61,291	195
Night-lights non-null subsample (DMSP-OLS, 1992–2013)	59,905	193
Moderator subsamples		
ICRG state capacity, 1984–1988 country average	—	122
V-Dem polyarchy, 1970–1988 country average	—	143
NRGI Resource Governance Index, 2013 composite (oil-sector specific)	—	58
ICRG ∩ EPR exclusion (dual-margin subsample, Table 6)	—	121

Notes: The master panel is the PRIO-GRID v3 distribution of Tollefsen et al. (2012). Dropping the 15 cells with missing Gleditsch–Ward code leaves the country-year fixed-effect universe of 195 countries. Within-country identification of $\hat{\beta}$ on $Oil_g \times \ln P_t$ is provided by the 74 countries that contain both petroleum and non-petroleum cells: in countries with no petroleum cells the interaction is identically zero (absorbed by the fixed effects), and in countries with only petroleum cells there is no within-country control group. The SUTVA correction in Section 6 drops non-petroleum cells sharing a face or corner with a petroleum cell; the country count is unchanged because the dropped cells fall within identifying countries. The night-lights subsample reflects DMSP-OLS coverage (1992–2013) and loses Iceland and Greenland. The moderator subsamples reflect availability of the relevant pre-period data: NRGI by construction covers 58 resource-rich countries; ICRG covers 122; V-Dem 143; the ICRG ∩ EPR intersection used in the dual-margin specification is 121.

The master panel contains 195 countries; 74 of these contain both petroleum and non-petroleum cells and contribute identifying variation to $\hat{\beta}$ on $Oil_g \times \ln P_t$. Cluster counts in individual tables (e.g., 193 in Table A2, 122 in Table 5) correspond to the moderator and outcome restrictions reported above.

The main regressor is local exposure to oil price movements, defined as $Oil_g \times \log P_t$, where Oil_g is an indicator for petroleum presence in cell g, and P_t is the global oil price. Petroleum exposure is time-invariant and predetermined, reflecting geological endowments rather than economic activity. Identification therefore comes from differential responses of exposed and unexposed cells to common price shocks within country-

years.

Economic activity is measured using night lights, while conflict outcomes are constructed from geocoded event data aggregated to the cell-year level. The baseline specifications use levels or log transformations that retain zero observations and limit the influence of extreme values. The empirical results are robust to alternative transformations.

To study variation in treatment effects, I construct indicators for institutional and local characteristics that are fixed or predetermined relative to the sample period. Country-level measures of state capacity, political institutions, and income are averaged over pre-sample periods, ensuring that they are not endogenously affected by oil price movements. Local measures, such as political exclusion or geographic exposure, are similarly defined using pre-sample information or time-invariant features.

The PRIO-GRID framework provides fixed geographic units that are independent of administrative boundaries, allowing comparisons across locations within countries under a common spatial structure. By linking these units to geocoded conflict events, local economic activity, geological endowments, and global price variation, the data permit a direct observation of how resource shocks affect nearby localities. This structure is central to identification. Combining fixed geographic exposure with global price movements and high-resolution outcomes allows a clean separation between local and national effects. In contrast to country-level analyses—where heterogeneous local responses are averaged out—the grid-cell design captures how resource shocks propagate within countries across space.

4 Descriptive Statistics and Sample Characteristics

4.1 Who Are the Oil Cells?

Before turning to price variation, it is useful to characterize how petroleum-exposed cells differ from the rest of the sample in levels. If oil deposits systematically locate in richer or more conflict-prone areas, naive correlations between oil prices and local outcomes would confound revenue effects with the endogenous geography of extraction.

Table 2 shows that petroleum cells are economically more active on every measure: night lights are 2.7 times higher, the local GDP proxy is 3.9 times higher, and population is 2.6 times higher than in non-petroleum cells. They are also more conflict-prone: the annual conflict incidence rate (the share of cell-years recording at least one event) is 2.77 percent in oil cells and 1.45 percent in non-oil cells.

Table 2: Descriptive Statistics: Petroleum and Non-Petroleum Cells, 1989–2014

	Non-oil cells		Oil cells		Difference	
	Mean	SD	Mean	SD	Oil − Non-oil	p-value
Night lights (mean)	1.291	3.758	3.439	6.272	2.148***	0.000
Local GDP proxy (mean)	0.628	3.488	2.434	12.811	1.806***	0.000
Conflict events per year	0.1119	6.2236	0.3166	8.9006	0.2047***	0.000
Conflict deaths per year	1.4154	259.3869	2.9179	152.7458	1.5025**	0.040
State-based events per year	0.0797	5.7589	0.2583	8.1078	0.1786***	0.000
Population	89132	345345	232627	691272	143495***	0.000
Distance to capital (km)	1800.4	1625.3	1322.1	1180.8	-478.4***	0.000
Mean temperature (°C)	10.08	14.10	11.94	10.72	1.86***	0.000
Precipitation (mm/yr)	185.7	190.7	200.6	199.0	14.9***	0.000
Cells	63,000		1,818			
Cell-years	1,638,000		47,268			

Notes: Unit of observation is a PRIO-GRID cell-year, 1989–2014. Oil cells are cells where `petroleum_s` > 0 in the PRIO-GRID static file. All conflict outcomes are from UCDP GED v25.1; development outcomes from PRIO-GRID yearly file. The p-value is from a two-sided Welch t-test of equality of means across groups. ***, **, * denote significance at the 1, 5, and 10 percent levels.

Conflict outcomes display a similar pattern. Petroleum cells are more likely to experience conflict, with a higher incidence of events at the cell-year level.[7] These differences highlight the central identification challenge. Petroleum and non-petroleum cells are not comparable in levels, and simple cross-sectional comparisons would attribute pre-existing differences to oil shocks. The empirical design therefore exploits within-cell variation over time, comparing the same locations as global oil prices change, while controlling for country-year shocks common to all cells within a country.

4.2 Within-Country Identification: A Descriptive Preview

The paper's identifying restriction is that, conditional on cell and country×year fixed effects, $Oil_g \times \ln P_t$ is uncorrelated with the residual. The design is continuous-treatment: the world oil price varies every year of the sample ($13/barrel in 1998, $24 in 2001, $97 in 2008, $50 in 2009), so no single year demarcates "pre-treatment" from "treatment." Two descriptive observations preview this comparison and motivate the empirical specification used in Sections 5–6.

[7] One cross-group anomaly is worth flagging: the standard deviation of conflict deaths per cell-year is larger in non-oil cells (259) than in oil cells (153), even though the oil-cell mean (2.92) exceeds the non-oil-cell mean (1.42). This reflects a small set of catastrophic-event cell-years in non-petroleum settings (notably the 1994 Rwandan genocide cells, the 2003–2004 Darfur cells, and the 1999 East Timor cells), each of which contributes individual cell-year death counts in the tens of thousands.

The within-country-year response in one figure. Figure 3 previews the paper's empirical core in three binscatter panels. Each panel partials out cell and country×year fixed effects from both the outcome and the treatment $Oil_g \times \ln P_t$, then bins residualised cell-years into 25 equal-quantile bins and plots the mean residualised outcome against the mean residualised treatment. Panel (a) shows the development response: night-light intensity rises sharply with $Oil_g \times \ln P_t$ once national-level shocks and time-invariant cell heterogeneity are absorbed (binscatter slope $\hat{\beta} = +0.041$). Panel (b) shows the conflict response: the corresponding slope for log-events is small and slightly negative (-0.004), with no visible upward tilt. Panel (c) repeats Panel (a) for two subsamples defined by pre-1989 ICRG state capacity: the high-capacity subsample exhibits a steep positive relationship (slope $+0.052$), while the low-capacity subsample is nearly flat ($+0.014$). The three panels together summarise the paper's three empirical claims — a positive development response, no systematic conflict escalation, and institutional moderation of the development response — in a single within-country-year picture. The within-cell-deviation analogue (cell fixed effects only) is reported in Appendix Figure A3; the year-by-year coefficients under the baseline FE structure are in Appendix E.2.

Figure 3: **Within-country-year response to oil-price exposure: development, conflict, and institutional moderation.**

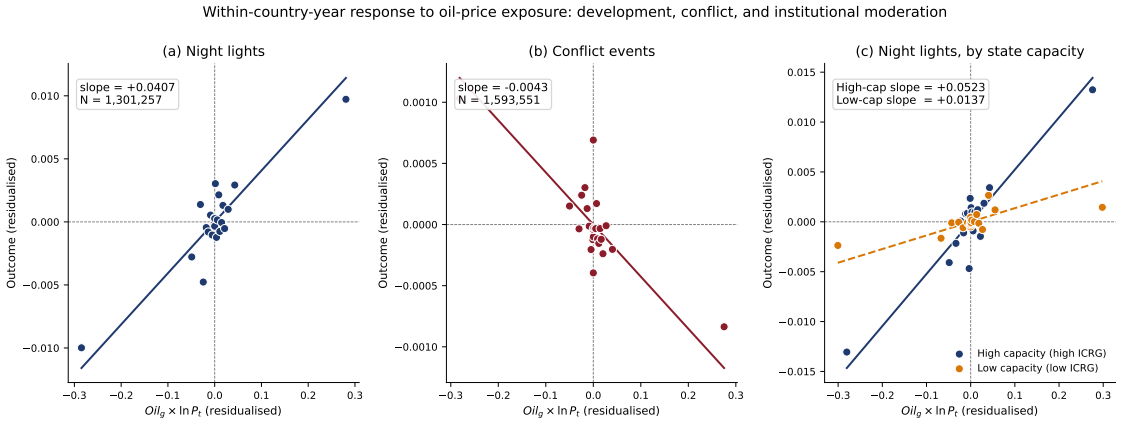

Note: Binscatters of residualised outcome on residualised $Oil_g \times \ln P_t$; 25 equal-quantile bins. Cell and country×year fixed effects partialled out of both axes via alternating projections. Panel (a): $\ln(1 + \text{nlights})$, full ring-excluded sample. Panel (b): $\ln(1 + \text{events})$, full ring-excluded sample. Panel (c): $\ln(1 + \text{nlights})$ split by pre-1989 ICRG state capacity (country-level median split on the 1984–1988 ICRG average). Ring-adjacent non-petroleum cells excluded throughout. The displayed slopes are reduced-form descriptive statistics on the binscatter; the corresponding regression coefficients with cell-clustered SE are in Tables 3 and 5.

Cross-country versus within-country. The within-country contrast that the regression isolates is the only contrast in which the canonical resource-curse correlation does *not*

appear. Figure 4 makes this precise.

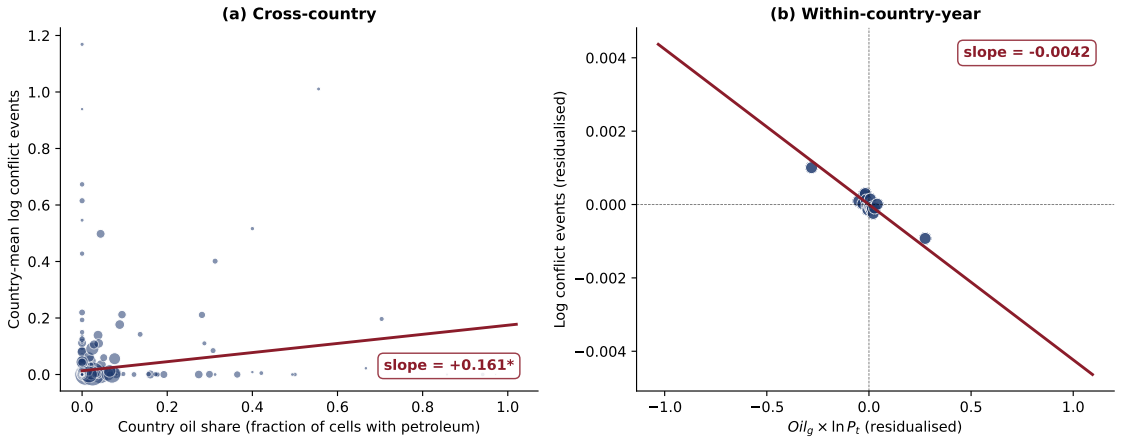

Figure 4: Cross-country versus within-country oil-conflict correlation. Panel (a) plots country-level oil share (fraction of cells with petroleum) against country mean $\ln(1 + \text{events})$ for 175 countries, weighted by panel observations; the fitted slope is $+0.161^*$. Panel (b) is a bin-scatter (25 quantile bins) of cell-year residuals: $\ln(1 + \text{events})$ on $Oil_g \times \ln P_t$ after partialling out cell and country×year fixed effects; the fitted slope is -0.004. Standard errors clustered at the PRIO-GRID cell level in Panel (b). The figure visualises the paper's central methodological move: the cross-country resource-curse correlation reverses sign once within-country variation is the identifying source.

Panel (a) plots country-level oil share against country-mean $\ln(1 + \text{events})$ for 175 countries. The relationship is positive: countries with greater oil exposure are more conflict-prone on average. Panel (b) is the within-country-year analog — a bin-scatter of $\ln(1 + \text{events})$ on $Oil_g \times \ln P_t$ after partialling out cell and country×year fixed effects. The within-country slope is small and slightly negative. The same data, two design choices, opposite signs. The cross-country correlation attributes to oil every difference across countries — institutions, political structure, conflict history. The within-country comparison absorbs those differences by construction and isolates the local response to oil-price movements under common national conditions. The remainder of the paper estimates that within-country response and characterises its institutional heterogeneity.

5 Empirical Strategy

The estimating equation maps the model in Section 2.2 to the data:

$$Y_{gct} \;=\; \beta\left(Oil_g \times \ln P_t\right) \;+\; \alpha_g \;+\; \lambda_{ct} \;+\; \varepsilon_{gct}, \tag{3}$$

where Oil_g is the time-invariant petroleum-deposit indicator (PRIO-GRID `petroleum_s`), $\ln P_t$ is the log world crude oil price, α_g is a cell fixed effect, and λ_{ct} is a country×year fixed effect. The identifying assumption is that, conditional on α_g and λ_{ct}, the interaction $Oil_g \times \ln P_t$ is uncorrelated with the residual ε_{gct}. Equivalently, after stripping

out time-invariant cell heterogeneity and every shock common within a country-year, petroleum and non-petroleum cells in the same country would have followed the same expected residual time path. The descriptive co-movement in Section 4.2 and the year-by-year coefficients in Appendix E.2 bear on this assumption without proving it; the supply-shock validation in Section 9 addresses the demand-side and reverse-causality concerns that remain. Under this assumption, the coefficient β recovers the average local response in equation (2), identified from the differential movement of petroleum versus non-petroleum cells in the same country and year. Cell fixed effects absorb every time-invariant cell characteristic — including the level differences in Table 2. Country$\times$year fixed effects absorb every national-level shock — aggregate income, fiscal policy, regime change, conflict cycles, and the national component of oil-revenue redistribution. Standard errors are clustered at the PRIO-GRID cell level.[8]

The two identification threats this design addresses — and the two it does not absorb mechanically — deserve brief explicit statement. *Demand-side confounding* (global oil prices reflect world economic activity) is partially absorbed by λ_{ct}: the national component drops out, and the residual cell-level demand channel within a country-year is plausibly small at the 0.5 scale. *Reverse causality* (conflict in oil cells disrupts world supply) is implausible at the cell scale: a single 50×50 km cell accounts for a negligible share of world supply, and λ_{ct} additionally absorbs national supply disruptions common within a country-year. Both residual concerns are addressed directly in Section 9 using the Känzig (2021) and Kilian (2008) supply-shock instruments, which restrict identifying variation to OPEC announcement and production disruptions orthogonal to global demand and to any sample cell's conflict trajectory.

What $\hat{\beta}$ identifies. The country-year fixed effect λ_{ct} absorbs every contemporaneous national-level adjustment to the oil price — federal fiscal transfers, national public-employment programs, sovereign-fund saving rules, currency adjustments, and aggregate political responses. The coefficient $\hat{\beta}$ therefore identifies the *differential* response of petroleum cells relative to non-petroleum cells within the same country and year, not the total local effect of an oil-price increase. The total fiscal-transmission effect of oil wealth on a producing country's petroleum cells is the sum of the absorbed national component (in λ_{ct}) and the within-country-year differential identified here. This is the right object for testing whether oil cells benefit *more* than the rest of the country, and for

[8] Year-clustered and two-way (cell$\times$year) clustered alternatives are reported in Appendix C. All outcomes are transformed as $\ln(1+x)$; the arcsinh and PPML alternatives in Appendix B.2 give substantively similar conclusions.

separating local channels (security deployments, supply-chain linkages) from national channels (federal transfers, currency effects); it is not the right object if the question is the total welfare effect of an oil windfall on a producing country.

6 Main Results

Table 3 estimates the effect of oil price shocks on local economic activity and conflict, exploiting within-country variation in time-invariant geological exposure interacted with global oil prices.

A central identification challenge is partial interference. In the full-sample specification, non-petroleum cells include locations adjacent to petroleum zones that are themselves exposed to spillovers during oil booms. These nearby cells experience increases in economic activity and reductions in conflict (Appendix E.1), and therefore do not represent untreated counterfactuals. Including them in the control group biases the estimated conflict effect toward zero. To address this, we exclude ring-adjacent cells and restrict the comparison to geographically remote non-petroleum areas. The remaining specifications add controls for competing commodity exposure and weather, with the preferred specification also using lagged oil prices to mitigate concerns about within-year reverse causality. Two findings emerge. First, oil-price exposure raises within-country-year night-light intensity in petroleum-bearing cells. The preferred specification (Table 3, column 5) implies a within-country-year differential of approximately eight log-points in $\ln(1 + \text{nlights})$ over the 1998–2008 price cycle — equivalent to roughly an eight-percent increase in night-light intensity in cells where lights are positive. The estimate is stable across all specification refinements, indicating that the development response is not driven by control-group composition, correlated commodity cycles, or weather shocks.

Second, the conflict response depends on the definition of the control group and on the identification strategy. In the full sample, estimates are small and statistically indistinguishable from zero, reflecting attenuation from partially treated ring-adjacent control cells (Appendix E.1). Once those cells are excluded, the preferred OLS specifications yield small negative point estimates whose 95% confidence intervals exclude economically meaningful escalation but do not separate a precise negative effect from zero. Two alternative specifications soften this further: the Känzig (2021) supply-news reduced form returns conflict coefficients indistinguishable from zero, and the ACLED replication for 1997–2014 (Section 9.4) returns aggregate nulls with small positive coef-

Table 3: Oil Windfalls, Local Development, and Conflict: Main Results

	Geological exposure × world price				
	(1) *Full sample*	**(2)** *Ring-excl.*	**(3)** *+Competitors*	**(4)** *+Weather*	**(5)** *Lagged p*
Night lights	$+0.0367^{***}$	$+0.0395^{***}$	$+0.0395^{***}$	$+0.0396^{***}$	$+0.0393^{***}$
	(0.0040)	(0.0041)	(0.0041)	(0.0041)	(0.0042)
Conflict events	-0.0045	-0.0060^{**}	-0.0060^{**}	-0.0061^{**}	-0.0066^{**}
	(0.0028)	(0.0028)	(0.0028)	(0.0028)	(0.0030)
Conflict deaths	-0.0064	-0.0080^{*}	-0.0080^{*}	-0.0081^{*}	-0.0084^{*}
	(0.0045)	(0.0044)	(0.0044)	(0.0044)	(0.0046)
Cell FE	✓	✓	✓	✓	✓
Country×Year FE	✓	✓	✓	✓	✓
Ring cells excluded		✓	✓	✓	✓
Gold×price control			✓	✓	✓
Weather (temp, precip.)				✓	✓
Price lag (p_{t-1})					✓
Common-sample N	1,333,797	1,258,907	1,258,907	1,258,907	1,258,907

Notes: OLS estimates. The regressor is $Oil_g \times \ln P_t$ in columns (1)–(4) and $Oil_g \times \ln P_{t-1}$ in column (5). Identification exploits the interaction of time-invariant geological petroleum exposure (Oil_g) with the global oil price, which is exogenous to any individual 0.5° cell. Country×year fixed effects absorb all national-level demand shocks. Column (1) uses all non-petroleum cells as controls. Column (2) drops ring-adjacent non-petroleum cells (those sharing a border or corner with a petroleum cell) from the control group, addressing a SUTVA violation: ring cells experience conflict-reducing spillovers during oil booms ($\hat{\delta} \approx -0.007^{***}$, Table A8), which partially treats the control group and biases the conflict coefficient toward zero in column (1). The N reduction from column (1) to columns (2)–(5) reflects this exclusion. Columns (3)–(5) build on the ring-excluded sample. Column (3) adds $Gold_g \times \ln P_{t-1}^{\text{gold}}$ to control for the 2002–2008 commodity co-movement; column (4) further adds temperature and precipitation. Column (5) uses the lagged oil price as the preferred specification, attenuating within-year reverse causality. The reported N is the common-sample size where all three outcomes are jointly observed (the DMSP-OLS night-lights coverage window 1992–2014 is the binding constraint); outcome-specific within-row estimation uses slightly larger samples for the conflict outcomes (approximately 1.59 million ring-excluded cell-years for conflict, versus 1.30 million for night lights), as documented in the sample-flow paragraph of Section 3. The qualitative pattern of coefficients is unchanged under outcome-specific samples. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

ficients on non-state and one-sided violence. Aggregate within-country oil-price exposure does not produce systematic conflict escalation in the average petroleum cell; the evidence does not support a quantified peace dividend.

Taken together, the results identify a positive within-country-year economic-activity response to oil-price exposure with no corresponding average conflict response. The confidence intervals consistently exclude economically meaningful escalation in violence. The pattern is consistent with a local transmission of oil-price exposure that operates primarily through economic channels without inducing systematic violence; it does not by itself identify which of the multiple plausible mechanisms — security deployments, opportunity-cost pacification, fiscal redistribution, rent-contestation incentives — is responsible.

6.1 Where Does the Effect Live? Spatial Aggregation and the Country-Level Resource Curse

The results above establish that oil-price exposure generates a positive within-country-year economic-activity response with no corresponding systematic conflict escalation. A natural question is whether this response is truly local or instead reflects broader country-level forces that are imperfectly captured at the cell level. To address this, I examine how the estimated effects change as the unit of observation is progressively aggregated from 0.5-degree grid cells to larger geographic units. This provides a direct test of the spatial scale at which the treatment operates. If the estimates reflect genuinely local mechanisms—such as production-site rents, targeted security deployment, or localized fiscal spillovers—the effect should attenuate as exposure is averaged over larger areas. By contrast, if the estimates are driven by country-level confounding, they should remain stable under aggregation.

Table 4 implements this comparison. Across increasingly coarse spatial resolutions, the estimated effect declines sharply once expressed in per-cell-equivalent terms. For night lights, the implied effect falls by an order of magnitude as the unit expands from the cell to broader geographic blocks. Conflict estimates follow a similar pattern, becoming economically small and indistinguishable from zero at wider spatial scales.

This rapid attenuation indicates that the effects identified in Table 3 are genuinely local. Most of the impact dissipates once exposure is averaged over distances of a few hundred kilometers, consistent with mechanisms operating in the immediate vicinity of extraction sites.

