Does Foreign Aid Buffer Oil Revenue Shocks? Geocoded Evidence from Nigeria

APEP Autonomous Research* @ai1scl

March 9, 2026

Abstract

Can foreign aid prevent resource-dependent countries from sliding into conflict when commodity revenues collapse? I test the "aid-as-stabilizer" hypothesis using geocoded data from Nigeria, combining 376 AidData AIMS projects across 37 states with UCDP conflict events and monthly Brent crude oil prices. A continuous difference-in-differences design exploits the September 2008 oil price crash as an exogenous shock, interacted with pre-determined (December 2007) subnational aid exposure. The main estimate is positive but statistically insignificant ($\hat{\beta} = 0.143$, RI $p = 0.207$): states with greater aid exposure experienced weakly *more* conflict after the oil shock, not less. This null is robust across conflict types, alternative specifications, and Poisson count models. The 95 percent confidence interval (-0.033 to 0.318) rules out large protective effects but cannot exclude small buffering. The results challenge the presumption that development aid substitutes for lost resource revenues in preventing violence.

JEL Codes: F35, O17, D74, Q34
Keywords: foreign aid, oil shocks, armed conflict, Nigeria, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

In July 2008, Brent crude oil traded above \$133 per barrel. By December, it had fallen below \$40—a 70 percent collapse in five months. For Nigeria, where petroleum revenues fund roughly 80 percent of the federal budget and 95 percent of export earnings, this was not merely a financial event. It was a direct shock to the fiscal capacity of the state and, potentially, to the social contract that holds violence at bay (Ross, 2012; Sala-i Martin and Subramanian, 2013).

A large literature connects negative income shocks to armed conflict in developing countries (Miguel et al., 2004; Blattman and Miguel, 2010; Dube and Vargas, 2013). A parallel literature asks whether foreign aid can reduce violence—by financing public goods, buying hearts and minds, or simply substituting for lost government revenue (Berman et al., 2011; Crost et al., 2014; Collier and Hoeffler, 2004). These two literatures intersect at a question that is policy-relevant yet empirically unresolved: when commodity prices crash, does prior aid exposure buffer recipient regions against the resulting uptick in conflict?

This paper tests that question. I combine three geocoded datasets covering Nigeria from 1997 to 2014: (i) the AidData Nigeria AIMS v1.3.2 Research Release, which provides precise subnational locations for 376 development projects across 37 states (Tierney et al., 2011; Strandow et al., 2011); (ii) the Uppsala Conflict Data Program Georeferenced Event Dataset v24.1, which records armed conflict events across Nigeria with coordinates and fatality counts (Sundberg and Melander, 2013); and (iii) daily Brent crude oil prices from the Federal Reserve Bank of St. Louis. I construct a balanced panel of 37 states observed monthly over 18 years (7,992 state-months) and estimate a continuous difference-in-differences model that interacts pre-determined aid exposure (cumulative projects as of December 2007) with a post-shock indicator for the September 2008 oil price collapse.

The identifying assumption is straightforward: conditional on state and year-month fixed effects, subnational aid exposure as of 2007 is orthogonal to the differential evolution of conflict after the oil shock. Several features of the setting support this assumption. First, aid project locations were determined years before the oil crash by donor agencies responding to local needs assessments, not by anticipated conflict dynamics. Second, the oil price shock was driven by the global financial crisis—exogenous to Nigerian subnational politics. Third, the continuous treatment exploits variation across all 37 states rather than relying on a binary split that could conflate aid with other regional characteristics.

The main result is a null. The preferred estimate of $\hat{\beta} = 0.143$ (standard error 0.086, randomization inference $p = 0.207$) implies that a one-unit increase in log aid exposure is associated with a 14.3 percent increase in log conflict events after the shock—weakly *more* conflict, not less. The 95 percent confidence interval runs from -0.033 to 0.318, comfortably spanning zero and the wrong-signed region. This null is not an artifact of a single specification. It survives across four alternative outcome measures (state-based conflict, non-state conflict, fatalities, civilian deaths), a binary treatment indicator, raw project counts, annual aggregation, Poisson pseudo-maximum-likelihood estimation, and exclusion of the Federal Capital Territory. The randomization inference p-value, which permutes aid exposure across states 1,000 times, confirms that the point estimate is well within the null distribution. Placebo tests using non-event years (2003, 2005, 2011) yield coefficients near zero ($\hat{\beta} \approx 0.02$), and leave-one-out analysis shows the estimate ranges from 0.085 to 0.173 as each state is dropped—never approaching significance.