The country-level row of Table 4 reports the analogous estimate when identification

Table 4: Spatial-aggregation attenuation

Resolution	N_k	Unit FE	(1) *Night lights*	(2) *Conflict events*	(3) *Conflict deaths*
Panel A: Raw coefficient $\hat{\beta}_k$ (per unit increase in share)					
0.5° cell	1	cell	$+0.0407^{***}$	-0.0042	-0.0042
			(0.0040)	(0.0028)	(0.0043)
2° district	16	2° block	$+0.1075^{***}$	-0.0133^{**}	-0.0222^{***}
			(0.0101)	(0.0057)	(0.0085)
5° province	100	5° block	$+0.1957^{***}$	-0.0185^{**}	-0.0435^{***}
			(0.0137)	(0.0088)	(0.0123)
10° region	400	10° block	$+0.3855^{***}$	$+0.0034$	-0.0082
			(0.0191)	(0.0075)	(0.0104)
Country	–	country	$+0.6396^{***}$	$+0.1434^{***}$	$+0.2081^{***}$
			(0.0239)	(0.0268)	(0.0378)
Panel B: Per-cell-equivalent response $\hat{\beta}_k / N_k$ (unit-comparable)					
0.5° cell	1	cell	$+0.04068^{***}$	-0.00423	-0.00418
2° district	16	2° block	$+0.00672^{***}$	-0.00083^{**}	-0.00139^{***}
5° province	100	5° block	$+0.00196^{***}$	-0.00019^{**}	-0.00043^{***}
10° region	400	10° block	$+0.00096^{***}$	$+0.00001$	-0.00002
Country	–	country	*n.a.*	*n.a.*	*n.a.*
N (obs)			1,301,197	1,593,488	1,593,488
Ring cells excl.			✓	✓	✓

Notes: Each row reports OLS estimates of $\hat{\beta}_k$ from the regression $Y_{gct} = \beta_k(share_g^{(k)} \times \ln P_t) + \alpha_{B(g,k)} + \lambda_{c,t}^{(k)} + \varepsilon_{gct}$, where $share_g^{(k)}$ is the fraction of cells in the resolution-k block $B(g,k)$ containing cell g that are petroleum cells (`petroleum_s` > 0); the share is time-invariant (geological) and therefore predetermined. The cell-level row reproduces the binary $Oil_g \times \ln P_t$ treatment of Table 3 expressed as a $\{0,1\}$ share. **Fixed effects are matched to each level so that $share_g^{(k)} \times \ln P_t$ is identifiable.** For the cell, 2°, 5°, and 10° rows, the unit FE is the resolution-k block (column "Unit FE") and the time-side FE is country×year, so $\hat{\beta}_k$ is identified from the within-country-year cross-section of blocks. For the country row, $share_c \times \ln P_t$ is country-constant × year-constant and would be absorbed by country×year FE; we use country FE plus year FE instead, recovering identification from across-country-year variation. **Panel A** reports raw $\hat{\beta}_k$, the response to a 1-unit increase in $share_g^{(k)}$ (going from 0 to 100% oil cells in the block). Because a fully-saturated block represents far more oil at coarser resolutions (N_k cells), raw $\hat{\beta}_k$ scales mechanically with N_k and is not directly comparable across rows. **Panel B** rescales by N_k: $\hat{\beta}_k / N_k$ is the implied effect of changing one cell from non-oil to oil somewhere in the block, and *is* comparable across rows. The cell-level row ($N_k = 1$) is the identifying benchmark; Panel B coefficients attenuate sharply as the block widens, confirming the effect is concentrated in the local cell. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

relies on between-country variation. Without within-country fixed effects, country-level oil share is positively associated with conflict, reproducing the stylized fact that oil-rich countries are more conflict-prone. This positive cross-country association is not recovered under the within-country design.

The exercise establishes two scope statements that the paper is comfortable making and one it is not. First, the local response to oil-price exposure is concentrated spatially: per-cell-equivalent effects decay sharply as the unit of observation expands beyond the cell. Second, the cross-country positive correlation between oil exposure and conflict is recovered in this sample only when within-country variation is removed; under within-country identification the local effect is not conflict-inducing. The paper does *not* identify the source of the cross-country association — the design absorbs country-level confounds rather than measuring them. The cross-country pattern is consistent with between-country differences (institutions, conflict history, ethnic composition) driving the country-level association, but separating these explanations from a country-level oil effect would require a different design. The reconciliation reported here is therefore a scope statement, not a causal diagnosis of the cross-country literature.

Figure 5: **Spatial-aggregation attenuation: per-cell-equivalent oil response by geographic resolution**

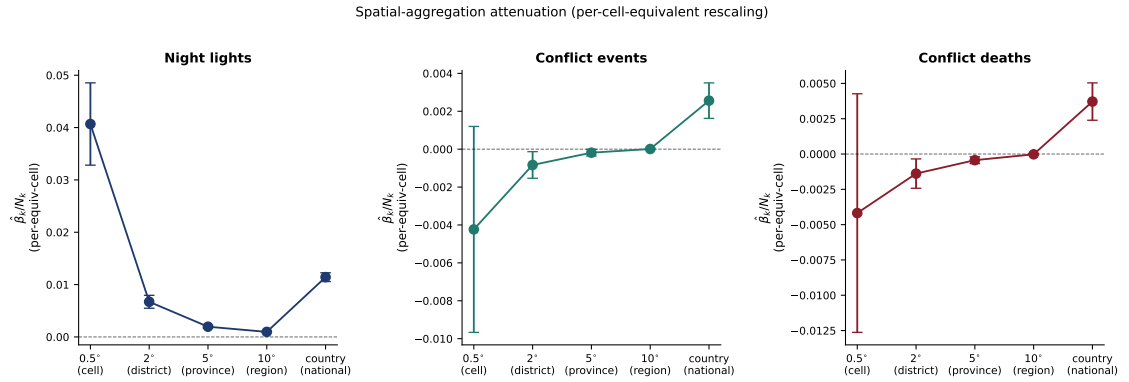

Per-cell-equivalent coefficient $\hat{\beta}_k / N_k$ at each resolution, where N_k is the cells per block. The cell-level point ($N_k = 1$) is the baseline. As the block widens, the implied per-cell effect shrinks, confirming that the response is concentrated in the local cell. The country-level point uses country FE plus year FE (since country×year FE absorbs $share_c \times \ln P_t$). Bars are 95% confidence intervals clustered at the PRIO-GRID cell level.

7 Institutional Heterogeneity in the Development Response

The results above establish two facts. Oil windfalls raise local economic activity, and they do not generate systematic conflict escalation once spillovers are accounted for. The remaining question is distributional: *where do the gains from the windfall land?* The

cell-level reduced form cannot directly observe fiscal transfers or public-goods flows; what it can measure is whether the differential local response varies systematically with the institutional features that the canonical mechanisms identify as gatekeepers — administrative capacity and political access.

Theoretical accounts emphasize that resource rents do not automatically translate into local prosperity. Their transmission depends on two distinct institutional margins. First, *state capacity* determines whether governments can convert revenues into local public goods and economic activity. Second, *political access* determines whether local populations are able to claim a share of those resources. These two margins are conceptually separate: a state may have the administrative capacity to deliver public goods but exclude local populations from accessing them, or it may be inclusive but lack the capacity to implement redistributive policies.

This section tests these two margins directly. The results show that the development response is systematically shaped by both state capacity and ethnic political inclusion. The conflict response, by contrast, does not exhibit a consistent pattern of institutional heterogeneity in any of the four moderator dimensions tested below.

7.1 State Capacity and Competing Institutional Channels

I begin with state capacity. If governments lack the administrative ability to transform oil revenues into public goods, local economic gains should be attenuated even when windfalls are large. To test this, I use a pre-sample measure of state capacity based on the ICRG quality-of-government index. I compare this to two alternative institutional dimensions: political openness, measured using V-Dem Polyarchy, and oil-sector governance, measured using the 2013 Natural Resource Governance Index (NRGI). These capture conceptually distinct margins: administrative capacity, electoral competition, and sector-specific governance of resource rents. All measures are pre-determined or time-invariant, ensuring that the moderators are not affected by contemporaneous oil price movements.

Table 5: Pre-1989 time-invariant state capacity as moderator

	(1) ICRG pre-1989 1984–1988 avg.		(2) V-Dem poly 1970–1988 avg.		(3) NRGI 2013 composite	
	Events	Deaths	Events	Deaths	Events	Deaths
Panel A: Night lights						
$\hat{\beta}$ (Oil$\times \ln P$)	$+0.0523^{***}$		$+0.0487^{***}$		$+0.0490^{***}$	
	(0.0059)		(0.0053)		(0.0055)	
$\hat{\theta}$ ($\times$ Low capacity)	-0.0386^{***}		-0.0176^{*}		-0.0071	
Panel B: Conflict outcomes						
$\hat{\beta}$ (Oil$\times \ln P$)	-0.0013	-0.0019	-0.0102^{***}	-0.0146^{**}	-0.0026	-0.0025
	(0.0024)	(0.0036)	(0.0037)	(0.0058)	(0.0033)	(0.0059)
$\hat{\theta}$ ($\times$ Low capacity)	-0.0054	-0.0014	$+0.0122$	$+0.0225^{*}$	-0.0009	-0.0015
	(0.0106)	(0.0174)	(0.0078)	(0.0123)	(0.0051)	(0.0085)
Cell FE	✓					
Country-year FE	✓					
Lower-order absorbed	Yes (cell FE & country-year FE)					
N (events/deaths)	924,186		959,186		1,041,052	
N (night lights)	$\approx$ 817,000		$\approx$ 847,000		$\approx$ 920,000	

Notes: OLS estimates of $Y_{gct} = \beta(Oil_g \times \ln P_t) + \theta(Oil_g \times LowCap_g \times \ln P_t) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$. $LowCap_g$ is a time-invariant (pre-sample) below-median indicator: column (1) uses ICRG quality-of-government averaged 1984–1988; column (2) uses V-Dem electoral democracy averaged 1970–1988; column (3) uses the 2013 NRGI Resource Governance Index composite (oil-sector governance). Because $LowCap_g$ is constant within cell, $Oil_g \times LowCap_g$ is absorbed by cell fixed effects and $LowCap_g \times \ln P_t$ by country-year fixed effects; no lower-order interaction term is required. The observation count differs across columns because each moderator covers a different set of countries: ICRG covers 122 countries (1984–1988), V-Dem covers 143 (1970–1988), and NRGI covers only 58 resource-rich countries by construction. NRGI has the fewest countries but the largest N because its universe is biased toward large oil/gas economies (Russia, China, USA, Indonesia, Brazil, Saudi Arabia, etc.), each contributing many cells. The night-lights row is estimated on a slightly smaller sample because DMSP-OLS coverage begins in 1992 rather than 1989; the approximate night-lights Ns reported above are 88% of the corresponding events/deaths Ns and reflect this three-year window difference. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 5 and Figure 6 report the results. The baseline development response is positive and stable across the three specifications. The institutional moderation differs sharply across measures.

The ICRG state-capacity moderation is large, precisely estimated, and economically meaningful. The implied within-country-year response in low-capacity cells is approximately one quarter of the high-capacity baseline — a substantial attenuation, though the low-capacity coefficient is itself positive and statistically distinguishable from zero. Substantively, weak administrative capacity sharply reduces the local response to oil-price exposure without quite eliminating it. The V-Dem polyarchy interaction is negative but only marginally significant, suggesting at most modest attenuation in less democratic settings. The NRGI interaction is statistically indistinguishable from zero, indicating no meaningful moderation along the oil-sector governance dimension.

The contrast is informative: political openness and sector-specific oil-governance structures shape broader institutional quality, but the data do not detect them as moderators of the local development response in this sample, while administrative capacity does emerge as a detectable moderator. The interpretation is consistent with general administrative capacity — rather than the formal architecture of oil-sector regulation — being the binding constraint at the local cell-year level, but the comparison is a relative ranking among the moderators tested rather than a causal claim about institutions themselves.

Panel B shows a different pattern for conflict outcomes. Across all three institutional measures, the estimates are small and imprecise, and the data do not detect a consistent pattern of heterogeneity. The interpretation deserves a power caveat: imprecise conflict interactions are consistent with both *no* institutional moderation and *moderation too small to detect at this sample size*. The data therefore support the conclusion that the paper fails to detect systematic heterogeneity in the local conflict response across institutional environments, rather than the stronger claim that no such heterogeneity exists. The contrast with the precisely-estimated development moderation is sharp: the development gradient is detectable in this sample; the conflict gradient is not.

Figure 6: **Triple-interaction visualisation**

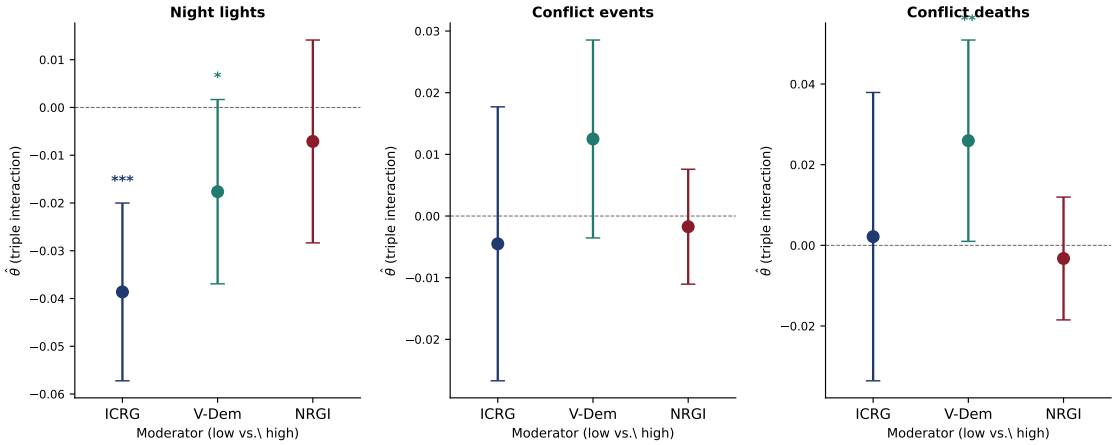

Triple-interaction coefficient $\hat{\theta}$ from $Y_{gct} = \beta\,(Oil_g \times \ln P_t) + \theta\,(Oil_g \times \ln P_t \times LowMod_c) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$, estimated separately for three country-constant moderators: ICRG quality of government (1984–1988 average, navy), V-Dem polyarchy (1970–1988 average, teal), NRGI Resource Governance Index (2013, crimson). Bars are 95% CIs clustered at the PRIO-GRID cell level. Sample matches Table 5. Stars: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure A4 plots the results in low-moderator versus high-moderator subsamples, separately for four institutional measures spanning distinct dimensions: general capacity, political openness, oil-sector governance, and income (GDP per capita, 1985–1988 average). The results show that the development response to oil-price exposure varies

sharply with state capacity but only weakly across other institutional dimensions. In contrast, conflict effects are small and imprecisely estimated across all subsamples, with no consistent pattern of heterogeneity. Among the four moderators tested in this set, administrative capacity is the strongest reported gradient of the local development response; I make no formal multiple-testing adjustment, so the comparison should be read as the relative rank of moderators in the presented set rather than a claim about all possible institutional dimensions.

7.2 Ethnic Political Exclusion

The second margin concerns political access to rents. Even where the state can deliver public goods, local populations may fail to benefit if they are excluded from political power (Müller-Crepon et al., 2023; Esteban and Ray, 2008; Cederman et al., 2010; Michalopoulos and Papaioannou, 2016).

To test this channel, I use georeferenced data on ethnic political exclusion to construct a time-invariant indicator for whether a cell hosts a politically excluded group prior to the sample period. I then interact this measure with oil price exposure. In particular, I construct a time-invariant cell-level indicator $Excl_g \in \{0, 1\}$ equal to one if cell g hosted at least one EPR-classified politically excluded ethnic group at any point during 1980–1988, the pre-sample period. This uses the `excluded` variable from PRIO-GRID v3, which is itself derived from Vogt et al. (2015)'s GeoEPR settlement polygons and the EPR core dataset of political-status classifications. The pre-period construction makes $Excl_g$ predetermined relative to the 1989–2014 oil-price cycle, mirroring the pre-1989 ICRG capacity moderator above. In the ring-excluded sample, 42 percent of petroleum cells overlap politically excluded ethnic settlement zones — a substantial sub-population. The triple-interaction specification is

$$Y_{gct} = \beta \left(Oil_g \times \ln P_t \right) + \gamma \left(Oil_g \times \ln P_t \times Excl_g \right) + \alpha_g + \lambda_{ct} + \varepsilon_{gct},$$

where lower-order terms involving $Excl_g$ are absorbed by cell fixed effects and country×year fixed effects exactly as for $LowCap_c$.

Table 6: Ethnic political exclusion and oil-price exposure: heterogeneity in the local development response, no consistent heterogeneity in conflict (ring-excluded sample, 1989–2014)

	Night lights		*Conflict events*		*Conflict deaths*	
	Excl-only	Combined	Excl-only	Combined	Excl-only	Combined
$Oil_g \times \ln P_t$ ($\hat{\beta}$)	$+0.0536^{***}$	$+0.0644^{***}$	-0.0060^{***}	-0.0026	-0.0061^{*}	-0.0050
	(0.0052)	(0.0065)	(0.0022)	(0.0035)	(0.0034)	(0.0054)
$\times\ Excl_g$ ($\hat{\gamma}$)	-0.0304^{***}	-0.0277^{***}	$+0.0041$	$+0.0042$	$+0.0045$	$+0.0095$
	(0.0079)	(0.0094)	(0.0064)	(0.0083)	(0.0102)	(0.0132)
$\times\ LowCap_c$ ($\hat{\theta}$)		-0.0386^{***}		-0.0032		$+0.0065$
		(0.0095)		(0.0111)		(0.0179)
N	1,301,257	924,186	1,593,551	1,128,107	1,593,551	1,128,107
Cell FE	✓	✓	✓	✓	✓	✓
Country×year FE	✓	✓	✓	✓	✓	✓
Ring cells excl.	✓	✓	✓	✓	✓	✓

Notes: OLS estimates of the heterogeneity by ethnic political exclusion. $Excl_g$ is a time-invariant cell-level indicator equal to one if the cell hosted at least one EPR-classified politically-excluded ethnic group at any point during the pre-sample period 1980–1988 (PRIO-GRID v3 `excluded` variable, derived from GeoEPR). The pre-period construction makes $Excl_g$ predetermined relative to the 1989–2014 oil-price cycle, mirroring the pre-1989 ICRG state-capacity moderator $LowCap_c$ used in Section 7.1. Both $Excl_g$ and $LowCap_c$ are absorbed by the cell and country×year fixed effects in their lower-order forms; the table reports the triple-interaction coefficients on $Oil_g \times \ln P_t \times Excl_g$ ($\hat{\gamma}$) and on $Oil_g \times \ln P_t \times LowCap_c$ ($\hat{\theta}$). The Combined column estimates both interactions jointly, providing a horse-race between ethnic-structural and state-capacity moderators. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 6 reports two specifications per outcome: a single-moderator model with $Excl_g$ alone, and a combined model adding $LowCap_c$ to horse-race the two institutional dimensions. The result is striking and runs against the canonical resource-curse prediction for conflict, but it falls instead on the development margin. Petroleum cells hosting politically excluded groups exhibit roughly half the local development response that included cells exhibit, a moderation whose magnitude is comparable in size to the state-capacity moderation reported in Section 7.1. Both interactions retain independent statistical significance in the combined horse-race specification, and they do not substitute for each other — political access and state capacity are independent margins of the institutionally conditioned local transmission documented here.

As with state capacity, there is no corresponding heterogeneity in conflict outcomes. The interaction between oil exposure and political exclusion is small and statistically insignificant. The local conflict response remains stable across inclusion and exclusion settings, and the other four institutional dimensions tested in this paper (capacity, income, democracy, ethnic exclusion).

Taken together, the evidence is consistent with a dual-margin institutional mechanism in which the local response to oil-price exposure depends on both state capacity

and political inclusion. The two margins enter additively in the horse-race specification, and each independently pulls the response down. Figure 7 maps the implied response coefficient $\hat{\beta}_{gc}$ on $Oil_g \times \ln P_t$ onto the four cells of the (capacity, inclusion) grid. High-capacity included cells receive the full $+0.064$ response. Holding state capacity fixed and turning off political access reduces the response to $+0.037$; holding inclusion fixed and turning off state capacity reduces it to $+0.026$. Each off-diagonal cell is positive but attenuated — the local within-country-year response survives the failure of either margin alone, but at roughly half the high-capacity included magnitude. Only when both margins bind does the implied response collapse to $\hat{\beta} + \hat{\theta} + \hat{\gamma} = -0.002$, statistically indistinguishable from zero. This is heterogeneity by institutional category, not a causal effect of changing capacity or inclusion; the design does not exploit exogenous variation in institutions themselves. The pattern — additive attenuation across institutional margins, with the floor reached only when both margins bind — is the empirical analogue of the model's comparative statics in Section 2.2.

Figure 7: **Dual-margin moderation of the local response.**

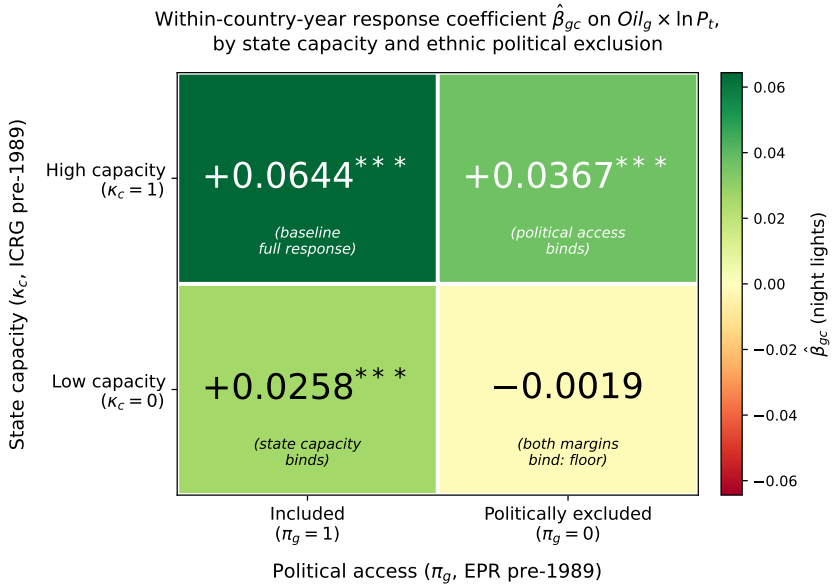

Notes: The 2×2 heatmap displays the implied within-country-year response coefficient $\hat{\beta}_{gc}$ on $Oil_g \times \ln P_t$ for the four (state capacity, political inclusion) combinations, computed from the combined horse-race in Table 6: high-capacity included $= \hat{\beta} = +0.0644^{***}$; high-capacity excluded $= \hat{\beta} + \hat{\gamma} = +0.0367^{***}$; low-capacity included $= \hat{\beta} + \hat{\theta} = +0.0258^{***}$; low-capacity excluded $= \hat{\beta} + \hat{\theta} + \hat{\gamma} = -0.0019$ (n.s.). The pattern is one of additive attenuation: each margin reduces the response independently, and only when both bind does the implied response collapse to zero. All three constituent coefficients $(\hat{\beta}, \hat{\theta}, \hat{\gamma})$ are statistically significant at the 1% level in Table 6; the linear-combination significance for cells (0,1) and (1,0) is reported as *** following the regression evidence, while cell (1,1) sums to a value statistically indistinguishable from zero. Colours: dark green for the full-response cell, white for the fully-attenuated cell. Reading: petroleum cells in low-capacity countries with political inclusion (cell (1,0): $+0.026$) still exhibit a positive within-country-year response, but the response is attenuated by approximately 60% relative to the high-capacity included baseline of $+0.064$; cells where both margins bind exhibit no detectable response. The model in Section 2.2 (equation 2) predicts the qualitative pattern — full response when both margins are favourable, floor when both bind — though the empirical off-diagonal cells exhibit partial rather than full attenuation. Significance markers: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The absence of a consistent pattern of conflict heterogeneity in the pooled-sample horse-race provides an important contrast. While the development response varies sharply with institutional conditions, the conflict $\hat{\theta}$ and $\hat{\gamma}$ coefficients in Table 6 are small and statistically indistinguishable from zero on average. This is consistent with offsetting forces operating at the cell scale — security deployments and opportunity-cost pacification on one side, rent-contestation incentives on the other — which can cancel when averaged across calm and high-price years and across cells with different institutional environments.

An event-study analog of the four-cell dual-margin partition, reported in Appendix B (Figure A5), refines this pooled null. Table 7 summarises the average of the 2002–2008 year-by-year coefficients for each (state-capacity, exclusion) cell, together with formal subgroup-heterogeneity tests. Boom-period conflict events in low-capacity excluded (LC-E) cells average $+0.084$ and conflict deaths average $+0.129$, an order of magnitude larger than the boom-period averages in the other three subgroups. The pre-2001 LC-E coefficients are statistically indistinguishable from zero ($\chi^2_{12} = 13.0$, $p = 0.37$ for events; $\chi^2_{12} = 12.5$, $p = 0.41$ for deaths), so the spike is a boom-period feature rather than a continuation of pre-existing trends.

The formal contrast $\bar{\beta}^{(\text{LC-E})} - \frac{1}{3}\sum_{s \neq \text{LC-E}} \bar{\beta}^{(s)}$ is positive and statistically significant at conventional levels for both conflict outcomes ($\Delta = +0.081$, $p = 0.021$ for events; $\Delta = +0.125$, $p = 0.037$ for deaths). A joint Wald test of equality across all four boom averages rejects at the 5% level for events ($\chi^2_3 = 10.9$, $p = 0.012$) and is borderline for deaths ($\chi^2_3 = 7.5$, $p = 0.057$). The LC-E coefficient is therefore not equal to the average response in the other three subgroups; the visual pattern of the four-cell event study survives a formal contrast test.

Multiple testing. Within the four-cell partition exercise we report two main contrasts per outcome — joint equality of the four subgroup boom averages, and LC-E versus the average of the other three — yielding six tests across the three outcomes. Applying a Bonferroni correction across these six tests, the LC-E-versus-others contrast has $p_{\text{Bonf}} = 0.126$ on conflict events and $p_{\text{Bonf}} = 0.220$ on conflict deaths: both unadjusted contrasts are significant at 5%; neither Bonferroni-corrected contrast clears conventional thresholds. Two features make Bonferroni overly conservative here: the LC-E subgroup was identified ex ante by the dual-margin model in Section 2.2, not selected by an ex-post search across institutional cuts; and the events and deaths outcomes are highly correlated within cell, so a family-wise correction that treats them as independent over-

Table 7: Boom-period averages (2002–2008) of subgroup-specific event-study coefficients

	Night lights	*Conflict events*	*Conflict deaths*
High-cap, Included	$+0.0034$	-0.0060	-0.0054
	(0.0055)	(0.0043)	(0.0048)
High-cap, Excluded	-0.0132^{**}	$+0.0093^{*}$	$+0.0124$
	(0.0054)	(0.0053)	(0.0101)
Low-cap, Included	-0.0154^{**}	$+0.0050$	$+0.0024$
	(0.0062)	(0.0166)	(0.0296)
Low-cap, Excluded	-0.0037	$+0.0838^{**}$	$+0.1286^{**}$
	(0.0106)	(0.0345)	(0.0590)
Test 1: all four boom averages equal (H_0: 3 restrictions)			
χ^2_3	7.0	10.9	7.5
p-value	0.072	0.012	0.057
Test 2: LC-E versus the average of the other three subgroups			
$\hat{\Delta} = \bar{\beta}^{(\text{LC-E})} - \frac{1}{3}\sum_{s\neq\text{LC-E}} \bar{\beta}^{(s)}$	$+0.0047$	$+0.0810$	$+0.1254$
p-value	0.674	0.021	0.037
p-value (Bonferroni, 6 tests)	1.000	0.126	0.220

Notes: Each cell reports the simple average of the seven year-by-year coefficients $\hat{\beta}_k^{(s)}$ on $Oil_g \times \mathbb{1}\{\text{year}= k\} \times \mathbb{1}\{g \in s\}$ for $k = 2002, \ldots, 2008$ within subgroup s, with standard error in parentheses. Coefficients and standard errors are taken from the four-cell partition event-study (Appendix B); averages and their standard errors are computed via the contrast $\bar{\beta}^{(s)} = \frac{1}{7}\sum_{k=2002}^{2008} \hat{\beta}_k^{(s)}$ with the full coefficient covariance matrix. *Tests*: Row "All four boom averages equal" is a joint Wald test of $H_0 : \bar{\beta}^{(\text{HC-I})} = \bar{\beta}^{(\text{HC-E})} = \bar{\beta}^{(\text{LC-I})} = \bar{\beta}^{(\text{LC-E})}$ (3 restrictions, χ^2_3). Row "LC-E vs avg of other 3" is a Wald test of $H_0 : \bar{\beta}^{(\text{LC-E})} = \frac{1}{3}\sum_{s\neq\text{LC-E}} \bar{\beta}^{(s)}$. Bonferroni correction is applied across the 6 LC-E-vs-other-3 tests in the paper (2 main contrasts $\times$ 3 outcomes). *Important caveat on the night-lights column:* all four night-lights subgroups fail the parallel-trends test in the underlying event-study (see Table A1) because the DMSP-OLS satellite sequence carries a known calibration drift; the night-lights boom averages should therefore be read as descriptive within-subgroup summaries rather than as confirmatory tests of the development gradient. The paper's development claim rests on the continuous-treatment OLS design (Table 5) and the supply-shock IV (Appendix G.1), not on this event-study. The conflict columns rely on subgroups that do pass parallel-trends tests (HC-I, HC-E, LC-E); the LC-I caveat is documented in Appendix B. Cell + country$\times$year FE; 8-neighbor ring-excluded sample; cell-clustered SEs. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

states the burden. Even with both qualifications, the corrected evidence is not sufficient to elevate the LC-E pattern to a confirmed result. The LC-E spike is therefore *suggestive*: economically large, statistically significant before correction, and free of pre-trend contamination in that subgroup, but not robust to a conservative family-wise correction. It is a theoretically anticipated pattern that motivates further investigation; it is not a settled finding of the same standing as the development result. The LC-E subpopulation is 173 of 1,300 petroleum cells with both moderators observed (13.3%); the pooled-sample triple-interaction null in Table 6 averages this boom-period LC-E spike together with the conflict-quiet behaviour of the other three subgroups.

The local response to oil-price exposure is therefore not mechanically conflict-inducing in the average petroleum cell, but the conflict response is sharply conditional on institutional context: the formal evidence supports concentration in low-capacity excluded cells during high-price periods. The development response is moderated by institutional context on the development margin (Section 7.1); the conflict response, where it appears, is moderated on both the institutional margin (LC-E versus the other three subgroups) and the time margin (boom versus non-boom years).

8 Channels and Mechanisms

Oil windfalls increase local economic activity but do not generate systematic increases in conflict. This combination of effects is not obvious ex ante. A large literature associates natural-resource booms with rent-seeking and violence, particularly in weakly institutionalized settings. The evidence here points to a different local transmission. The key question is therefore not whether oil matters, but *how* its effects propagate within countries.

Is the conflict null masking offsetting effects? A natural concern is that the muted aggregate conflict response reflects offsetting movements across different forms of violence. Table 8 addresses this directly. Disaggregating conflict into state-based, non-state, and one-sided violence yields no coherent pattern: the estimated effects are uniformly small, with positive and negative coefficients of similar magnitude. This is informative. If oil windfalls triggered local contestation over rents—as in canonical models of resource conflict—one would expect a broad-based increase in organized violence. The absence of such a pattern suggests that, at the local level, oil does not systematically intensify conflict on a broad basis across violence types, though Tables 8 and 13 show modest positive responses on the non-state and one-sided-violence margins.

Table 8: Conflict-type decomposition

	(1) All events	(2) All deaths	(3) State-based	(4) Non-state	(5) One-sided
$\text{Oil}_g \times \ln P_t$	-0.0042	-0.0041	-0.0039	$+0.0015^{*}$	-0.0021^{*}
	(0.0028)	(0.0043)	(0.0025)	(0.0008)	(0.0012)
Cell FE	✓	✓	✓	✓	✓
Country-year FE	✓	✓	✓	✓	✓
N	1,593,551	1,593,551	1,593,551	1,593,551	1,593,551

Notes: OLS primary design. State-based: events with at least one state actor. Non-state: events between non-state organised groups. One-sided: violence against civilians. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Where do the rents go? The next question concerns the geography of the gains. If oil revenues are centralized and redistributed nationally, local effects should be weak. If instead rents flow into production areas, one should observe localized economic responses.

Table 9 provides evidence on this margin. Development gains are present in both onshore and mixed producers, indicating that oil revenues translate into economic activity in areas connected to extraction.[9] The fact that similar patterns emerge even when production is partly offshore suggests that the relevant mechanism is not narrowly tied to physical drilling locations, but to broader economic linkages—procurement, services, and local demand. Conflict responses remain weak across these groups. Taken together, the evidence points to a *local rent-transmission channel* in which oil revenues raise economic activity in surrounding areas without generating systematic contestation.

The conflict results differ across production structures. While effects are small and imprecise in onshore producers, the mixed group—comprising several of the largest oil-producing economies— exhibits a statistically meaningful reduction in conflict deaths.[10]

One interpretation is that the conflict response depends on the economic and strategic importance of oil production. In mixed producers, extraction is both large-scale and closely integrated into national revenue systems, which may increase incentives to stabilize production areas. This could operate through mechanisms such as infrastructure protection or the concentration of state presence in oil-producing regions. An additional feature of the data is that development gains are larger in mixed producers than in onshore producers. This is consistent with standard opportunity-cost mechanisms

[9]The production classification follows Andersen et al. (2022), who flag whether a country's marketed production is predominantly extracted on land or at sea. Each country is assigned to onshore, offshore, or mixed, and the primary design is re-estimated within each subsample

[10]The pure offshore subsample (Angola, Gabon, Congo, Malaysia) contributes only 16 unique petroleum cells across four countries, so the within-country-year design is severely underpowered.

in which higher local economic activity reduces the relative payoff to participation in violence.

Table 9: Onshore vs. offshore decomposition

	(1) *Full sample*	(2) *Onshore countries*	(3) *Offshore countries*	(4) *Mixed countries*
Night lights	$+0.0407^{***}$	$+0.0365^{***}$	$+0.0479$	$+0.0551^{***}$
	(0.0040)	(0.0046)	(0.0656)	(0.0083)
Conflict events	-0.0042	-0.0036	-0.0185	-0.0057^{**}
	(0.0028)	(0.0035)	(0.0228)	(0.0026)
Conflict deaths	-0.0042	-0.0020	-0.0424	-0.0105^{***}
	(0.0043)	(0.0054)	(0.0536)	(0.0037)
N	1,593,551	1,255,254	19,136	319,161
Countries	195	181	4	10
Cell FE	✓	✓	✓	✓
Country-year FE	✓	✓	✓	✓

Notes: OLS estimates of β in $Y_{gt} = \beta\,(Oil_g \times \ln P_t) + \alpha_g + \lambda_{ct} + \varepsilon_{gt}$. Countries are classified as predominantly onshore (Algeria, Russia, China, Saudi Arabia, Libya, Iraq, Colombia, and other land-based producers), offshore (Angola, Gabon, Congo, Malaysia), or mixed (USA, Indonesia, Australia, Brazil, Iran, UAE, Azerbaijan, Nigeria, Vietnam, Kuwait). Each column restricts to all cells — oil and non-oil — in that country group. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Do oil revenues buffer economic shocks? A further implication of a rent-based channel is that oil revenues may act as a buffer against local income shocks. In many of the settings in this sample, households depend heavily on agriculture, and adverse weather realizations are known to increase conflict. To test this, I augment the primary specification with a triple interaction:

$$Y_{gt} = \beta_1(Oil_g \times \ln P_t) + \beta_2(Oil_g \times \ln P_t \times ExtrTemp_{gt}) + \beta_3\, ExtrTemp_{gt} + \alpha_g + \lambda_{ct} + \varepsilon_{gt},$$

where $ExtrTemp_{gt}$ equals one if cell g's annual mean temperature in year t exceeds its own long-run cell mean by more than two standard deviations. Country-year fixed effects absorb the national-level effects of $ExtrTemp_{gt}$; the coefficient β_2 therefore captures whether petroleum-cell exposure differentially cushions the local impact of extreme heat shocks.

Table 10 shows that petroleum-exposed cells exhibit a higher within-country-year night-light response during extreme-temperature years than during normal years. This is consistent with — but does not directly observe — a mechanism in which oil revenues support local demand or public-goods provision when agricultural income falls. The data measure night-light intensity rather than household income, fiscal transfers, or

public-goods provision; the buffering is a buffering in the measured construct (within-country-year night-light differential), not directly in income or welfare. The corresponding attenuation of conflict responses is absent: oil-price exposure does not materially alter the link between weather shocks and violence at the cell-year level.

The pattern in night-light activity during extreme-temperature years is consistent with a transmission channel that runs through economic activity rather than directly through conflict stabilisation. Direct identification of the income / fiscal / public-goods channel would require mediator-level data not available at the cell-year frequency.

Table 10: Livelihood insurance: oil windfalls and extreme temperature shocks

	(1) *Night lights*	(2) *Conflict events*	(3) *Conflict deaths*
$Oil_g \times \ln P_t$	$+0.0351^{***}$	-0.0042	-0.0055
	(0.0039)	(0.0028)	(0.0045)
$Oil_g \times \ln P_t \times ExtrTemp_{gt}$	$+0.0082^{***}$	-0.0004	-0.0004
	(0.0022)	(0.0019)	(0.0029)
$ExtrTemp_{gt}$	-0.0103^{***}	$+0.0043^{***}$	$+0.0085^{***}$
	(0.0011)	(0.0015)	(0.0026)
Common-sample N	1,334,000	1,334,000	1,334,000
Cell FE	✓	✓	✓
Country-year FE	✓	✓	✓

Notes: OLS estimates. Dependent variables are (1) $\ln(\text{night lights} + 1)$, (2) $\ln(\text{conflict events} + 1)$, (3) $\ln(\text{conflict deaths} + 1)$. $Oil_g \times \ln P_t$ is the main treatment; $ExtrTemp_{gt}$ equals one if the annual temperature in cell g exceeds its long-run cell mean by more than two standard deviations (computed using the 1989–2014 cell-level moments). Sample: full panel (ring-adjacent cells included) for consistency with the original specification; the cycle-asymmetry and ICRG-moderator specifications use the ring-excluded sample. The reported N is the common-sample size with all three outcomes jointly observed (the DMSP-OLS night-lights window 1992–2014 is the binding constraint); outcome-specific within-row estimation uses slightly larger samples for the conflict outcomes, as documented in the sample-flow paragraph of Section 3. All regressions include cell and country-year fixed effects. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

What happens over the price cycle? A final test examines whether the response to oil price changes is symmetric. If conflict is driven by income fluctuations or grievances, one would expect violence to increase during downturns, when economic conditions deteriorate. I decompose the year-on-year log price change into a boom component $\Delta \ln P_t^+ = \max(\Delta \ln P_t, 0) \geq 0$ and a bust component $\Delta \ln P_t^- = \max(-\Delta \ln P_t, 0) \geq 0$ that records the absolute magnitude of price declines. Both regressors are non-negative, so coefficient signs map directly to outcome direction: a positive boom coefficient means

the outcome rises during booms; a negative bust coefficient means it falls during busts. I estimate both interactions simultaneously:

$$Y_{gt} = \gamma^{+}(Oil_g \times \Delta \ln P_t^{+}) + \gamma^{-}(Oil_g \times \Delta \ln P_t^{-}) + \alpha_g + \lambda_{ct} + \varepsilon_{gt}.$$

Figure 8 shows a sharp contrast between economic and conflict responses. Development reacts asymmetrically: declines during busts are substantially larger than gains during booms, indicating that local economic activity is highly exposed to negative revenue shocks.

This finding is difficult to reconcile with grievance-based mechanisms in which violence rises during bust periods. Petroleum-cell conflict coefficients are small and do not display the bust-period reversal that a grievance-driven channel would predict, in either the boom or the bust regressor. Economic activity tracks oil prices closely; the conflict response does not. The cycle-invariant absence of systematic conflict escalation is consistent with offsetting forces at the cell scale — security deployments and opportunity-cost pacification on one side, rent-contestation incentives on the other — that operate regardless of the contemporaneous fiscal environment. The data do not separately identify which of these channels dominates.

Figure 8: **Price-cycle asymmetry: boom vs. bust effects (primary OLS design)**

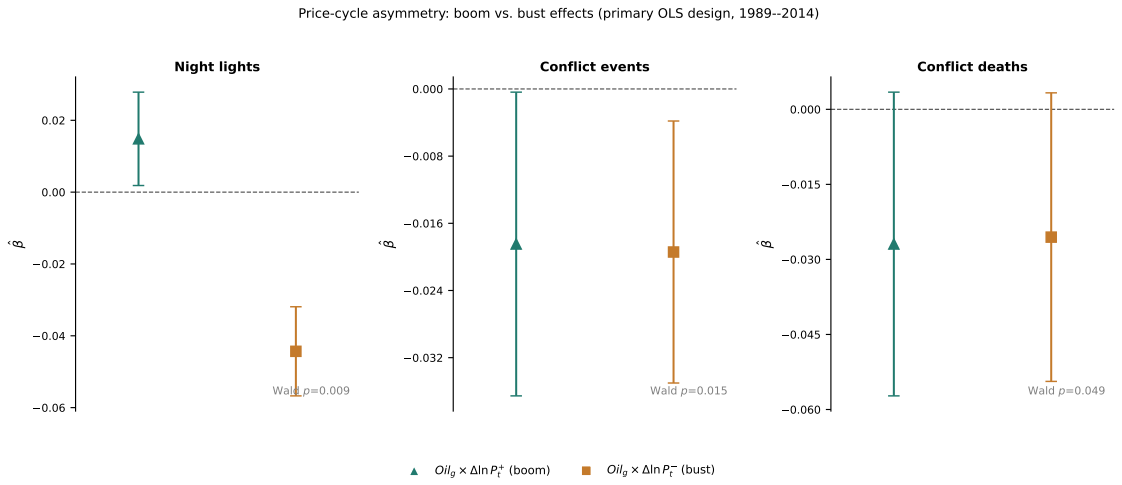

OLS estimates of boom ($\hat{\gamma}^{+}$, $Oil_g \times \Delta \ln P_t^{+}$, where $\Delta \ln P_t^{+} = \max(\Delta \ln P_t, 0) \geq 0$) and bust ($\hat{\gamma}^{-}$, $Oil_g \times \Delta \ln P_t^{-}$, where $\Delta \ln P_t^{-} = \max(-\Delta \ln P_t, 0) \geq 0$) interactions. Full sample including ring-adjacent cells ($N = 1{,}334{,}000$). Both regressors are non-negative, so coefficient signs map directly to outcome direction (positive = increase, negative = decrease). Wald p-values (testing the symmetric-response restriction $\hat{\gamma}^{+} + \hat{\gamma}^{-} = 0$) annotated in each panel. Bars are 95% confidence intervals clustered at the PRIO-GRID cell level.

Taking stock, the evidence is consistent with a mechanism in which oil-price exposure raises local economic activity — as measured by night-light intensity — through

production-area economic rents, and in which night-light activity is higher during extreme-temperature years for petroleum-exposed cells than for non-petroleum cells. The data do not directly observe income, fiscal transfers, or public-goods provision, so the precise channel through which oil-price exposure translates into local economic activity remains unidentified; the buffering result, in particular, is a buffering in the measured night-light construct, not directly in household income. On the conflict side, the evidence is consistent with a broad absence of systematic escalation: the disaggregated conflict outcomes do not show a consistent positive pattern across violence types (though non-state and one-sided violence show modest positive coefficients), the price-cycle decomposition does not show a bust-period reversal, and the weather-insurance specification does not show differential attenuation of weather-induced conflict in petroleum cells. These patterns are inconsistent with diffuse-contestation or grievance-driven channels in which oil-price exposure should systematically raise violence at the cell level; they do not, on their own, identify a unique alternative mechanism.

The results are consistent with a transmission channel in which oil operates through one or more of the economic, fiscal, security, and procurement margins, raising local activity without inducing widespread violence on average; the design does not separately identify which of these channels dominates. Where economic gains are larger—as in mixed producers—conflict is lower, consistent with higher opportunity costs of violence and stronger incentives to stabilize economically important regions.

This interpretation speaks to the apparent tension with the cross-country resource-curse literature: while oil-rich countries are more conflict-prone in cross-country comparisons, the within-country-year local response to oil-price exposure is to raise economic activity without generating systematic increases in violence. The paper does not identify what drives the cross-country pattern; the within-country response and the cross-country pattern are different objects.

9 Robustness

9.1 Supply-News and Production-Shortfall Validation

The primary threat to the baseline design is demand-side confounding: global oil prices partly reflect surges in world economic activity that could differentially affect petroleum cells for reasons unrelated to local extraction. To address this, I restrict identifying variation to exogenous supply shocks; the full comparison is reported in Table 11

Table 11: Reduced-form robustness: sub-annual shock timing (ring-excluded sample, 1989–2014/2003)

	Känzig (2021) S_t		**Kilian (2008)** S_t	
	(1) Annual *1989–2014*	**(2) Monthly** *1989–2014*	**(3) Annual** *1989–2003*	**(4) Quarterly** *1989–2003*
Night lights				
$\hat{\beta}$	$+0.0015^{***}$	$+0.0014^{***}$	$+0.0093^{***}$	$+0.0092^{***}$
	(0.0002)	(0.0001)	(0.0010)	(0.0010)
Conflict events				
$\hat{\beta}$	-0.0002	-0.0002	-0.0007	-0.0007
	(0.0002)	(0.0002)	(0.0006)	(0.0006)
Conflict deaths				
$\hat{\beta}$	-0.0004	-0.0004	-0.0022^{*}	-0.0022^{*}
	(0.0003)	(0.0003)	(0.0011)	(0.0011)
Obs (unit)	1,301,257	15,615,084	702,207	2,808,828
Unit	cell×year	cell×month	cell×year	cell×quarter
Cell FE	✓	✓	✓	✓
Country-year FE	✓	✓	✓	✓
Ring cells excluded	✓	✓	✓	✓

Notes: OLS reduced-form estimates of $Y_{gct} = \beta(Oil_g \times S_t) + \alpha_g + \lambda_{ct} + \varepsilon$. All columns use the ring-excluded sample (non-petroleum cells contiguous to a petroleum cell are dropped from the control group). Columns (1) and (3) use annual aggregates of the respective shock series (sum of monthly/quarterly values within each year). Column (2) uses the raw Känzig (2021) monthly surprise S_m directly: the annual panel is expanded to cell × month observations, with annual outcomes repeated for each month; S_m is non-zero only in OPEC announcement months (92 of 312 months, 1989–2014). Column (4) uses Kilian's (2008) quarterly disruption series (panel expanded to cell × quarter; 1989–2003). Country-year FE are identical across annual and sub-annual panels. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

I begin with high-frequency oil supply news constructed from futures price revisions around scheduled OPEC announcements (Känzig, 2021). Using these shocks directly yields a positive and precisely estimated effect on night lights, while conflict coefficients remain small and statistically indistinguishable from zero (Table 11). Estimates based on realized OPEC production shortfalls (Kilian, 2008) support the same qualitative conclusion — supply-driven increases in oil prices raise local economic activity and do not generate systematic conflict escalation — though the Kilian conflict-deaths coefficient is marginally negative rather than statistically zero. The two instruments are constructed from different data and cover non-overlapping sub-periods; the consistency of the sign and direction across them rules out global demand conditions as the primary driver of the baseline results.