This paper contributes to three literatures. First, it speaks to the growing body of work on whether foreign aid reduces conflict at subnational scales. Crost et al. (2014) find that a development program in the Philippines increased insurgent violence—not because aid was bad, but because rebels targeted projects to undermine state legitimacy. Berman et al. (2011) show that small-scale reconstruction spending in Iraq reduced violence, but only when paired with adequate troop levels. Nunn and Qian (2014) demonstrate that US food aid—instrumented with American wheat production—increased civil conflict in recipient countries. Blair and Samii (2023) examine how foreign aid shapes rebel group behavior more broadly, while Gehring et al. (2022) provide subnational evidence linking aid to conflict escalation in certain settings. My contribution is to test a specific mechanism—aid as a buffer against commodity revenue shocks—that is often invoked in policy circles but rarely tested directly.

Second, the paper contributes to the commodity shocks and conflict literature. Dube and Vargas (2013) show that oil price increases fuel conflict in oil-producing municipalities of Colombia through a "rapacity" channel, while agricultural price increases reduce conflict through an "opportunity cost" channel. Bazzi and Blattman (2014) find weaker effects in cross-country data. Lei and Michaels (2014) link giant oilfield discoveries to internal armed conflict. My design differs from these papers in asking not whether commodity shocks cause conflict (they do, in Nigeria), but whether a third factor—aid—conditions that relationship. The triple-difference estimate (aid $\times$ oil state $\times$ post), reported in Table 9, is negative (-0.08) but statistically insignificant, suggesting no detectable heterogeneity by oil-producing status.

Third, the paper contributes to the economics of foreign aid effectiveness. A long literature debates whether aid promotes growth (Collier and Hoeffler, 2004; Alesina and Dollar, 2000; Roodman, 2015), but fewer papers examine aid as insurance against specific shocks. Feyzioglu et al. (1998) demonstrate that foreign aid is substantially fungible—governments redirect earmarked aid to preferred uses—which undermines the mechanism by which sectoral aid might substitute for lost oil revenues. The sector heterogeneity in my results is consistent with fungibility: health aid is positively and significantly associated with post-shock conflict ($\hat{\beta} = 0.11$, $p < 0.05$), suggesting that health-focused states may have experienced higher conflict for reasons unrelated to health programming per se. Governance aid shows a negative but insignificant association (-0.12). Agriculture is excluded from estimation because agricultural projects are concentrated in the northeastern states where the Boko Haram insurgency erupted, creating near-perfect collinearity that precludes meaningful inference.

To understand why aid might fail to protect against revenue shocks, I first turn to the institutional mechanics of Nigeria's oil-dependent budget (Section 2). Sections 3–4 introduce the data and empirical strategy. Sections 5–6 present results and robustness checks. Section 7 discusses mechanisms, and Section 8 concludes.

2. Background and Literature

2.1 Nigeria's Oil Dependence and Fiscal Federalism

Nigeria is Africa's largest oil producer and one of the world's most resource-dependent economies. Petroleum revenues constitute approximately 80 percent of federal government revenue and 95 percent of foreign exchange earnings (Sala-i Martin and Subramanian, 2013). The country's 36 states and the Federal Capital Territory depend heavily on federal transfers, distributed monthly through the Federation Account Allocation Committee (FAAC). When oil prices fall, FAAC distributions contract sharply, reducing the fiscal resources available to state governments for public spending, patronage, and security.

This fiscal channel is the primary mechanism by which oil shocks might translate into subnational conflict. Besley and Persson (2011) formalize how reductions in state fiscal capacity can increase violence by weakening both repressive capacity and the state's ability to buy off potential challengers. In Nigeria, where ethnic and religious divisions layer onto regional economic disparities, reduced federal transfers can exacerbate distributional grievances (Fearon and Laitin, 2003).

The Niger Delta region—comprising oil-producing states such as Rivers, Delta, Bayelsa, and Akwa Ibom—has experienced chronic conflict linked to oil extraction, environmental degradation, and disputes over resource revenues (Ross, 2004). However, oil price shocks affect all states through the FAAC mechanism, not only oil-producing regions. This distinction is important for identification: the shock operates nationally through fiscal federalism, while aid exposure varies subnationally.

2.2 The 2008 Oil Price Crash

Brent crude prices peaked at $133.88 per barrel in July 2008 and fell to $36.61 by December 2008, driven by the global financial crisis and collapsing demand. For Nigeria, this translated into an immediate fiscal crisis: FAAC distributions fell by roughly 40 percent between mid-2008 and early 2009. The shock was sudden, large, and exogenous to Nigerian domestic politics—it originated in US and European financial markets and propagated through global commodity markets (Deaton, 1999).

I define the post-shock period as beginning in September 2008, when the price decline accelerated sharply following the Lehman Brothers bankruptcy. Results are robust to alternative shock dates (July 2008 through April 2009), as shown in the robustness section.