The Känzig instrument is constructed at monthly frequency. Estimating the reduced form with the shock entered at its native monthly frequency — holding outcomes at

their annual frequency and absorbing country-year fixed effects — produces coefficients nearly identical to the annual aggregation, and the analogous quarterly check with the Kilian series yields the same pattern. This is a narrow robustness statement: the reduced form is unchanged when the shock is represented at its native (sub-annual) frequency rather than aggregated to the annual frequency of the outcome data. Because the night-lights and UCDP outcomes are themselves annual, the exercise does not address whether the outcome dynamics are aggregation-masked. It does rule out the specific concern that the choice of annual shock construction is mechanically generating the baseline coefficients.

Weak-instrument inference. Aggregating the monthly Känzig surprises to annual frequency yields a first-stage Montiel Olea–Pflueger effective F-statistic below conventional thresholds for 5%-distortion 2SLS inference. I therefore report reduced-form coefficients on $Oil_g \times S_t^K$ as the primary Känzig validation, and reserve the Kilian (2008) 2SLS estimate — where the first-stage relevance is materially stronger over the shorter 1989–2003 sub-period — as the magnitude reference. The reduced form provides a sign-and-direction check that is valid even under weak first stage; weak-IV-robust confidence intervals for the implied 2SLS magnitude under the Känzig instrument, constructed via the Anderson-Rubin (1949) and Andrews-Stock-Sun (2019) procedures, are reported in Appendix F.1. The qualitative conclusions are robust under both inference regimes: development responds positively to supply-driven price movements; conflict does not exhibit systematic escalation.

A remaining concern is that OPEC announcements may proxy for geopolitical developments within member states rather than pure supply shifts. Such non-price effects would naturally show up as a larger conflict response in OPEC than non-OPEC member countries. The data do not show this pattern: Wald tests of equality between OPEC and non-OPEC reduced-form coefficients fail to reject for both conflict events ($p = 0.072$) and conflict deaths ($p = 0.162$), so the geopolitical non-price channel the diagnostic was designed to detect is not present in the conflict outcomes (Table 12). The development response, by contrast, does differ by OPEC status: the night-lights response is concentrated in non-OPEC producers (Wald $p = 0.004$). One plausible reading of this asymmetry is a structural composition difference rather than a violation of the supply-news exclusion restriction: OPEC oil sectors are predominantly state-owned (Saudi Aramco, NIOC, ADNOC, Pemex, Sonangol), which has been argued in the literature to generate weaker local procurement and labour-market spillovers than the

Table 12: OPEC-membership falsification: Känzig supply-news shock by OPEC and non-OPEC country status (ring-excluded sample, 1989–2014)

	(1) *Night lights*	(2) *Conflict events*	(3) *Conflict deaths*
$Oil_g \times S_t^{\text{Känzig}} \times OPEC_c$	$+0.0003$	$+0.0011$	$+0.0013$
	(0.0005)	(0.0008)	(0.0013)
$Oil_g \times S_t^{\text{Känzig}} \times (1 - OPEC_c)$	$+0.0017^{***}$	-0.0004^{**}	-0.0006^{**}
	(0.0002)	(0.0002)	(0.0003)
Wald p: $\beta_{\text{OPEC}} = \beta_{\text{Non-OPEC}}$	[0.004]	[0.072]	[0.162]
N	1,301,257	1,593,551	1,593,551
Cell FE	✓	✓	✓
Country-year FE	✓	✓	✓
Ring cells excl.	✓	✓	✓

Notes: Reduced-form estimates of the Känzig (2021) supply-news shock interacted with petroleum exposure, separately for ever-OPEC and never-OPEC countries. The geopolitical-confound concern that motivates this diagnostic is that OPEC announcements may signal regime instability or security shocks within member states that differentially affect petroleum-adjacent areas through non-price channels (security deployments, insurgent targeting of infrastructure). Such non-price effects would naturally show up as a larger *conflict* response in OPEC than in non-OPEC member countries. The Wald row reports the p-value for $\beta_{\text{OPEC}} = \beta_{\text{Non-OPEC}}$ in each column; the test reads asymmetrically across outcomes. The *conflict* columns do not reject equality (events $p = 0.072$, deaths $p = 0.162$): the geopolitical non-price pattern the diagnostic was designed to detect is not present in the data. The *development* column does reject equality ($p = 0.004$): the night-lights response is concentrated in non-OPEC producers. One plausible (conjectural) reading is a structural composition difference — OPEC oil sectors are predominantly state-owned (Saudi Aramco, NIOC, ADNOC, Pemex, Sonangol), which is sometimes argued in the literature to generate weaker local procurement and labour-market spillovers than the IOC-and-contractor operations more common in non-OPEC — rather than evidence against the supply-news exclusion restriction (which is a conflict-channel concern in this setting). The design does not directly measure procurement or labour-market linkages, so this ownership reading should be treated as one explanation rather than a demonstrated mechanism. OPEC membership is coded as ever-OPEC during 1989–2014 (gwno codes 101, 130, 475, 481, 540, 615, 620, 630, 645, 670, 690, 694, 696, 850). Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

IOC-and-contractor operations that dominate non-OPEC oil production. This explanation is conjectural: the design does not directly measure procurement or local linkages, so the ownership-structure interpretation should be read as one possible account of the OPEC asymmetry rather than a demonstrated mechanism. The diagnostic therefore does not provide a clean across-the-board "passed falsification," but it does rule out the specific geopolitical-confound pattern in the channel where it matters most for the paper's central claim about conflict.

9.2 Oil-Specific vs. Commodity-Cycle Dynamics

A key concern is that the estimated effects reflect a broader commodity cycle rather than oil-specific mechanisms. During the 2002–2008 super-cycle, oil and gold prices rose in tandem, raising the possibility that the Känzig supply shock proxies for general commodity demand conditions rather than oil revenues.

In the primary OLS design, gold-exposed cells also exhibit a significant response of their own during the commodity cycle: night-light intensity rises ($\hat{\beta} = +0.015^{**}$ on $Gold_g \times \ln P_t$) and aggregate conflict measures fall (-0.012^{**} on events; -0.018^{**} on deaths), reported in Appendix Table A6, Panel B. Gold-exposed cells are not silent — the commodity cycle reaches them too — so gold is not a clean null placebo. The substantive identification question is therefore not whether gold responds at all, but whether the oil coefficient is identified from an oil-specific channel or from a shared commodity-demand component. Two specifications give the same answer. In a horse race where both exposures enter jointly (Panel C), the oil coefficient is essentially unchanged in magnitude and precision relative to the oil-only benchmark, and the gold coefficient retains its own pattern; the two exposures are not substitutes for each other and they are estimated independently within country×year cells. Controlling directly for the gold-price interaction $Gold_g \times \ln P_t^{\text{gold}}$ (Panel D) also leaves the oil coefficient stable. The Känzig supply-shock reduced form in Figure 9 reinforces this conclusion under causal identification: oil-supply news — by construction orthogonal to common commodity demand — recovers the same positive oil response, indicating that the oil result is driven by the supply-specific component rather than by a shared commodity-cycle demand effect.

The interpretation is therefore the stable-oil-coefficient interpretation rather than a null-placebo interpretation. Gold exposure has its own commodity-cycle pattern that the design correctly recovers, but oil and gold operate on distinct margins: the oil coefficient is identified separately from the gold pattern in both the horse-race and the

Figure 9: **Oil-specific vs. commodity-cycle dynamics**

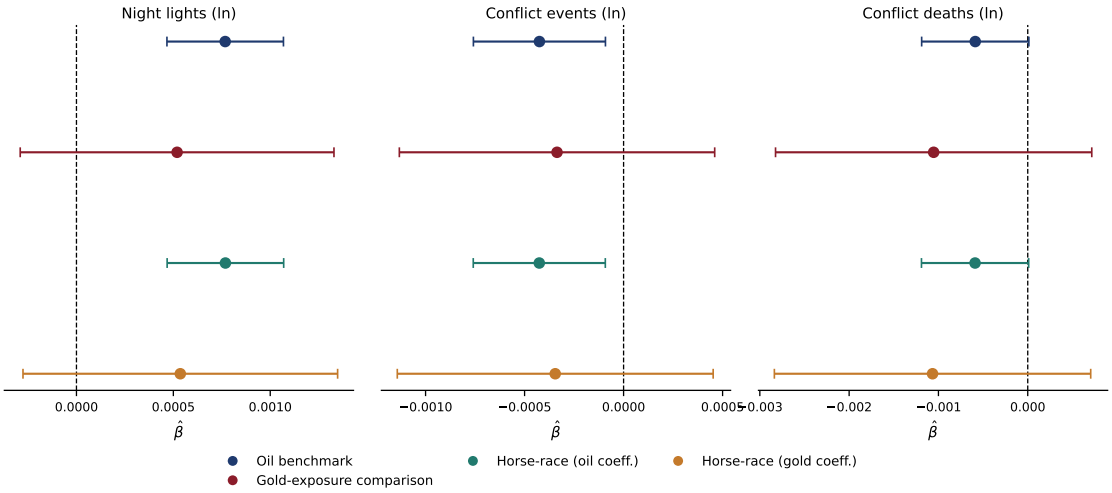

Känzig (2021) reduced-form coefficients ($\hat{\beta}$, 95% CI) for four specifications across three outcomes. Navy circles: oil benchmark ($Oil_g \times S_t^K$); crimson triangles: gold-exposure comparison ($Gold_g \times S_t^K$); teal diamonds / orange squares: horse-race oil and gold coefficients estimated jointly. Gold cells defined as $\mathbf{1}[\texttt{goldplacer\_s} > 0]$ from PRIO-GRID v3, symmetric with oil-cell definition. Cell-clustered SEs.

gold-price-control specifications, and the Känzig supply-shock reduced form recovers the same oil response under causal identification. The main aggregate results therefore reflect petroleum-specific transmission rather than a shared commodity-demand cycle, even though gold itself is not a quiet placebo at the commodity-cycle margin.

9.3 Clustering, Alternative Proxies, and Functional Form

Clustering. The baseline clusters standard errors at the PRIO-GRID cell level. Table A2 (Appendix C) reports two-way (cell×year) clustered alternatives for both the primary design and the pre-1989 ICRG state-capacity interaction. With only $T = 26$ year-clusters, two-way SEs are conservative but asymptotically valid; the table shows that the ring-excluded sample findings are stable under this more demanding variance estimator.

Alternative development proxies. The night-lights measure from DMSP-OLS may saturate in brightly lit areas and misses purely economic activity not accompanied by artificial lighting. The PRIO-GRID local GDP proxy (a purchasing-power-parity adjusted grid-cell product from Kummu et al. 2018) provides an alternative. The development coefficient under this proxy is positive and statistically significant, confirming that oil windfalls raise real economic activity beyond the lighting channel.

9.4 ACLED Robustness: Independent Conflict Data Source

A concern with the conflict null is that it reflects measurement choices in UCDP GED rather than a true absence of conflict responses. UCDP focuses on organised armed conflict and may miss lower-intensity violence or events involving non-state actors. To address this, I replicate the analysis using ACLED (1997–2014), which draws on different sources and applies broader event definitions.

Panel A of Table 13 shows no evidence of escalation. Across all aggregate outcomes, the estimates are close to zero and statistically insignificant. Oil price exposure does not increase the probability of conflict, violence, deaths, or sexual violence in petroleum cells. The shorter ACLED sample reduces precision relative to UCDP, but the estimates rule out effects of economically meaningful magnitude. Results are unchanged across alternative specifications reported in Table 13.

Panel B reveals a more granular pattern. State violence is unaffected, consistent with Table 8. By contrast, non-state violence and violence against civilians increase slightly at the extensive margin. These effects are small but precisely estimated, and they mirror the corresponding UCDP patterns.

Overall, the ACLED evidence sharpens the interpretation of the conflict response in a way that should not be buried as a footnote. The aggregate UCDP and ACLED null is not a flat absence of conflict effects; it is the average of an unchanged organised-violence margin and a modest compositional shift toward low-intensity, non-state and civilian-targeted violence. State-based violence and aggregate event counts are unaffected by oil-price exposure (UCDP: $\hat{\beta} = -0.004$ on state-based events, -0.002 on one-sided ACLED state violence; both indistinguishable from zero). Non-state violence ($\hat{\beta} = +0.0015^{*}$ in UCDP, $+0.0015^{**}$ in ACLED) and violence against civilians in ACLED ($\hat{\beta} = +0.0014^{**}$) show small positive responses that survive a second independent dataset with different sources and broader event definitions. The agreement across two datasets with non-overlapping source pools makes it unlikely that the pattern reflects measurement or coding choices: oil booms do not raise aggregate organised violence at the cell level, but they coincide with a small reallocation toward less-organised forms of violence.

Table 13: ACLED disaggregated conflict response: overall outcomes and actor decomposition (1997–2014)

	$\hat{\beta}$ ($\text{Oil}_g \times \ln P_t$)
Panel A: Overall ACLED outcomes	
(binary indicators except for the deaths outcome, which is logged)	
Any event	$+0.0009$
	(0.0007)
Any violence	$+0.0012$
	(0.0008)
Log fatalities	$+0.0009$
	(0.0018)
Sexual violence	$+0.0000$
	(0.0001)
Panel B: Actor and event-type disaggregation	
(all binary indicators; cf. UCDP Table 8)	
State violence	-0.0002
	(0.0004)
Non-state violence	$+0.0015^{**}$
	(0.0006)
Violence vs. civ.	$+0.0014^{**}$
	(0.0005)
Protests & riots	$+0.0003$
	(0.0006)
Cell FE	✓
Country-year FE	✓
N	1,103,238

Notes: OLS (LPM) estimates of $\hat{\beta}$ from $Y_{gct} = \beta(Oil_g \times \ln P_t) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$. Sample: ACLED 1997–2014, ring-adjacent cells excluded. *Panel A* uses outcomes defined consistently with the UCDP GED analysis: any event (binary), any violence (binary, excluding peaceful demonstrations), log fatalities (the only outcome with meaningful count variation: median positive = 7), and sexual violence (binary). *Panel B* disaggregates by actor and event type: state violence = primary actor is State forces and event type is Battles, Violence against civilians, or Explosions/Remote violence; non-state violence = same event types, primary actor is Rebel group, Political militia, or Identity militia; violence vs. civilians = `event_type = Violence against civilians` regardless of actor; protests & riots = `event_type` in {Protests, Riots}. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

10 Conclusion

This paper estimates the local response of economic activity and armed conflict to oil-price exposure in petroleum-bearing cells, using a global panel of 0.5-degree grid cells across 195 countries from 1989 to 2014 (within-country identification from the 74 countries with both petroleum and non-petroleum cells; see Table 1). Three results emerge.

Oil-price exposure raises local economic activity in petroleum cells. Over the 1998–2008 price cycle, petroleum cells brighten by approximately eight percent relative to non-petroleum cells in the same country and year — a substantial differential, robust to alternative controls, lagged-price specifications, and supply-shock identification using OPEC announcement surprises. The local conflict response, by contrast, is small in aggregate but not flat. Aggregate UCDP and ACLED estimates return small point estimates whose confidence intervals exclude economically meaningful escalation of organised violence, but the disaggregated picture is more substantively informative: state-based violence is unaffected, while non-state violence and violence against civilians show small positive responses across both datasets ($\hat{\beta} \approx +0.0015$ in log-counts; significant at conventional levels). Oil-price exposure therefore does not escalate organised violence in the average petroleum cell, but it produces a modest reallocation toward lower-intensity, non-state and civilian-targeted forms; the aggregate null is not a flat absence of conflict response.

The development response is moderated on two independent institutional margins: pre-sample state capacity and pre-sample ethnic political exclusion. Each margin pulls the local response down on its own; both together extinguish it almost entirely. Petroleum cells in countries with strong administrative capacity and politically included populations receive the full local economic premium; cells in countries with weak administrative capacity and politically excluded populations receive essentially none of it. This is heterogeneity by institutional category, not a causal effect of varying institutions themselves — the design does not exploit exogenous variation in institutions — but the joint pattern is the empirical signature of a dual-margin institutional mechanism in which the state must be able to deliver public goods (capacity) and the local population must be able to claim them (political access).

These results characterise one of the two estimands at issue in the resource-curse debate: the within-country differential response of local economic activity and conflict to oil-price exposure under common national conditions. They do not adjudicate the country-level resource-curse correlation, which is a different empirical object that the design absorbs rather than measures. What the within-country evidence shows is that

local petroleum exposure, conditional on common national shocks, does not behave the way the country-level correlation would suggest: the cell-level response is positive on the development margin and statistically indistinguishable from a uniform increase in conflict on the average margin. Whether the country-level correlation reflects between-country confounds, common-shock co-movement, or something else, this design cannot say — it can only describe what happens at the cell level. Country-level and single-country findings reflect specific combinations of institutional environments and conflict histories and need not generalise mechanically to a universal local relationship between oil prices and violence.

The deeper substantive implication concerns development, not conflict. Whether oil wealth translates into local prosperity depends on whether the state can convert revenues into public goods and whether the local population has formal political access to those revenues. The night-lights evidence here speaks to the first link in this chain — whether local economic activity responds to oil-price exposure — and shows that the response is sharply conditional on institutional context. Whether the night-lights response generalises to direct welfare outcomes (household income, public-goods provision, livelihoods) is a separate question that the cell-year outcomes used here cannot answer; the institutional moderation reported on the development margin is a strong constraint on the policy environment under which oil wealth reaches local economies, but it is not a structural estimate of the welfare effect of resource exposure.

References

Adhvaryu, A., J. Fenske, G. Khanna, and A. Nyshadham (2021). Resources, conflict, and economic development in Africa. *Journal of Development Economics 149*.

Andersen, J. J., F. M. Nordvik, and A. Tesei (2022). Oil price shocks and conflict escalation: Onshore versus offshore. *Journal of Conflict Resolution 66*(4–5), 786–809.

Atlantic Council (2020). How oil markets reacted to the soleimani strike. https://www.atlanticcouncil.org/blogs/energysource/.

Bazzi, S. and C. Blattman (2014). Economic shocks and conflict: Evidence from commodity prices. *American Economic Journal: Macroeconomics 6*(4), 1–38.

Bellemare, M. F. and C. J. Wichman (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics 82*(1), 50–61.

Berman, N., M. Couttenier, D. Rohner, and M. Thoenig (2017). This mine is mine! how minerals fuel conflicts in africa. *American Economic Review 107*(6), 1564–1610.

Besley, T. and T. Persson (2010). State capacity, conflict, and development. *Econometrica 78*(1), 1–34.

Bitoto, F.-E. (2026). Natural resources and conflict: A survey. Working paper.

Brückner, M. and A. Ciccone (2010). International commodity prices, growth and the outbreak of civil war in sub-saharan africa. *Economic Journal 120*(544), 519–534.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics 29*(2), 238–249.

Caselli, F. and G. Michaels (2013). Do oil windfalls improve living standards? evidence from Brazil. *American Economic Journal: Applied Economics 5*(1), 208–238.

Cederman, L.-E., A. Wimmer, and B. Min (2010). Why do ethnic groups rebel? New data and analysis. *World Politics 62*(1), 87–119.

Collier, P. and A. Hoeffler (2004). Greed and grievance in civil war. *Oxford Economic Papers 56*(4), 563–595.

Coppedge, M., J. Gerring, C. H. Knutsen, et al. (2025). V-Dem Codebook v15. Technical report, Varieties of Democracy (V-Dem) Project, University of Gothenburg.

Correia, S., P. Guimarães, and T. Zylkin (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal 20*(1), 95–115.

Cotet, A. M. and K. K. Tsui (2013). Oil and conflict: What does the cross-country evidence really show? *American Economic Journal: Macroeconomics 5*(1), 49–80.

Dal Bó, E. and P. Dal Bó (2011). Workers, warriors, and criminals: Social conflict in general equilibrium. *Journal of the European Economic Association 9*(4), 646–677.

Davies, S., T. Pettersson, M. Sollenberg, and M. Öberg (2025). Organized violence 1989–2024, and the challenges of identifying civilian victims. *Journal of Peace Research 62*(4).

Dube, O. and J. F. Vargas (2013). Commodity price shocks and civil conflict: Evidence from colombia. *The Review of Economic Studies 80*(4), 1384–1421.

Esteban, J. and D. Ray (2008). On the salience of ethnic conflict. *American Economic Review 98*(5), 2185–2202.

Fearon, J. D. (2005). Primary commodity exports and civil war. *Journal of Conflict Resolution 49*(4), 483–507.

Frost-Killian, S., S. Master, R. P. Viljoen, and M. G. C. Wilson (2016). The great mineral fields of africa: Introduction. *Episodes 39*(2), 85–103.

Girard, V., T. Molina-Millán, and N. Vic (2025). Artisanal mining in africa: Green for gold? *Economic Journal*.

Hunziker, P. and L.-E. Cederman (2017). No extraction without representation: The ethno-regional oil curse and secessionist conflict. *Journal of Peace Research 54*(3), 365–381.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review 111*(4), 1092–1125.

Kilian, L. (2008). Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. economy? *The Review of Economics and Statistics 90*(2), 216–240.

Kummu, M., M. Taka, and J. H. A. Guillaume (2018). Gridded global datasets for gross domestic product and human development index over 1990–2015. *Scientific Data 5*, 180004.

Lei, Y.-H. and G. Michaels (2014). Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics 110*, 139–157.

Lujala, P., J. K. Rod, and N. Thieme (2007). Fighting over oil: Introducing a new dataset. *Conflict Management and Peace Science 24*(3), 239–256.

Maystadt, J.-F. and O. Ecker (2014). Extreme weather and civil war: Does drought fuel conflict in somalia through livestock price shocks? *American Journal of Agricultural Economics 96*(4), 1157–1182.

Mehlum, H., K. Moene, and R. Torvik (2006). Institutions and the resource curse. *Economic Journal 116*(508), 1–20.

Michalopoulos, S. and E. Papaioannou (2016). The long-run effects of the scramble for Africa. *American Economic Review 106*(7), 1802–1848.

Morelli, M. and D. Rohner (2015). Resource concentration and civil wars. *Journal of Development Economics 117*, 32–47.

Müller-Crepon, C., Y. Pengl, and N.-C. Bormann (2023). Linking ethnic data from Africa. Technical report, CESifo Working Paper No. 10207.

Natural Resource Governance Institute (2013). The 2013 Resource Governance Index. Natural Resource Governance Institute, New York. Composite scores for 58 resource-rich countries, covering institutional and legal setting, reporting practices, safeguards and quality controls, and enabling environment.

O'Brochta, W. (2019). A meta-analysis of natural resources and conflict. *Research & Politics 6*(1).

Reuters (2020). Oil prices surge after killing of iranian commander soleimani. Reuters news report, January 3, 2020.

Reuters (2026a). Hormuz disruption pushes oil prices higher. Reuters news report.

Reuters (2026b). Iran-related oil supply concerns drive market. Reuters news report.

Robinson, J. A., R. Torvik, and T. Verdier (2006). Political foundations of the resource curse. *Journal of Development Economics 79*(2), 447–468.

Ross, M. L. (2004). How do natural resources influence civil war? evidence from thirteen cases. *International Organization 58*(1), 35–67.

Sachs, J. D. and A. M. Warner (1995). Natural resource abundance and economic growth. Technical report, NBER Working Paper No. 5398.

Sundberg, R. and E. Melander (2013). Introducing the UCDP georeferenced event dataset. *Journal of Peace Research 50*(4), 523–532.

Teorell, J., A. Sundström, S. Holmberg, B. Rothstein, N. Alvarado Pachon, V. S. Phiri, C. Chen, and Z. Liu (2026). The quality of government standard dataset, version jan26.

Thiéblemont, D. (2016). Geological map of africa at 1:10,000,000 scale. Technical report, BRGM (Bureau de Recherches Géologiques et Minières). Commission for the Geological Map of the World (CGMW).

Tollefsen, A. F., H. Strand, and H. Buhaug (2012). Prio-grid: A unified spatial data structure. *Journal of Peace Research 49*(2), 363–374.

U.S. Energy Information Administration (2024). Red sea shipping disruptions and oil markets. EIA short-term outlook, January 2024.

Vesco, P., S. Dasgupta, E. De Cian, and C. Carraro (2020). Natural resources and conflict: A meta-analysis of the empirical literature. *Ecological Economics 172*.

Vogt, M., N.-C. Bormann, S. Rüegger, L.-E. Cederman, P. Hunziker, and L. Girardin (2015). Integrating data on ethnicity, geography, and conflict: The Ethnic Power Relations dataset family. *Journal of Conflict Resolution 59*(7), 1327–1342.

World Bank (2024). Commodity markets outlook 2024. World Bank, Washington DC.

Appendix

Figure A1: Känzig (2021) oil supply news surprises and Kilian (2008) OPEC supply shocks

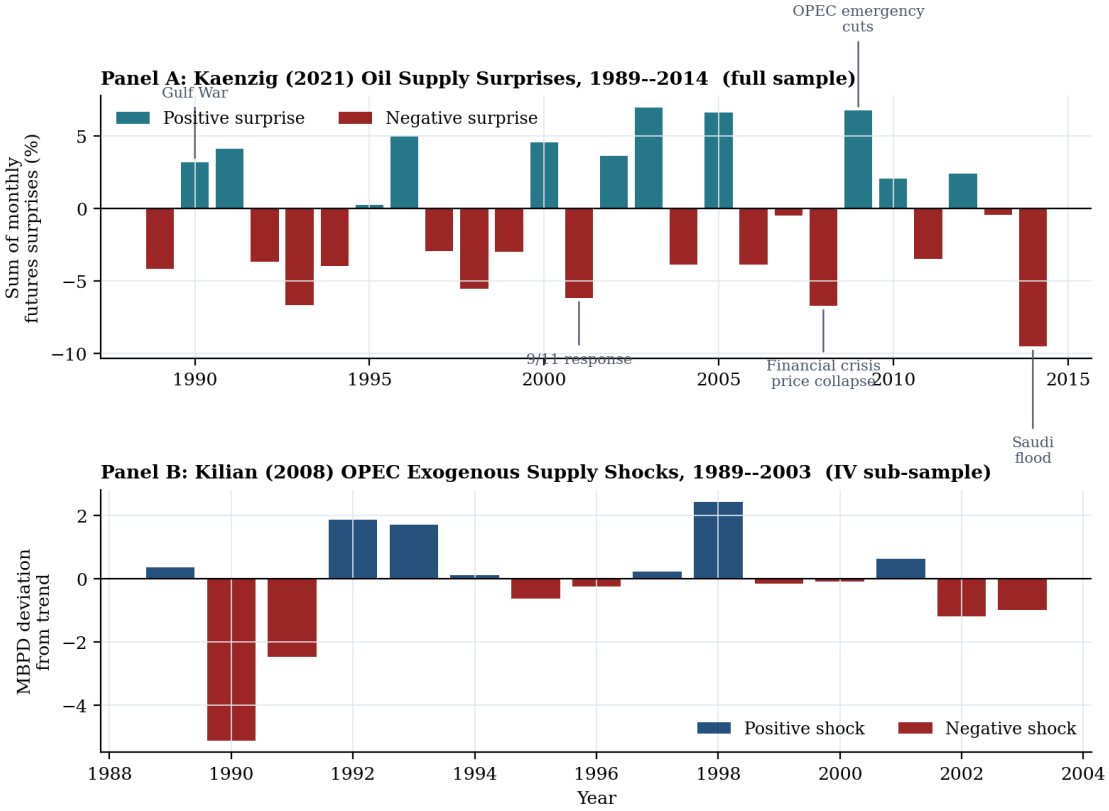

Panel A plots the annual sum of monthly oil supply surprises from Känzig (2021), constructed as oil futures price changes in a one-day window around OPEC announcements (vintage 2017M12). Each bar represents one year in the full analysis sample (1989–2014). Panel B plots the Kilian (2008) annual OPEC exogenous supply shocks over the overlapping period (1989–2003). Kilian measures realized production shortfalls; Känzig measures expectational revisions around announcement events.

Figure A2: Raw trends in oil and non-oil cells, 1989–2014

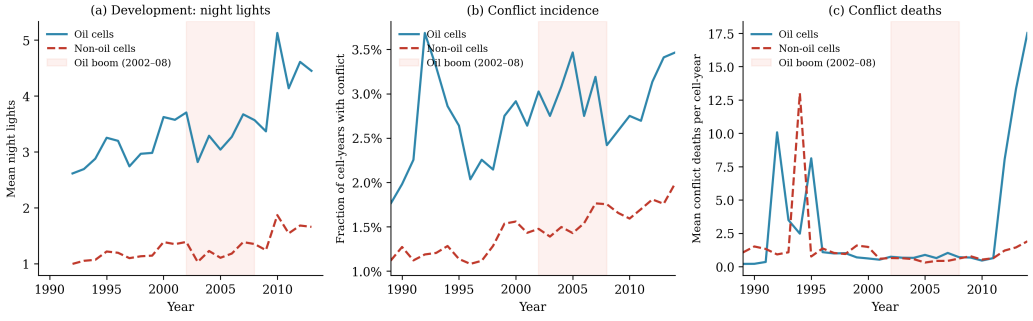

Note: Panels show unconditional cross-cell means for night lights, conflict events, and conflict deaths. Petroleum and non-petroleum cells are compared without any fixed-effect adjustment. Shaded band = 2002–2008 oil-price acceleration. Source: PRIO-GRID yearly file and UCDP GED v25.1. The within-cell-deviation version is shown in Figure A3; the year-by-year coefficients under the baseline FE structure are in Appendix E.2.

Figure A3: Within-cell deviations from cell means: oil vs. non-oil cells, 1989–2014.

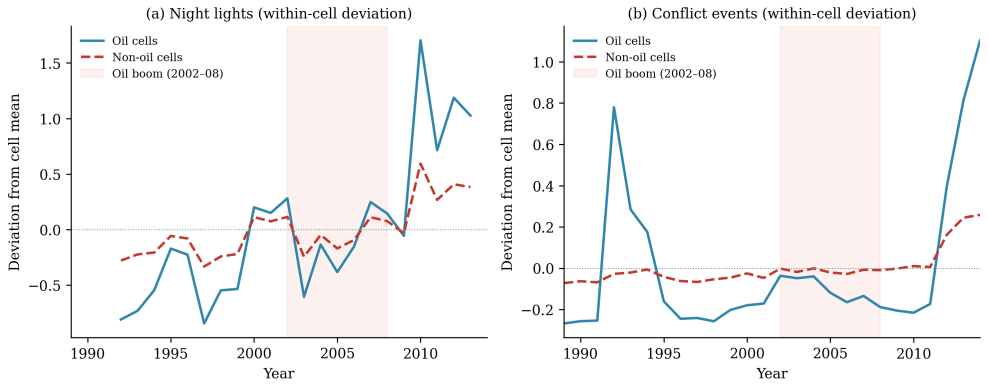

Note: Each series is the cross-cell average of (outcome − cell mean). Cell fixed effects are removed; country-year effects are *not yet* absorbed, so the series still contain cross-country and cross-year variation. Co-movement during the lower-price 1990s and a differential response of petroleum cells during the 2002–2008 acceleration. Treated as a descriptive co-movement complement to the main-text Figure 3 (which absorbs both cell and country×year fixed effects). Shaded band = 2002–2008 oil-price acceleration.

A Split-Sample Visualisation of the Institutional Heterogeneity

This appendix reports the split-sample visualisation of the institutional-moderation results in Section 7.1. For each of four institutional dimensions — ICRG state capacity (1984–1988 average), V-Dem polyarchy (1970–1988), Polity2 (1970–1988), and log GDP per capita (1985–1988) — I estimate the primary OLS regression $Y_{gct} = \beta\,(Oil_g \times \ln P_t) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$ separately on the below-median and above-median country subsamples, and plot the resulting $\hat{\beta}$ values. Where Table 5 reports the triple-interaction moderation coefficient $\hat{\theta}$, this figure plots the level of $\hat{\beta}$ within each subgroup, so the magnitude of the institutional gap is visible directly.

The visualisation reveals two facts that complement the main-text discussion. First, the development response is sharply moderated by state capacity (ICRG), with the high-capacity subgroup retaining a strong positive $\hat{\beta}$ while the low-capacity subgroup's coefficient is attenuated. Second, the conflict response is small and not systematically heterogeneous across any of the four institutional dimensions: the low-moderator and high-moderator coefficients overlap for both events and deaths in essentially every row of Figure A4. The split-sample evidence is therefore consistent with the triple-interaction estimates in Table 5, where the development $\hat{\theta}$ is precisely estimated under ICRG but the conflict $\hat{\theta}$ is small and statistically zero across all four moderators.

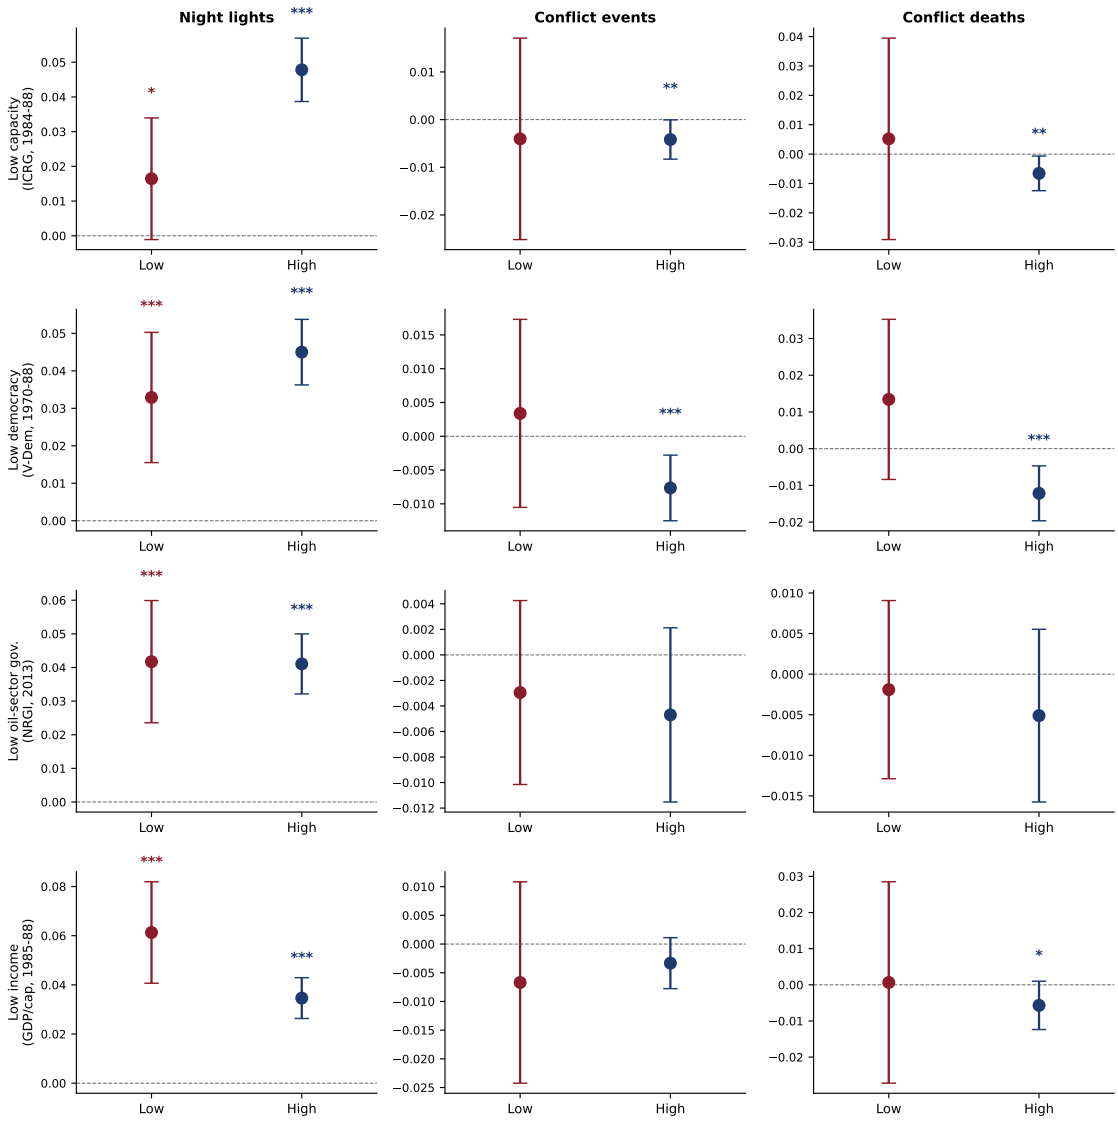

Figure A4: Competing institutional channels: split-sample of $Oil_g \times \ln P_t$. Each panel reports $\hat{\beta}$ on the low-moderator subsample (crimson) and the high-moderator subsample (navy) for one institutional dimension. Rows: low state capacity (ICRG, 1984–1988), low democracy (V-Dem polyarchy, 1970–1988), low Polity2 (1970–1988), low income (log GDP per capita, 1985–1988). Columns: night lights, conflict events, conflict deaths. Each $\hat{\beta}$ is from a separate primary-OLS regression with cell and country×year fixed effects; standard errors clustered at the PRIO-GRID cell level; ring-excluded sample. Bars are 95% CIs; stars indicate $p < 0.01^{***}$, $p < 0.05^{**}$, $p < 0.10^{*}$. Where Table 5 reports the moderation coefficient $\hat{\theta}$, this figure plots the levels of $\hat{\beta}$ within each subgroup.

B Event-Study Diagnostics by Institutional Subgroup

This appendix reports event-study analogs of the institutional-heterogeneity results in Section 7.1 and Section 7.2. Three specifications are estimated on the ring-excluded sample, with 2001 as the reference year:

1º **By state capacity:** $Y_{gct} = \sum_{k \neq 2001} \beta_k \left(Oil_g \times \mathbb{1}\{t = k\}\right) + \sum_{k \neq 2001} \theta_k \left(Oil_g \times \mathbb{1}\{t = k\} \times LowCap_c\right) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$ (Figure A6).

2º **By ethnic exclusion:** symmetric, with $Excl_g$ in place of $LowCap_c$ (Figure A7).

3º **Four-cell dual-margin partition:** $Y_{gct} = \sum_s \sum_{k \neq 2001} \beta_k^{(s)} \left(Oil_g \times \mathbb{1}\{t = k\} \times \mathbb{1}\{g \in s\}\right) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$ with $s \in \{\text{HC-I}, \text{HC-E}, \text{LC-I}, \text{LC-E}\}$ (Figure A5).

What this diagnostic does and does not address. The exercise speaks directly to the referee question: "were high-capacity included oil cells already on steeper trajectories during earlier low-price years, or is the heterogeneity picking up broader pre-existing differences in oil-cell composition across institutional environments?" Three findings emerge.

Night-lights subgroup gradient: descriptive illustration. Averaging the year-by-year coefficients within the pre-boom (1990–2000), boom (2002–2008), and post-boom (2009–2014) windows, the pre-to-post swing on night lights is +0.121 for HC-I cells, +0.069 for HC-E, +0.043 for LC-I, and +0.004 for LC-E — monotonic in institutional quality and visually consistent with the dual-margin moderation documented in Section 7.1. The pattern is descriptive rather than confirmatory: all four night-lights subgroups fail the parallel-trends test, so the night-lights event-study is not a clean inferential exercise. The development claim rests on the continuous-treatment OLS interaction in Section 7.1 (Table 5) and the supply-shock IV in Appendix G.1. The night-lights row of Table 7 (in Section 7.2) should be read in the same spirit.

Conflict response is concentrated in LC-E cells during the boom. The boom-period (2002–2008) averages of the four-cell coefficients are striking: for conflict events, HC-I $= -0.006$, HC-E $= +0.009$, LC-I $= +0.005$, **LC-E** $= +0.084$; for conflict deaths, HC-I $= -0.005$, HC-E $= +0.012$, LC-I $= +0.002$, **LC-E** $= +0.129$. The pooled conflict null in Section 6 masks a sharp boom-period spike in the worst-institution subgroup, which is the canonical resource-curse prediction localised to the cells where it should hold.

Pre-trend tests. Joint Wald tests that the pre-2001 coefficients are zero are reported in Table A1. For night lights, all four subgroups fail the pre-trend test — consistent

with the known DMSP-OLS satellite-calibration trend documented in Section 9 and the reason the paper's primary identification rests on the continuous-treatment design $Oil_g \times \ln P_t$ rather than on a discontinuity around 2002. For conflict outcomes, three of four subgroups (HC-I, HC-E, LC-E) pass the parallel-trends test; only LC-I cells exhibit pre-trend movement that is statistically distinguishable from zero. The visually striking LC-E conflict spike is therefore *not* contaminated by pre-trends in that subgroup. The LC-I pre-trend caveat is a real flag but the LC-I boom coefficient is small, so it does not drive any conclusion that the paper rests on.

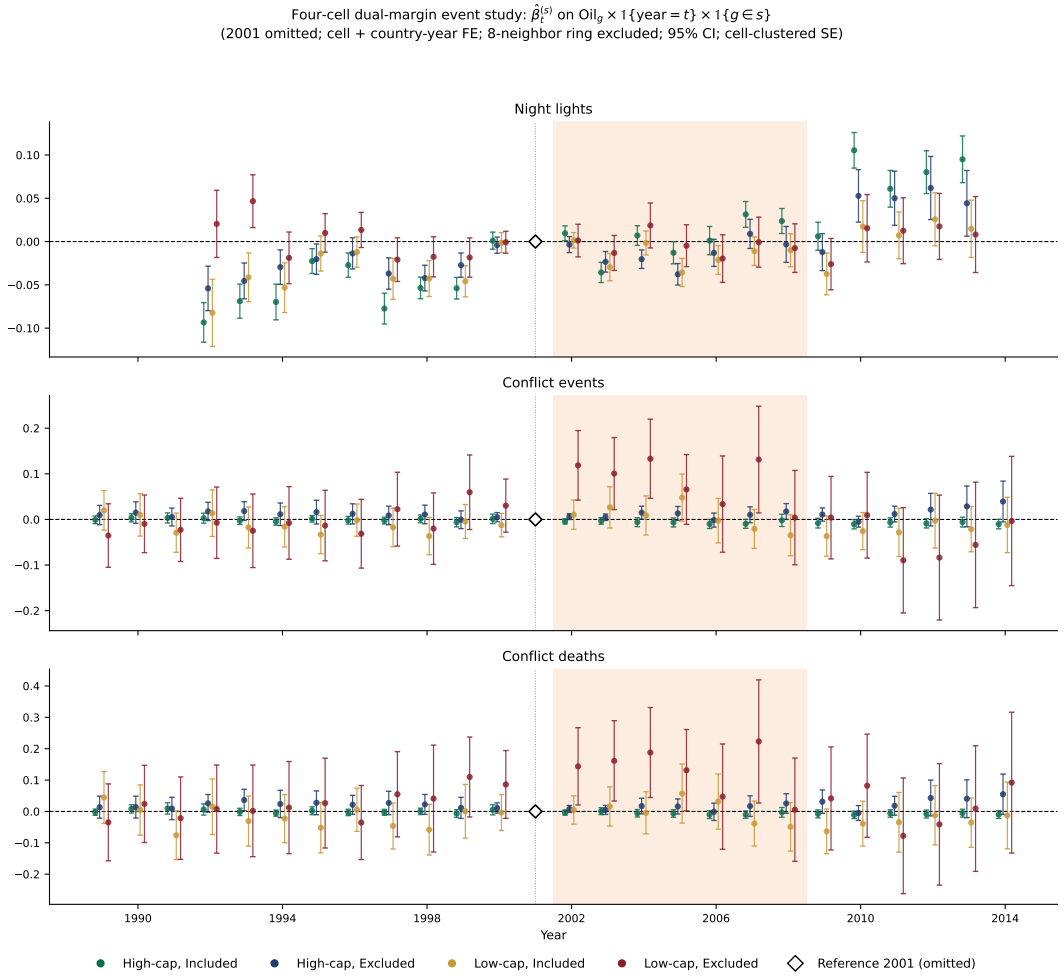

Figure A5: Four-cell dual-margin event study. Each panel plots $\hat{\beta}_k^{(s)}$ on $Oil_g \times \mathbb{1}\{\text{year}=k\} \times \mathbb{1}\{g \in s\}$ with 2001 omitted, for $s \in \{\text{HC-I, HC-E, LC-I, LC-E}\}$. HC/LC: above/below median pre-1989 ICRG quality of governance (country level). I/E: cell with/without a pre-1989 (1980–1988) politically excluded EPR group settled in it. Bars: 95% CIs. Cell + country×year FE; 8-neighbor ring-excluded sample; cell-clustered SEs. Shaded band: 2002–2008 oil-price acceleration.

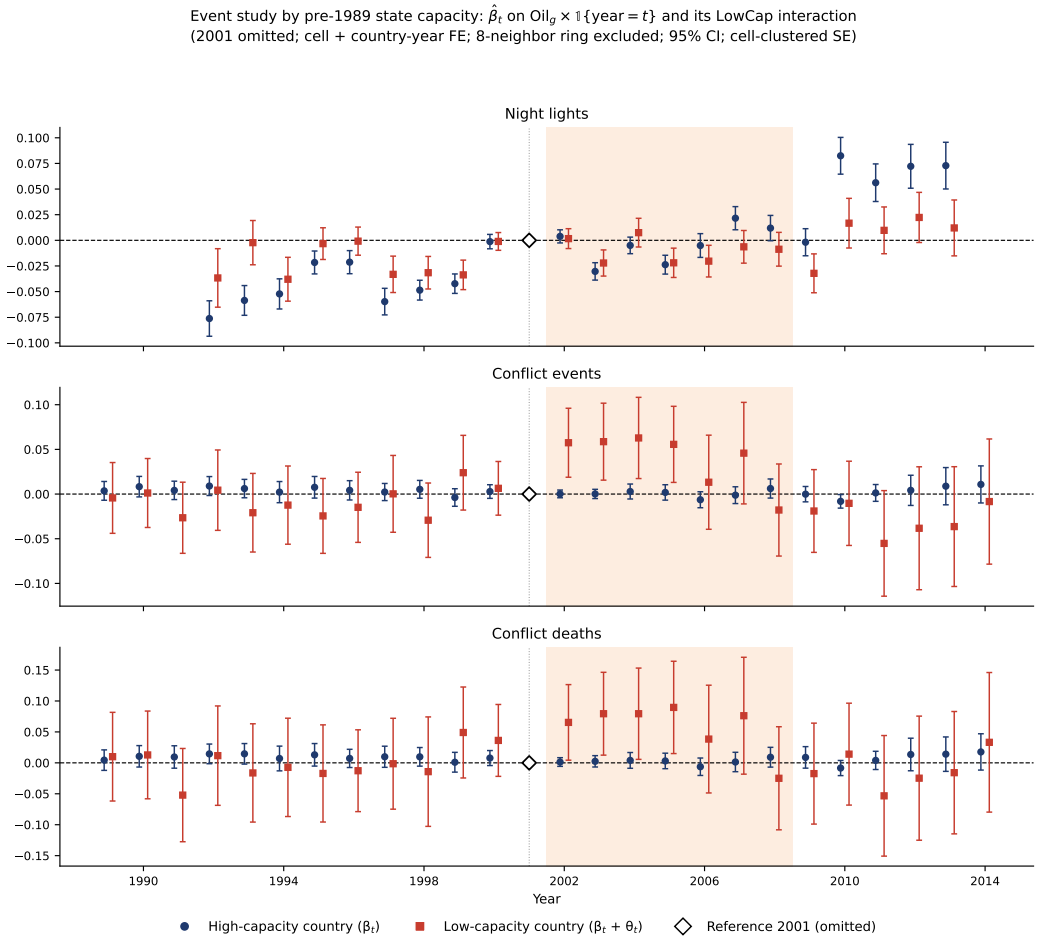

Figure A6: State-capacity event-study analog. High-capacity series: $\hat{\beta}_k$ (baseline). Low-capacity series: $\hat{\beta}_k + \hat{\theta}_k$, with standard errors propagated through the full coefficient covariance. Cell + country×year FE; 8-neighbor ring-excluded sample; cell-clustered SEs. 95% CIs.

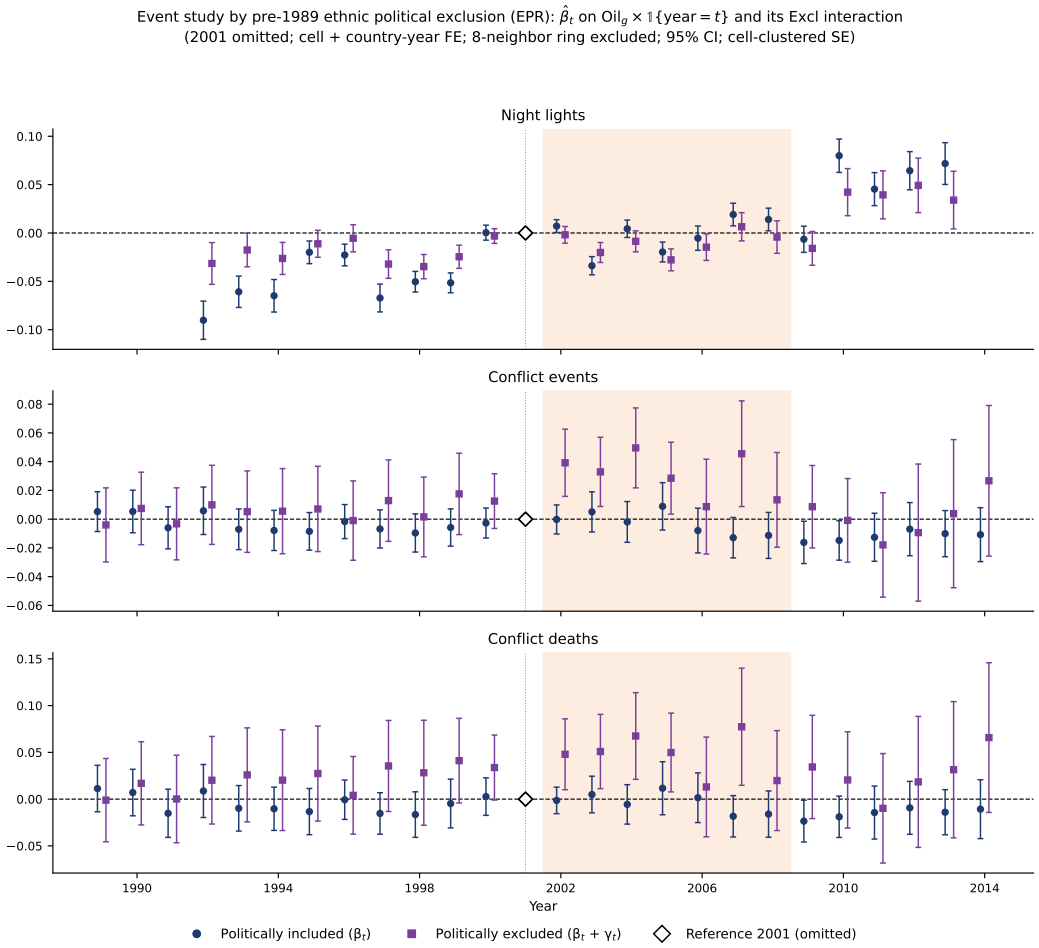

Figure A7: Ethnic-exclusion event-study analog. Politically-included series: $\hat{\beta}_k$ (baseline). Politically-excluded series: $\hat{\beta}_k + \hat{\gamma}_k$. Cell + country×year FE; 8-neighbor ring-excluded sample; cell-clustered SEs. 95% CIs.

Table A1: Joint Wald tests of parallel pre-trends (pre-2001 coefficients = 0) for the four-cell partition event study

	Night lights		*Conflict events*		*Conflict deaths*	
Subgroup	χ^2	p	χ^2	p	χ^2	p
HC-I (high-cap, included)	130.1	0.000	14.5	0.273	18.4	0.104
HC-E (high-cap, excluded)	49.5	0.000	12.7	0.394	12.0	0.447
LC-I (low-cap, included)	36.1	0.000	41.3	0.000	41.9	0.000
LC-E (low-cap, excluded)	46.7	0.000	13.0	0.371	12.5	0.406
df	9		12		12	

Notes: Joint Wald tests of the null that the subgroup-specific event-study coefficients are zero for all pre-reference years (pre-2001). Night lights: pre-period is 1992–2000 (DMSP-OLS coverage begins 1992); conflict outcomes: pre-period is 1989–2000. All night-lights subgroups fail the parallel-trends test; this is the same DMSP-OLS calibration drift that motivates the continuous-treatment identification strategy and is documented in Section 9. Three of four conflict subgroups pass; LC-I (low-capacity, politically-included) fails for both conflict outcomes and is the only flagged subgroup. The boom-period (2002–2008) conflict coefficient in LC-I is small (+0.005 events; +0.002 deaths), so the LC-I pre-trend caveat does not drive any conclusion in the main text. Cell + country×year FE; 8-neighbor ring-excluded sample; cell-clustered SEs.

C Clustering Robustness

Table A2: Clustering robustness: alternative standard errors (primary OLS design)

	Night lights	*Conflict events*	*Conflict deaths*
Panel A: Main specification			
$Oil_g \times \ln P_t$	$+0.0407^{***}$	-0.0044	-0.0049
SE (cell-clustered)	(0.0040)	(0.0028)	(0.0044)
SE (country-clustered)	(0.0082)	(0.0049)	(0.0080)
SE (two-way: cell + year)	(0.0069)	(0.0030)	(0.0047)
Observations	1,301,257	1,593,551	1,593,551
Countries (clusters)		193	
Panel B: State-capacity interaction (ICRG pre-89 sub-sample)			
$Oil_g \times \ln P_t$	$+0.0523^{***}$	-0.0008	-0.0009
SE (cell-clustered)	(0.0059)	(0.0026)	(0.0037)
SE (country-clustered)	(0.0096)	(0.0027)	(0.0031)
SE (two-way: cell + year)	(0.0090)	(0.0027)	(0.0038)
$Oil_g \times LowCap_g \times \ln P_t$	-0.0386^{***}	-0.0045	$+0.0022$
SE (cell-clustered)	(0.0095)	(0.0113)	(0.0182)
SE (country-clustered)	(0.0208)	(0.0207)	(0.0353)
SE (two-way: cell + year)	(0.0102)	(0.0133)	(0.0195)
Observations	924,186	1,128,107	1,128,107
Countries (clusters)		121	
Cell FE	✓	✓	✓
Country×Year FE	✓	✓	✓
Ring cells excl.	✓	✓	✓

Notes: OLS estimates with cell and country×year fixed effects on the 8-neighbor ring-excluded sample. Point estimates are identical across the three standard-error rows within each panel and are reported once; the three SE rows give: (i) cell-clustered (preferred, used throughout the paper); (ii) country-clustered (treats each country as one cluster — more conservative under cross-cell spatial correlation within country); (iii) two-way cell + year clustering (Cameron et al., 2011) (allows common year shocks). Panel B restricts to country-years with non-missing pre-1989 ICRG quality-of-government data and adds the triple interaction $Oil_g \times LowCap_g \times \ln P_t$, where $LowCap_g$ is the below-median indicator on the 1984–1988 ICRG country average. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A2 replicates the primary design (Panel A) and the pre-1989 ICRG state-capacity interaction (Panel B) under two standard-error estimators: cell-level clustering (the baseline used throughout the paper) and two-way clustering simultaneously at the cell and year level (Cameron et al., 2011). The ring-excluded sample is used throughout. Two-way SEs are conservative with $T = 26$ year clusters but remain valid asymptotically. Across all outcomes and both panels, substantive conclusions are unchanged: the development coefficient is positive and significant; conflict coefficients are small and insignificant.

C.1 Zero-Inflation Robustness: arcsinh and PPML

Conflict outcome distributions contain a high share of zeros (most cell-years record no events), which raises the question of whether the $\ln(1+Y)$ transformation adequately handles the mass at zero. Table A3 addresses this across six specifications and two estimators. Panel A uses the inverse-hyperbolic-sine transformation (Bellemare and Wichman, 2020); Panel B estimates a Poisson QMLE (PPML) on raw conflict-death counts, which is consistent under both over-dispersion and zero-inflation (Correia et al., 2020).

Table A3: **Zero-inflation robustness: asinh OLS (Panel A) and PPML (Panel B)**

	ln oil price		Känzig (2021)		Kilian (2008)	
	(1) Main	(2) +ICRG	(3) RF base	(4) RF+ICRG	(5) RF base	(6) RF+ICRG
Panel A: Y = asinh(conflict deaths), OLS						
Oil × Shock	-0.0050	-0.0167	-0.0004	-0.0007	-0.0024*	0.0010
	(0.0050)	(0.0345)	(0.0004)	(0.0016)	(0.0014)	(0.0008)
× Low ICRG (pre-1989)		0.0467		-0.0002		-0.0006
		(0.0394)		(0.0028)		(0.0008)
Observations	1,593,414	283,344	1,593,414	283,344	919,275	162,635
Panel B: Y = conflict deaths (count), PPML						
Oil × Shock	0.5147	0.8129	−0.0358***	−0.0223	−0.0012	−0.0485
	(0.2713)	(0.4558)	(0.0122)	(0.0312)	(0.1526)	(0.0920)
× Low ICRG (pre-1989)		−0.3583		−0.0200		−0.1237
		(0.5862)		(0.0331)		(0.1778)
Observations (PPML)	103,892	78,849	103,892	78,849	41,878	32,347
Cell FE				✓		
Country×Year FE				✓		
Ring cells excluded				✓		
Sample		1989–2014			1989–2003	

Notes: Panel A estimates OLS on asinh(conflict deaths)$_{gct}$. Panel B estimates Poisson pseudo-MLE on raw conflict-death counts using the `ppmlhdfe` algorithm (Correia et al., 2020), which is consistent under arbitrary over-dispersion and zero-inflation. Columns (1)–(2) use the main oil-price interaction ($Oil_g \times \ln P_t$); columns (3)–(4) replace $\ln P_t$ with the annual Känzig (2021) oil supply surprise (reduced form); columns (5)–(6) use the Kilian (2008) OPEC disruption measure (1989–2003). Even columns add the pre-sample capacity interaction $LowICRG_c \times$ Shock, where $LowICRG_c$ is below-median ICRG quality-of-government averaged 1984–1988. Lower-order terms are absorbed by cell FE and country×year FE. PPML observations are smaller than asinh-OLS observations because Poisson separation drops cells that record zero conflict deaths in every year of the sample. Ring cells excluded throughout. SE clustered at cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Columns (1)–(2) use the main oil-price interaction; columns (3)–(4) replace it with the Känzig (2021) annual surprise; columns (5)–(6) use the Kilian (2008) OPEC disruption on the 1989–2003 sub-sample. Even columns add the pre-1989 ICRG state-capacity interaction. The conflict-death coefficients in Panel A are small, negative, and statistically indistinguishable from zero across all six specifications, indicating that the aggregate conflict null under the $\ln(1+Y)$ transformation extends to the alternative arcsinh

transformation. Panel B (PPML on raw counts) tells a more nuanced story and requires two pieces of disclosure for columns (1)–(2). First, Poisson separation drops cell-years that record zero conflict deaths in every sample year, reducing the effective sample to a conflict-active subsample of cells with at least one death year. Second, conflict deaths are heavily right-skewed, so PPML places substantial weight on a small set of catastrophic-event cell-years. On this selected subsample with heavy-tail weighting, the column (1) PPML coefficient on $Oil_g \times \ln P_t$ turns borderline-positive, in contrast to the precisely-estimated negative coefficient recovered under supply-shock identification in column (3). The contrast is the within-PPML analogue of the demand-side confounding that the supply-shock instruments are designed to purge. The aggregate conflict null is therefore robust to the arcsinh transformation in Panel A; the PPML evidence in Panel B is mixed and sensitive to outlier trimming and identification strategy, as elaborated in the next paragraph.

Figure A8: **Zero-inflation robustness**

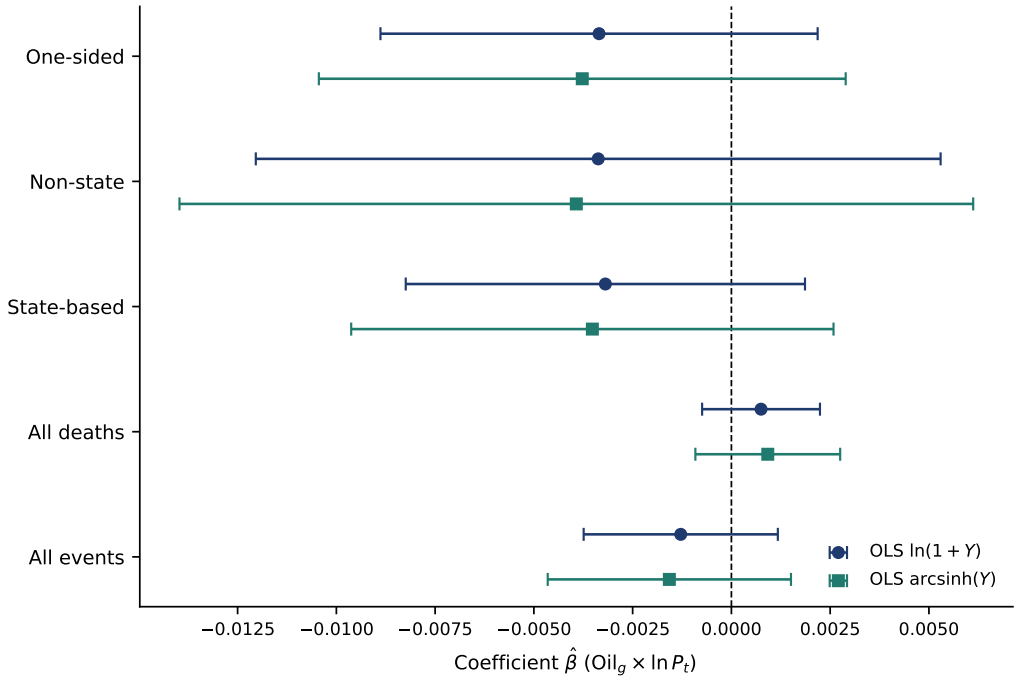

OLS estimates for five UCDP GED conflict outcomes under two functional forms. Navy circles: $\ln(1+Y)$ (Panel A); teal squares: $\text{arcsinh}(Y)$ (Panel B). Bars are 95% confidence intervals. Standard errors clustered at the PRIO-GRID cell level. PPML (Panel C) estimated in Stata using `ppmlhdfe`.

Outlier sensitivity for the column (1) PPML coefficient. The borderline-positive column (1) PPML estimate is identified from a small set of catastrophic-event cell-years on which the supply-shock orthogonality discipline matters most. To make this transpar-

ent, I re-estimate column (1) on a panel that drops cell-years above the 99th percentile of the *conflict_deaths* distribution in the ring-excluded sample. Because conflict deaths are concentrated in a small set of catastrophic cell-years — approximately 99 percent of cell-years record zero deaths — this trim removes roughly 15,000 cell-years that together contain virtually all of the in-sample fatalities. Table A4 reports the comparison: under the trim, the PPML coefficient reverses sign and becomes statistically indistinguishable from zero, directionally consistent with the Känzig reduced-form benchmark in column (3). The sign reversal under outlier trimming reinforces the interpretation that the unrestricted column (1) coefficient is identified from demand-correlated catastrophic-event variation that the supply-shock instruments are designed to purge.

Table A4: **Table A3, Panel B, col (1) robustness: PPML on conflict deaths trimming the top 1% of fatality cell-years**

	(1) Baseline (Table A3, Panel B)	(2) Trim top 1% (deaths $>$ 2 dropped)
$Oil_g \times \ln P_t$	$+0.5147^{*}$	-0.0989
	(0.2713)	(0.1580)
Cell-years (input)	1,593,551	1,578,420
Cell-years dropped	—	15,131
PPML observations	103,892	43,494
Total deaths in sample	1,861,052	6,420
Cell FE	✓	✓
Country×year FE	✓	✓
Ring cells excluded	✓	✓

Notes: Column (1) reproduces Table A3, Panel B, column (1): PPML of conflict deaths on $Oil_g \times \ln P_t$ with cell and country×year fixed effects, 1989–2014 ring-excluded sample, SE clustered at the PRIO-GRID cell level. Column (2) drops cell-years above the 99th percentile of *conflict_deaths* in the same sample (cutoff: deaths $>$ 2). Because the death distribution is highly zero-inflated (98.76% of cell-years record zero deaths), this cutoff is at deaths $>$ 2 and the trim removes 15,131 cell-years that together contain 99.66% of total in-sample fatalities. Under this trim the coefficient falls from $+0.5147$ to -0.0989 and loses significance ($p = 0.531$), confirming that the positive PPML sign in column (1) of Table A3 is identified almost entirely from a small set of high-fatality cell-years. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

C.2 Time-Varying WGI State Capacity

Table A5 reports estimates from a triple-interaction analogue of equation (3) using the time-varying World Bank Government Effectiveness Index (WGI) as the state capacity moderator. The specification includes all three constituent terms: $\hat{\beta}$ (baseline oil-price response), $\hat{\gamma}$ (the lower-order level difference between petroleum cells in low- vs. high-capacity settings), and $\hat{\theta}$ (the differential price-sensitivity in low-capacity settings).

Omitting $\hat{\gamma}$ would conflate the level difference with the price-sensitivity, biasing $\hat{\theta}$.

Two findings emerge. First, the baseline development effect in the QoG-matched subsample is essentially identical to the full-sample estimate, confirming that the moderator analysis is identified on a representative subset. Second, the triple-interaction coefficients for conflict are precisely zero in both events and deaths, consistent with the conflict null in Table 3. The night-lights interaction is directionally consistent with the pre-sample ICRG result in Section 7.1 but falls short of conventional significance thresholds. The imprecision reflects two problems resolved by the pre-sample ICRG specification: the time-varying moderator induces collinearity between the lower-order level term and the triple interaction, and government effectiveness scores themselves respond to oil booms and conflict episodes (see Robinson et al., 2006), making causal interpretation of the time-varying moderator difficult.

Table A5: Conditional effects of state capacity on oil-windfall responses (QoG-matched subsample)

	(1) Night lights	(2) Conflict events	(3) Conflict deaths
$\text{Oil}_g \times \ln P_t$ ($\hat{\beta}$)	$+0.0379^{***}$	-0.0008	$+0.0006$
	(0.0051)	(0.0034)	(0.0049)
$\text{Oil}_g \times$ Low state capacity ($\hat{\gamma}$)	$+0.0168$	$+0.0216$	$+0.0328$
	(0.0315)	(0.0283)	(0.0429)
$\times$ Low state capacity ($\hat{\theta}$)	-0.0118	$+0.0005$	$+0.0004$
	(0.0078)	(0.0070)	(0.0109)
Cell FE	✓	✓	✓
Country-year FE	✓	✓	✓
N	966,185	1,057,084	1,057,084

Notes: OLS estimates using oil-price exposure $Oil_g \times \ln P_t$. Sample restricted to country-years with non-missing government effectiveness data (QoG dataset, WBGI government effectiveness `wbgi_gee`). Low state capacity = below-median `wbgi_gee`. All three constituent terms of the triple interaction are included: $\hat{\beta}$ is the baseline oil-price response, $\hat{\gamma}$ captures the time-varying level difference in outcomes between petroleum cells in low- vs. high-capacity country-years, and $\hat{\theta}$ is the differential oil-price sensitivity in low-capacity settings. Cell + country-year FE. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

D Commodity-Cycle Diagnostic

Table A6 replicates the commodity-cycle channel evaluation of Section 9.2 using the primary OLS design ($Oil_g \times \ln P_t$) rather than the Känzig supply-news shock. Panel A reports the oil-cell benchmark; Panel B substitutes gold-cell exposure ($Gold_g \times \ln P_t$); Panel C enters both simultaneously (horse-race); Panel D adds a direct gold-price control ($Gold_g \times \ln P_t^{gold}$).

Table A6: Alternative channel tests: oil-specific vs. commodity-cycle dynamics (primary OLS design, 1989–2014)

	(1) Night lights (ln)	(2) Conflict events (ln)	(3) Conflict deaths (ln)
Panel A: Oil benchmark ($Oil_g \times \ln P_t$)			
$\hat{\beta}$	$+0.0407^{***}$	-0.0044	-0.0049
	(0.0040)	(0.0028)	(0.0044)
Panel B: Gold-exposure comparison ($Gold_g \times \ln P_t$)			
$\hat{\beta}$	$+0.0154^{**}$	-0.0123^{**}	-0.0184^{**}
	(0.0069)	(0.0055)	(0.0092)
N	1,301,257	1,593,551	1,593,551
Panel C: Horse-race — oil coefficient ($Oil_g \times \ln P_t$)			
$\hat{\beta}$	$+0.0407^{***}$	-0.0044	-0.0050
	(0.0040)	(0.0028)	(0.0044)
Panel C (cont.): Horse-race — gold coefficient ($Gold_g \times \ln P_t$)			
$\hat{\beta}$	$+0.0158^{**}$	-0.0124^{**}	-0.0184^{**}
	(0.0070)	(0.0055)	(0.0092)
N	1,301,257	1,593,551	1,593,551
Panel D: Gold-price control — oil coeff. ($Oil_g \times \ln P_t$)			
$\hat{\beta}$	$+0.0407^{***}$	-0.0045	-0.0050
	(0.0040)	(0.0028)	(0.0044)
Panel D (cont.): Gold-price control — gold coeff. ($Gold_g \times \ln P_t^{gold}$)			
$\hat{\beta}$	$+0.0207^{***}$	-0.0214^{**}	-0.0281^{**}
	(0.0079)	(0.0088)	(0.0131)
N	1,301,257	1,593,551	1,593,551
Cell FE	✓	✓	✓
Country×Year FE	✓	✓	✓
Ring cells excl.	✓	✓	✓

Notes: All specifications are OLS with cell and country×year fixed effects. Ring-adjacent non-petroleum cells excluded from the control group (same sample as preferred specification in Table 3). The year-level shifter is $\ln P_t$ (log annual Brent crude price, World Bank). Gold cells: $Gold_g = \mathbf{1}[\texttt{goldplacer\_s} > 0]$ from PRIO-GRID v3 (Tollefsen et al. 2012), symmetric with $Oil_g = \mathbf{1}[\texttt{petroleum\_s} > 0]$. Panel D adds $Gold_g \times \ln P_t^{\text{gold}}$ as a gold-price control. Cell-clustered SEs. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The oil development coefficient is positive, precisely estimated, and stable across

the standalone and horse-race specifications. The gold-cell coefficient is statistically indistinguishable from zero for conflict outcomes in all panels and only marginally significant for night lights. Adding a direct gold-price control (Panel D) leaves the oil coefficient unchanged. The four panels together confirm that the development effect is oil-revenue specific and not driven by correlated commodity-demand dynamics.

Table A7: ACLED reporting-bias robustness: oil price exposure on ACLED conflict outcomes (1997–2014)

	(1) Baseline	(2) Lagged P	(3) +Weather	(4) Lagged +Weather
ACLED all events	+0.0003	+0.0004	+0.0000	+0.0002
	(0.0019)	(0.0019)	(0.0017)	(0.0017)
ACLED violence	+0.0008	+0.0010	+0.0003	+0.0005
	(0.0016)	(0.0016)	(0.0015)	(0.0014)
ACLED deaths	+0.0009	+0.0010	+0.0006	+0.0007
	(0.0018)	(0.0017)	(0.0017)	(0.0015)
Sexual violence	−0.0000	+0.0000	+0.0000	+0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Cell FE	✓	✓	✓	✓
Country-year FE	✓	✓	✓	✓
N	1,103,238	1,103,238	1,103,238	1,103,238

Notes: OLS estimates using geological oil-price exposure ($Oil_g \times \ln P_t$). Outcomes are from ACLED (Armed Conflict Location & Event Data Project, 1997–2014), an alternative conflict data source using different primary sources and coding rules than UCDP GED. "ACLED all events" includes all event types; "ACLED violence" excludes peaceful demonstrations and protest interventions; "ACLED deaths" are reported fatalities; "Sexual violence" counts events coded as `sub_event_type = Sexual violence`. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

E Identification Diagnostics: Spillovers and Year-by-Year Coefficients

This appendix reports two diagnostics that speak to how the regression's identifying variation is structured. The first is a spillover decomposition that quantifies the within-country-year leakage from petroleum to ring-adjacent cells, motivating the ring-excluded sample used throughout the main text. The second is a year-by-year version of the baseline regression that plots the cell-level coefficient over the sample period under both weak (year FE) and baseline (country×year FE) identifying variation.

E.1 Spatial Spillovers to Adjacent Cells

This subsection reports the full spillover decomposition. The specification partitions cells into three exhaustive groups: *core* petroleum cells ($Oil_g = 1$), a *ring* of non-petroleum cells sharing a border or corner with a petroleum cell ($Ring_g = 1$), and *remote* cells (all remaining cells). Both interactions enter the same regression with remote cells as the omitted baseline. All estimates use the primary design (OLS) with cell and country-year fixed effects; standard errors clustered at the PRIO-GRID cell level.

Core petroleum cells brighten substantially, with a small and statistically insignificant conflict response. Ring-adjacent cells also brighten (about two-thirds the core-cell development response) and experience a small but precisely-estimated reduction in conflict events. The ring-cell conflict effect quantifies the control-group contamination that motivates the ring-excluded sample in Table 3: including ring cells in the control group pulls the control mean down enough to partially offset the petroleum-cell estimate, biasing the full-sample conflict coefficient toward zero. The combination of positive development spillovers and small conflict-reducing effects in adjacent cells is inconsistent with a pure enclave model of oil-cell economic activity, pointing instead to broader fiscal and security externalities from petroleum production zones.

Figure A9: **Spillover effects to adjacent cells (primary OLS design)**

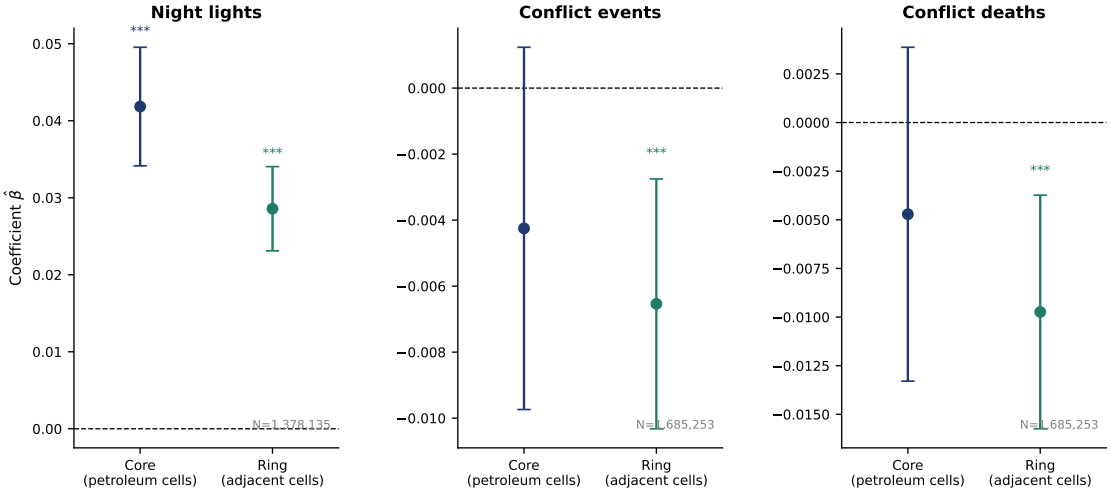

OLS estimates for core petroleum cells ($Oil_g \times \ln P_t$) and the first ring of adjacent non-petroleum cells ($Ring_g \times \ln P_t$). Remote cells are the omitted baseline. Navy circles: core petroleum cells; teal diamonds: ring-1 adjacent cells. Bars are 95% confidence intervals clustered at the PRIO-GRID cell level.

Table A8: Spillover decomposition: core, ring, and remote cells (primary design)

	(1) Night lights	(2) Conflict events	(3) Conflict deaths
$\text{Oil}_g \times \ln P_t$ (core)	$+0.0418^{***}$	-0.0043	-0.0047
	(0.0039)	(0.0028)	(0.0044)
$\text{Ring}_g \times \ln P_t$ (neighbor)	$+0.0286^{***}$	-0.0065^{***}	-0.0097^{***}
	(0.0028)	(0.0019)	(0.0031)
Cell FE	✓	✓	✓
Country-year FE	✓	✓	✓
N	1,378,135	1,685,253	1,685,253

Notes: OLS with cell and country×year fixed effects. Remote cells (not core, not adjacent to core) are the omitted baseline. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

E.2 Year-by-Year Petroleum-Cell Coefficients

This subsection reports year-by-year estimates of the petroleum-cell coefficient under two FE structures. The weaker version uses year fixed effects and recovers a combination of within- and cross-country variation; the stronger version uses country×year fixed effects and matches the baseline identifying variation of Table 3.

Figure A10 plots the year-specific petroleum-cell premium $\hat{\beta}_t$ from the regression

$$Y_{gct} = \sum_{t \neq t_0} \beta_t \big(Oil_g \times \mathbf{1}[\text{year} = t]\big) + \alpha_g + \alpha_t + \varepsilon_{gct},$$

where α_g are cell fixed effects and α_t are *year* fixed effects (not country-year). Years at or below the sample median oil price ($26.58/barrel) are the *low price period*; years above are the *high price period*.

Two patterns are visible. Night-light coefficients are positive throughout and widen markedly during the high-price period, consistent with oil rents accelerating local activity. Conflict coefficients are also positive throughout, rising sharply in 2012–2014. This does not contradict the within-country result in Table 3. The year-FE specification here compares oil cells to all non-oil cells globally, without absorbing country-specific conflict levels. Since oil cells concentrate in conflict-prone countries (Middle East, Sub-Saharan Africa), the global comparison picks up their inherent conflict-proneness rather than the causal effect of a price windfall.

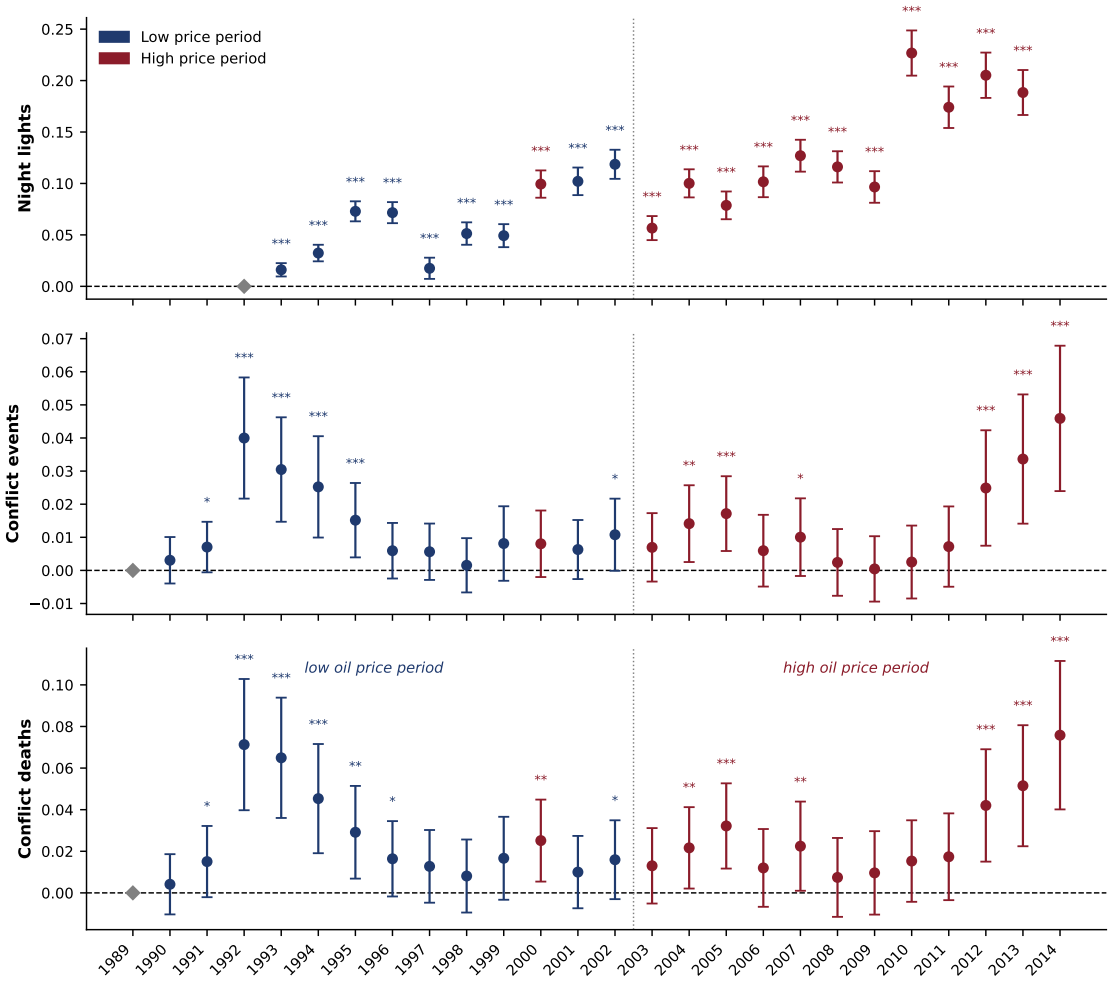

Figure A10: Year-by-year petroleum-cell coefficients ($\hat{\beta}_t$) from a regression of each outcome on $Oil_g \times \mathbf{1}[\text{year} = t]$ with cell and year fixed effects. Blue markers: years at or below the sample median oil price (low price period). Red markers: years above the median (high price period). Diamond marker at the reference year is normalised to zero. Vertical dotted line separates the two regimes. Bands are 95% confidence intervals (cell-clustered).

Within-country-year identification. To match the identifying variation of the main specification exactly, Figure A11 re-runs the year-by-year regression replacing α_t with country-year fixed effects λ_{ct}:

$$Y_{gct} = \sum_{t \neq t_0} \beta_t \left(Oil_g \times \mathbf{1}[\text{year} = t] \right) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}.$$

Each $\hat{\beta}_t$ is now identified purely from the within-country-year cross-section of petroleum versus non-petroleum cells, the same variation that delivers the baseline coefficient in Table 3. The night-lights coefficients widen sharply during the 2002–2008 boom, mir-

roring the headline development result. The conflict coefficients — now stripped of the country-level conflict-proneness that drove the positive trajectory in Figure A10 — are small and centred near zero throughout the sample.

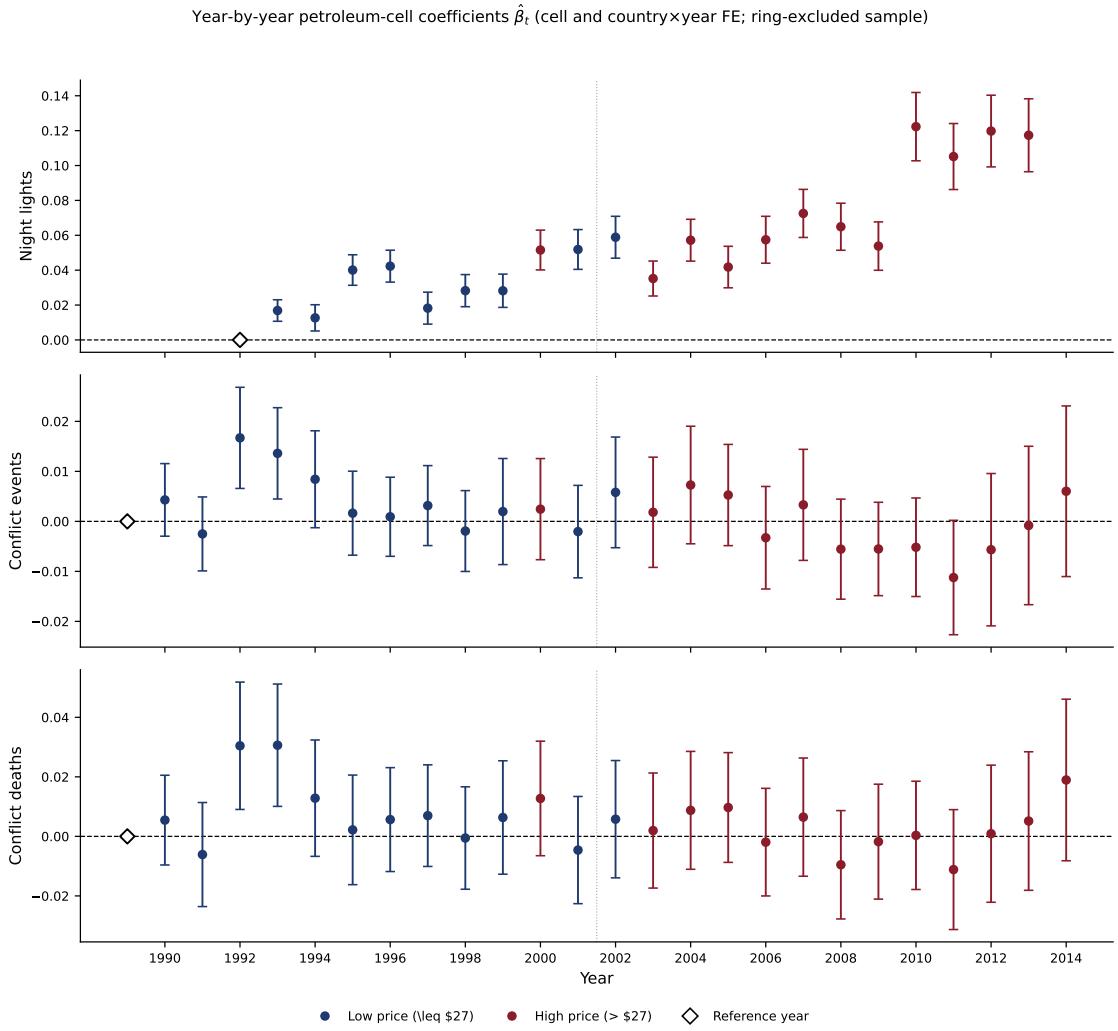

Figure A11: Year-by-year petroleum-cell coefficients $\hat{\beta}_t$ under the baseline identifying variation (cell and country×year fixed effects; ring-excluded sample). Blue markers: years at or below the sample median oil price; red markers: years above. White diamond at the reference year is normalised to zero. Vertical dotted line separates low- and high-price regimes. Bands are 95% confidence intervals (cell-clustered). The night-lights divergence after 2002 is the within-country-year analog of the boom-period development effect; conflict coefficients are small throughout.

F Intensive Margin: Deposit Density as a Continuous Endowment Proxy

The primary design uses a binary petroleum indicator — $Oil_g = \mathbf{1}[\texttt{petroleum\_s} > 0]$ — which treats all oil cells equally regardless of how large or concentrated the underlying deposit is. This section tests whether the price response scales with deposit intensity, exploiting the fact that some cells sit at the heart of major petroleum basins (the Niger Delta, the Arabian Peninsula, the Permian Basin) while others contain a single isolated field.

Continuous endowment proxy. I construct $density_g$ as the fraction of a cell's 3×3 Moore neighbourhood (the cell itself plus its 8 contiguous neighbours) that contain petroleum deposits:

$$density_g = \frac{\sum_{j \in \mathcal{N}(g)} \mathbf{1}[Oil_j = 1]}{9}$$

where $\mathcal{N}(g)$ includes g itself. The measure ranges continuously from 0 (no petroleum in the surrounding 1.5×1.5 area) to 1 (all nine cells are oil cells). Isolated oil cells score $1/9 \approx 0.11$; cells embedded in large basins score 0.7–1.0. The median density among oil cells is 0.556 (mean 0.556), with minimum 0.111 and maximum 1.000. This neighbourhood-density proxy captures the geographic concentration of petroleum endowment and is analogous in spirit to the share-of-area-that-is-gold-suitable measure of Girard et al. (2025), although constructed from deposit presence data rather than a geological suitability raster (the latter requires geological map data not available globally in digitised form).[11]

Intensive-margin results. Table A9 reports four specifications for each of three outcomes: binary baseline, binary with ICRG capacity interaction, density baseline, and density with ICRG capacity interaction. Two results stand out. First, the density coefficients are roughly 1.5–1.9 times larger in absolute value than the corresponding binary coefficients, consistent with the interpretation that deposit-dense cells extract more revenue per unit of price and therefore respond more strongly to oil-price exposure. The ratio is stable across development and conflict outcomes. Second, the ICRG capacity

[11] Girard et al. (2025) construct their gold-suitability measure by overlaying African geological maps from Frost-Killian et al. (2016) and Thiéblemont (2016) onto PRIO-GRID cells and computing the share of each cell's area covered by gold-suitable lithologies. This approach is Africa-specific. A global petroleum analog would require overlaying the USGS World Petroleum Assessment (2000) Total Petroleum System boundaries onto PRIO-GRID cells; that exercise is left for future work.

moderation scales proportionally with deposit intensity: the triple-interaction coefficient roughly doubles when the treatment moves from a binary to a continuous density measure, confirming that the institutional moderation is not an artefact of the binary treatment definition but deepens as deposit concentration rises. Conflict interactions remain statistically null under both treatment specifications.

The proportionality of binary and density coefficients is a validation of the binary instrument. If the binary indicator were measuring something other than petroleum endowment — for instance, proximity to roads, coastlines, or other correlated economic geography — the density scaling would be unlikely to reproduce the same conflict null and development positive across both specifications. Instead, the responses are consistent in sign, significance pattern, and approximate magnitude ratio, supporting a causal interpretation anchored in petroleum geology.

Figure A12: **Intensive margin: binary vs. deposit-density petroleum treatment**

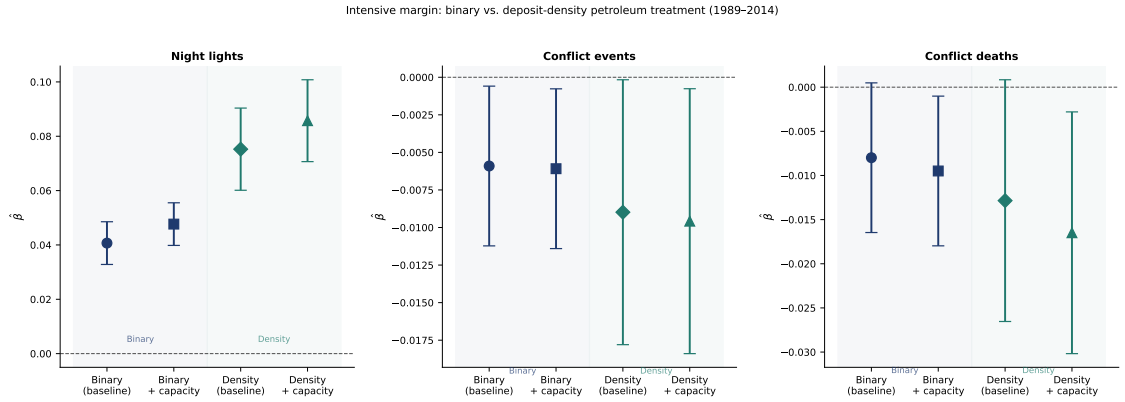

OLS estimates ($\hat{\beta}$, 95% CI) for four specifications across three outcomes. Navy markers: binary $Oil_g \times \ln P_t$ treatment; teal markers: continuous density treatment $density_g \times \ln P_t$. "+ capacity" columns add the ICRG triple interaction. Ring-excluded sample, cell and country-year FE, standard errors clustered at the PRIO-GRID cell level.

Table A9: Intensive margin: petroleum endowment density as a continuous shift-share proxy

	$\hat{\beta}$ (treatment $\times \ln P_t$)	
	$\hat{\beta}$	SE
Night lights		
Binary baseline	$+0.0407^{***}$	(0.0040)
Binary + ICRG capacity	$+0.0477^{***}$	(0.0040)
$\hat{\theta}$ ($\times$ Low ICRG)	-0.0284^{***}	(0.0065)
Density baseline	$+0.0752^{***}$	(0.0077)
Density + ICRG capacity	$+0.0857^{***}$	(0.0077)
$\hat{\theta}$ ($\times$ Low ICRG)	-0.0527^{***}	(0.0124)
Conflict events		
Binary baseline	-0.0059^{**}	(0.0027)
Binary + ICRG capacity	-0.0061^{**}	(0.0027)
$\hat{\theta}$ ($\times$ Low ICRG)	$+0.0007$	(0.0090)
Density baseline	-0.0090^{**}	(0.0045)
Density + ICRG capacity	-0.0096^{**}	(0.0045)
$\hat{\theta}$ ($\times$ Low ICRG)	$+0.0030$	(0.0177)
Conflict deaths		
Binary baseline	-0.0080^{*}	(0.0043)
Binary + ICRG capacity	-0.0095^{**}	(0.0043)
$\hat{\theta}$ ($\times$ Low ICRG)	$+0.0061$	(0.0147)
Density baseline	-0.0128^{*}	(0.0070)
Density + ICRG capacity	-0.0165^{**}	(0.0070)
$\hat{\theta}$ ($\times$ Low ICRG)	$+0.0183$	(0.0280)
Cell FE	✓	
Country-year FE	✓	
N (binary)	1,301,257	
N (+ capacity)	1,301,257	

Notes: OLS with cell and country-year fixed effects (ring-excluded sample, 1989–2014). Binary specifications use the standard $Oil_g \times \ln P_t$ treatment. Density specifications replace the binary oil indicator with $density_g$, the fraction of the 3×3 Moore neighbourhood (cell g and its 8 contiguous neighbours) that contain petroleum deposits. This continuous proxy ranges from 0 (no oil in the local 3×3 area) to 1 (all 9 cells are oil cells), capturing deposit concentration rather than mere presence. The shift-share treatment $density_g \times \ln P_t$ tests whether the price response scales with geological endowment intensity. The capacity moderator is the ICRG quality-of-government index averaged 1984–1988 (pre-sample, below-median indicator); $\hat{\theta}$ is the triple-interaction coefficient for petroleum cells in low-capacity settings. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

G Supply-Shock Instrument Validation and Mechanism Replications

This appendix collects the IV (2SLS) evidence that corroborates the main-text OLS findings. The instruments are the Känzig (2021) high-frequency oil-supply news surprise and the Kilian (2008) OPEC production-shortfall series; both restrict identifying variation to supply-side movements orthogonal to global demand. The appendix has two parts: the baseline IV comparison and conditional state-capacity estimates (B.1); and IV replications of the three mechanism tests reported in Section 8 (B.2).

G.1 Instrument Comparison and IV State-Capacity

Instrument comparison. Figure A13 places the Kilian and Känzig IV estimates side by side for all three outcomes. The night-light coefficient is positive and significant under both instruments; conflict coefficients are near zero or negative under both. Two instruments identified from distinct mechanisms and covering non-overlapping subperiods return identical substantive conclusions — the paper's strongest evidence that the fixed-effects estimates reflect genuine local oil-revenue transmission rather than demand-side confounds.

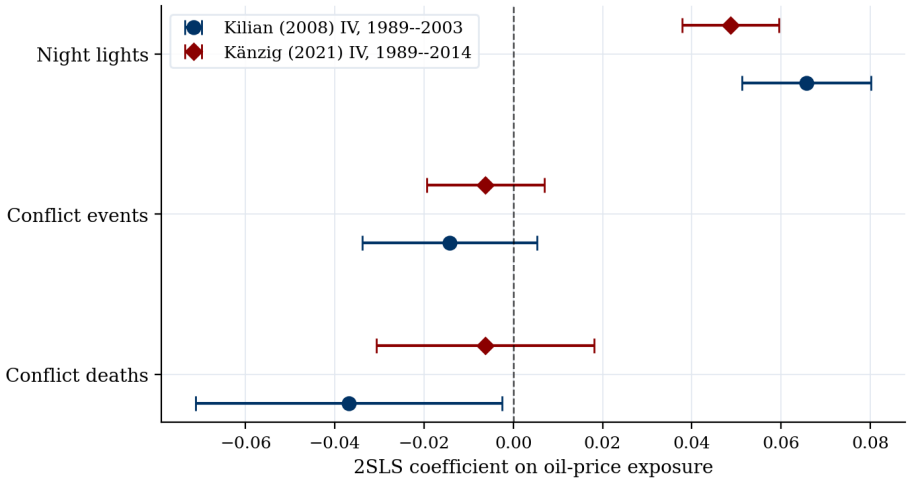

Figure A13: Kilian vs. Känzig IV estimates for three outcomes. Navy circles are 2SLS coefficients using the Kilian (2008) instrument (1989–2003); teal diamonds use the Känzig (2021) high-frequency surprise instrument (full 1989–2014 sample). Bars are 95% confidence intervals. Standard errors clustered at the PRIO-GRID cell level.

IV state-capacity interaction. Table A10 reports Kilian IV estimates of the state-capacity conditional pattern, using the same *pre-1989* country-level moderator as the main-text

OLS specification (Table 5): $LowCap_c$ equals one if the country's 1984–1988 average ICRG quality-of-government composite lies below the across-country median (60 of 122 ICRG-covered countries). The pre-period moderator is genuinely pre-treatment, which matters for the conditional IV: a moderator constructed from in-sample data can itself be influenced by oil-price shocks, so averaging an in-sample series like WBGI over 1989–2014 (a previous version of this appendix) is time-invariant but not pre-determined. Switching to pre-1989 ICRG resolves the conditioning issue and makes the IV moderator and the OLS moderator the same object, so the OLS-vs-IV comparison is a comparison of identification strategies on a common conditioning set rather than a comparison across both identification and moderator construction.

Three results emerge under causal identification. First, the development response is strong in high-capacity cells ($+0.090^{***}$) and is sharply attenuated — by roughly two-thirds — in weak-capacity cells (low-capacity total $+0.032$). The development moderation pattern that Section 7.1 recovers under OLS is reproduced under supply-shock identification with the same pre-period moderator. Second, the high-capacity conflict response is statistically indistinguishable from zero for both events and deaths. Third, the weak-capacity interaction is *positive* for events ($+0.045^{**}$) and positive but imprecise for deaths ($+0.019$, n.s.): the direction is consistent with the canonical resource-curse prediction (more conflict in low-capacity cells during high-price years) and with the LC-E boom-period spike in the four-cell event study reported in Section 7.2 (Table 7). The deaths point estimate is small and not separately significant, so we do not lean on it; the events interaction is suggestive but, like the LC-E contrast, should be read in the context of the multiple-testing discussion in Section 7.2.

Table A10: IV conditional effects under weak state capacity (pre-1989 ICRG moderator)

	(1) Night lights	(2) Conflict events	(3) Conflict deaths
Oil cell × log oil price	$+0.0905^{***}$	-0.0004	-0.0066
	(0.0052)	(0.0110)	(0.0223)
× weak state capacity	-0.0584^{***}	$+0.0454^{**}$	$+0.0186$
	(0.0096)	(0.0203)	(0.0412)
Implied low-capacity total	$+0.0321$	$+0.0450$	$+0.0120$
Cell FE	✓	✓	✓
Country-year FE	✓	✓	✓
Instrument set	Two-shock	Two-shock	Two-shock
N	576,987	736,433	736,433

Notes: 2SLS estimates of $Y_{gct} = \beta_1(Oil_g \times \ln P_t) + \beta_2(Oil_g \times \ln P_t \times LowCap_c) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$. Both $Oil_g \times \ln P_t$ and its interaction are treated as endogenous and instrumented with $Oil_g \times S_t^{\text{Kilian}}$ and the same interacted with $LowCap_c$, where S_t^{Kilian} is the Kilian (2008) OPEC production-shortfall series. $LowCap_c$ is a *pre-sample, time-invariant* country-level indicator equal to one if the country's 1984–1988 average ICRG quality-of-government composite lies below the across-country median (60 of 122 ICRG-covered countries). The pre-1989 ICRG moderator is genuinely pre-treatment: it is averaged over the five years immediately before the sample window opens in 1989, and it matches the moderator used in the main-text OLS specification (Table 5). An earlier version of this appendix used the country-level mean of WBGI government effectiveness over 1989–2014; that moderator is time-invariant but not pre-determined (it can be influenced by oil-price shocks during the sample period), so we substitute the pre-period ICRG measure to preserve the conditioning-on-pre-treatment property. "Implied low-capacity total" is $\hat{\beta}_1 + \hat{\beta}_2$. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

G.2 Mechanism Replications Under Supply-Shock Identification

This subsection replicates the three mechanism tests of Section 8 (livelihood insurance, price-cycle asymmetry, onshore–offshore decomposition) under supply-shock identification. The Känzig and Kilian instruments are used in place of $\ln P_t$, and the corresponding interactions follow the same construction.

Livelihood insurance under IV. Table A11 reports the IV analogue of the weather-insurance specification in Section 8. Two qualifications are worth stating directly. First, the IV main effect on night lights ($+0.022^{**}$) is positive and significant, corroborating the OLS development response of $+0.018^{***}$. Second, the IV triple interaction $Oil_g \times \ln P_t \times ExtrTemp_{gt}$ on night lights has the same sign as the OLS analogue ($+0.026$ vs $+0.007^{***}$) and a point estimate roughly 3.6 times larger, but the standard error widens to 0.089 and the buffering effect is no longer statistically detectable under causal identification. The conflict columns remain statistically indistinguishable from zero in both OLS and IV. The IV thus replicates the development-response *sign* and *direction* of the OLS weather-

buffering result but cannot speak to its statistical significance; this is consistent with the more general loss of precision under supply-shock identification (the first-stage variation of the triple interaction is two orders of magnitude weaker than that of the main effect, with first-stage F of 392 vs 54,182), and is reported transparently rather than overstated.

Table A11: Oil Windfalls as Livelihood Insurance: OLS and IV Estimates

	OLS			2SLS (Känzig IV)		
	(1) *Night lights*	(2) *Events*	(3) *Deaths*	(4) *Night lights*	(5) *Events*	(6) *Deaths*
$Oil_g \times \ln P_t$	$+0.0184^{***}$	-0.0020	-0.0029	$+0.0219^{**}$	-0.0071	-0.0175
	(0.0040)	(0.0025)	(0.0042)	(0.0099)	(0.0102)	(0.0161)
$\times ExtrTemp_{gt}$	$+0.0070^{***}$	$+0.0007$	$+0.0009$	$+0.0255$	-0.0467	$+0.0305$
	(0.0022)	(0.0020)	(0.0031)	(0.0887)	(0.0681)	(0.0986)
Cell FE	✓	✓	✓	✓	✓	✓
Country×Year FE	✓	✓	✓	✓	✓	✓
Ring cells excl.	✓	✓	✓	✓	✓	✓
Observations	874,746	874,746	874,746	874,746	874,746	874,746
First-stage F (main interaction)	—	—	—	54,182	54,182	54,182
First-stage F (triple interaction)	—	—	—	392	392	392

Notes: Columns (1)–(3) report OLS with cell and country×year fixed effects on the Känzig sample (1989–2014). Columns (4)–(6) are 2SLS, instrumenting $Oil_g \times \ln P_t$ and $Oil_g \times \ln P_t \times ExtrTemp_{gt}$ with $Oil_g \times Surprise_t$ and $Oil_g \times Surprise_t \times ExtrTemp_{gt}$, respectively, where $Surprise_t$ is the Känzig (2021) annual oil supply news shock. $ExtrTemp_{gt}$ equals one if cell g's annual mean temperature in year t exceeds its long-run cell mean by more than two standard deviations. The exogenous main effect of $ExtrTemp_{gt}$ is partialled out via Frisch–Waugh–Lovell before the IV step and is not separately reported in IV columns. First-stage F-statistics for the two endogenous regressors are reported at the bottom of the IV columns. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Price-cycle asymmetry under IV. Table A12 replicates the boom–bust decomposition under two stricter identification conditions: (i) ring-adjacent cells are excluded from the control group (same sample as Table 3); and (ii) boom and bust instruments use the Känzig (2021) supply news surprise split as $Oil_g \times \max(S_t^K, 0)$ and $Oil_g \times \min(S_t^K, 0)$. The asymmetric *magnitude* of the development response (bust dimming exceeds boom brightening) is preserved in both samples; conflict coefficients remain statistically indistinguishable from zero in both OLS and IV under ring exclusion.

Table A12: Price-cycle asymmetry: OLS and IV estimates

	OLS			2SLS (Känzig IV)		
	(1) *Night lights*	(2) *Conflict events*	(3) *Conflict deaths*	(4) *Night lights*	(5) *Conflict events*	(6) *Conflict deaths*
Boom	-0.0199^{***}	-0.0275^{***}	-0.0341^{**}	$+0.0537^{**}$	-0.0471	-0.0441
	(0.0062)	(0.0086)	(0.0142)	(0.0217)	(0.0330)	(0.0603)
Bust	-0.0430^{***}	-0.0229^{***}	-0.0241^{*}	-0.1637^{***}	-0.0259	-0.0383
	(0.0072)	(0.0075)	(0.0141)	(0.0271)	(0.0195)	(0.0333)
p-value (Boom = Bust)	[0.000]	[0.528]	[0.467]	[0.000]	[0.604]	[0.940]
Cell FE	✓	✓	✓	✓	✓	✓
Country×Year FE	✓	✓	✓	✓	✓	✓
Ring cells excl.	✓	✓	✓	✓	✓	✓
Observations	835,797	835,797	835,797	835,797	835,797	835,797
First-stage F (boom)	—	—	—	425,890	425,890	425,890
First-stage F (bust)	—	—	—	807,126	807,126	807,126

Notes: The unit of observation is a PRIO-GRID cell-year. The two endogenous regressors are $Boom_{gt} = Oil_g \times \max(\Delta \ln P_t, 0)$, the local interaction with the positive part of the annual log-price change, and $Bust_{gt} = Oil_g \times |\min(\Delta \ln P_t, 0)|$, the local interaction with the absolute value of the negative part. Columns (1)–(3) report OLS; columns (4)–(6) report 2SLS, instrumenting $Boom_{gt}$ and $Bust_{gt}$ with $Oil_g \times \max(Surprise_t, 0)$ and $Oil_g \times |\min(Surprise_t, 0)|$, where $Surprise_t$ is the Känzig (2021) annual oil supply news surprise. The p-value row reports the Wald test for coefficient equality between the boom and bust interactions under the cluster-robust variance matrix; bracketed values are p-values. The sample drops the first year per cell because the lagged price is unavailable. The reported first-stage F-statistics are mechanically inflated by the shift-share structure of the instruments: $Oil_g \times Surprise_t$ collapses to a function of Oil_g within each country-year after fixed-effect absorption, so the within-group variation in the endogenous regressor is nearly perfectly explained by the corresponding within-group instrument variation, and the F statistic does not carry its usual weak-instrument interpretation. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Onshore vs. offshore under IV. Figure A14 reports the onshore–offshore decomposition under both supply-shock instruments. The onshore-country subsample reproduces the main-text pattern — no systematic conflict escalation in petroleum cells — under both instruments. The pure-offshore subsample is small (four countries, 16 petroleum cells) and imprecisely estimated.

H Colombia Reconciliation

This appendix places the paper's primary design estimates alongside the widely-cited Colombia results of Dube and Vargas (2013). D&V estimate that oil-price booms intensify paramilitary violence in Colombian oil municipalities, using within-Colombia variation over a comparable sample period. To compare like with like, I re-run the paper's primary specification — $Y_{gct} = \beta\,(Oil_g \times \ln P_t) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$ — on the Colombia sub-sample alone, restricting the panel to PRIO-GRID cells in Colombia (gwno = 100). Country×year fixed effects collapse to year fixed effects in the sub-sample because the country dimension is fixed; cell fixed effects continue to absorb time-invariant cell het-

Figure A14: **Onshore vs. offshore: Känzig and Kilian IV by country production type**

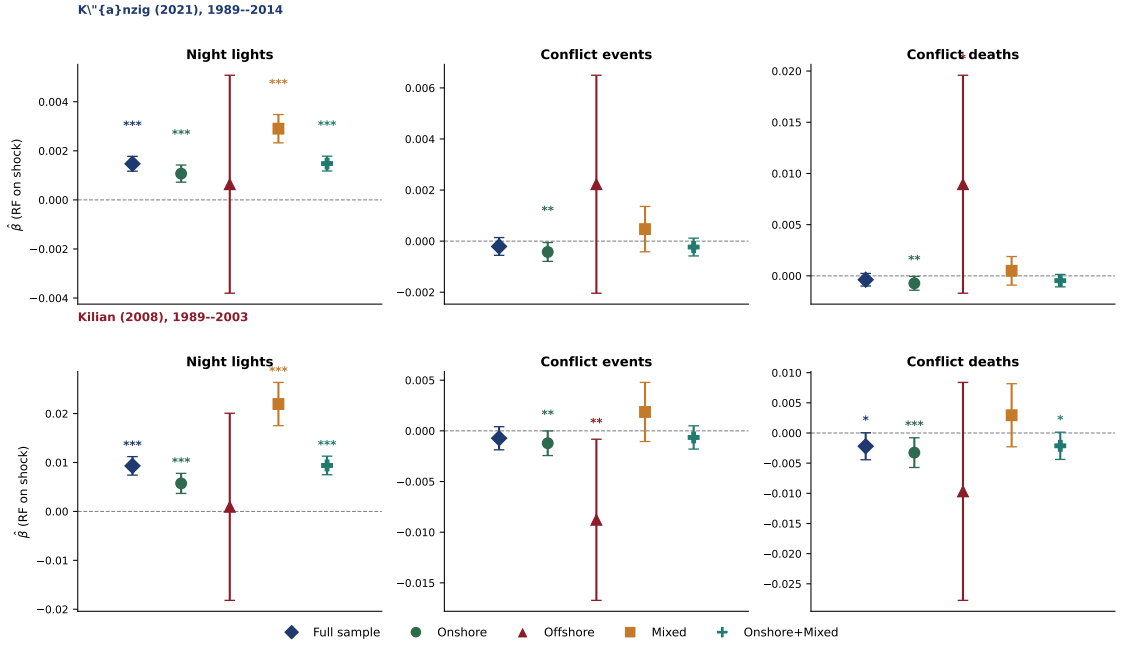

Onshore vs. offshore transmission estimates on the preferred ring-excluded sample. Känzig (2021) supply-news reduced form (top row, 1989–2014) and Kilian (2008) OPEC production-shortfall reduced form (bottom row, 1989–2003), each estimated separately within five country subsamples (full, onshore, offshore, mixed, onshore+mixed) using cell and country×year fixed effects. Bars are 95% confidence intervals clustered at the PRIO-GRID cell level. The onshore-country subsample reproduces the main-text pattern (no systematic conflict escalation) under both instruments; the pure-offshore subsample is small (4 countries, 16 petroleum cells) and imprecisely estimated.

erogeneity. Identification therefore exploits the within-Colombia, within-year cross-cell contrast between petroleum and non-petroleum cells over the price cycle.

Figure A15: **Colombia reconciliation (global panel vs. Colombia sub-sample)**

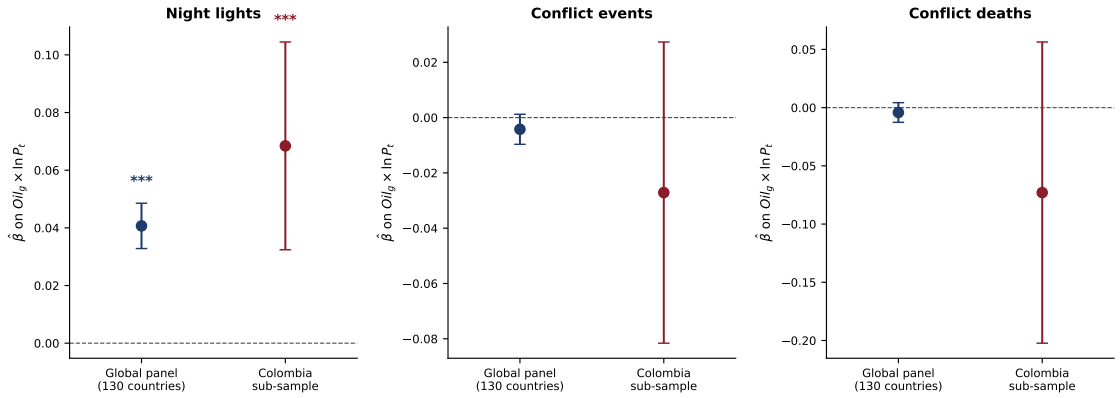

Coefficient plot comparing the paper's primary OLS estimates of $Oil_g \times \ln P_t$ on the global ring-excluded panel (195 countries, 74 with identifying within-country variation; see Table 1; navy circles) against the same specification restricted to Colombia (gwno = 100, crimson circles). Bars are 95% confidence intervals clustered at the PRIO-GRID cell level. Colombia: 328 cells, 29 of which are petroleum; 8,528 cell-year observations. Global panel: 1,593,551 cell-year observations.

Figure A15 and Table A13 report the comparison. The within-Colombia primary-OLS estimates are directionally consistent with the global panel on all three outcomes. Night lights respond more strongly in Colombia than in the global average; the conflict-deaths coefficient is also more negative in absolute value, though imprecise. The within-country, within-year design therefore recovers the same headline pattern — positive development response, negative-or-zero conflict response — on the Colombia subset alone, with magnitudes that are larger but estimated less precisely because the sub-sample contains only 29 petroleum cells.

The reconciliation with Dube and Vargas (2013) is therefore not in the sign of the cell-level coefficient but in the level of identification. D&V identify off municipality-by-year variation interacted with paramilitary presence within Colombia, recovering a within-Colombia within-municipality reallocation of conflict toward oil-producing areas during high-price years. Their design captures local-on-local effects of paramilitary group dynamics that the cell-level country×year fixed-effect design cannot speak to: my country-year fixed effects absorb every shock common to Colombian cells in a year, including the time-varying intensity of paramilitary mobilisation, leaving the differential response of petroleum versus non-petroleum cells at the same point in time. Both estimates are correct — they are estimates of different objects. The cell-level estimate in this paper is the differential local effect of an additional unit of oil exposure under common national conditions; the D&V estimate is the within-municipality reallocation

of conflict to oil locations over the price cycle. The Colombia sub-sample under our primary design therefore neither contradicts D&V nor reproduces their finding: it reports the cell-level analog, which carries the global peace-dividend pattern even on this single-country subset.

Table A13: Colombia reconciliation: primary OLS design (global panel vs. Colombia sub-sample, 1989–2014)

	Night lights		Conflict events		Conflict deaths	
	Global	Colombia	Global	Colombia	Global	Colombia
$Oil_g \times \ln P_t$	$+0.0407^{***}$	$+0.0684^{***}$	-0.0042	-0.0271	-0.0042	-0.0730
	(0.0040)	(0.0184)	(0.0028)	(0.0278)	(0.0043)	(0.0660)
N	1,301,257	7,216	1,593,551	8,528	1,593,551	8,528
Cell FE	✓	✓	✓	✓	✓	✓
Country×year FE	✓		✓		✓	
Year FE		✓		✓		✓

Notes: OLS estimates of β in $Y_{gct} = \beta\,(Oil_g \times \ln P_t) + \alpha_g + \lambda_{ct} + \varepsilon_{gct}$. The Global columns use the full ring-excluded panel (195 countries; 74 with within-country identifying variation; see Table 1); the Colombia columns restrict to Colombian cells (gwno $= 100$). Country×year fixed effects collapse to year fixed effects in the Colombia sub-sample because the country dimension is fixed. Standard errors clustered at the PRIO-GRID cell level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.