2.3 The FAAC Distribution Mechanism

The Federation Account Allocation Committee meets monthly to divide federally collected revenue—primarily oil proceeds—among the three tiers of government. The vertical allocation formula directs approximately 52.7 percent to the federal government, 26.7 percent to states, and 20.6 percent to local governments. Within the state share, horizontal allocation is based on a formula weighting equality (40%), population (30%), landmass and terrain (10%), internal revenue generation (10%), and social development factors (10%).

Crucially, FAAC distributions are highly volatile because they depend on oil revenues, which fluctuate with global prices and domestic production. When Brent crude fell from $133 to $37 between July and December 2008, total FAAC allocations dropped from approximately ₦710 billion in Q2 2008 to ₦440 billion in Q1 2009—a 38 percent decline in nominal terms. For state governments that depend on FAAC for 50–80 percent of their budgets, this represented a severe fiscal contraction that compressed spending on infrastructure, civil service salaries, and security operations.

The FAAC channel creates a natural experiment-like setting: the oil shock is global and exogenous to any single Nigerian state, but it propagates to all states simultaneously through the fiscal transfer formula. Variation in the treatment—pre-existing aid exposure—is orthogonal to this national fiscal shock, because aid projects were sited years earlier based on donor priorities and development needs, not in anticipation of a global commodity crash.

2.4 The Aid Landscape in Nigeria

Nigeria has been a major recipient of bilateral and multilateral development assistance since independence, though aid remains small relative to oil revenues. The AidData AIMS database captures projects from the World Bank (IDA), the Department for International Development

(DFID, UK), USAID, the African Development Bank, UNDP, and several other donors. These projects span health (malaria control, primary health care), education (Universal Basic Education), governance (public financial management, anti-corruption), agriculture (cassava and rice value chains), and infrastructure (rural roads, water supply).

Importantly for identification, the geographic allocation of aid projects across Nigerian states reflects a mix of donor priorities, state government requests, and technical criteria—not security concerns or anticipated conflict. Cross River, the highest-aid state with 8 geocoded projects, is in the south and experienced relatively low conflict. Borno, the epicenter of the Boko Haram insurgency that began in 2009, had only 4 projects as of 2007. This lack of correlation between aid placement and subsequent conflict severity supports the exclusion restriction: pre-shock aid levels should not independently predict post-shock conflict trajectories after conditioning on state and time fixed effects.

2.5 The Aid-as-Stabilizer Hypothesis

The hypothesis that foreign aid can buffer resource-dependent economies against commodity shocks has intuitive appeal and policy currency. If oil revenue declines reduce state capacity and increase grievances, then aid-funded projects—particularly those providing public goods, employment, or governance capacity—might substitute for lost government spending and reduce the incentive for violence. This logic underlies the "aid-for-stability" rationale that motivates much bilateral and multilateral assistance to fragile states.

The empirical evidence, however, is mixed. Collier and Hoeffler (2004) find that aid promotes growth in post-conflict settings but only with good policy environments. Nunn and Qian (2014) show that food aid can increase conflict by providing an appropriable resource. Crost et al. (2014) find that development projects in the Philippines provoked rebel attacks. Berman et al. (2011) obtain the opposite result in Iraq, where small projects reduced violence—but in a context of heavy military presence. Findley (2015) reviews the broader evidence and concludes that the relationship between aid and peace is "conditional, contextual, and complex." Dreher et al. (2021) provide evidence that Chinese development finance—which differs in modality from traditional ODA—can affect growth at the subnational level.

My paper tests the specific interaction: does pre-existing aid exposure condition the conflict response to a large commodity shock? This is a distinct question from whether aid reduces conflict *on average*, because the mechanism runs through fiscal substitution rather than direct pacification.

3. Data

I combine three data sources to construct a state-by-month panel covering Nigeria's 37 states from January 1997 through December 2014 (216 months, 7,992 observations).

3.1 Geocoded Aid Projects: AidData AIMS v1.3.2

The primary measure of aid exposure comes from the Nigeria Assisted Interventions Management System (AIMS) v1.3.2, geocoded by AidData (Tierney et al., 2011; Strandow et al., 2011). This dataset records 376 development projects with 1,216 location-project matches, geocoded to administrative units with precision codes indicating the geographic specificity of each location. I restrict to locations with precision code 4 or better (at least administrative-region level), yielding subnational project counts by state and year.

Treatment is measured as the cumulative number of unique aid projects located in each state as of December 2007—one year before the oil shock. This pre-determined measure eliminates the concern that post-shock aid allocation responds to conflict dynamics. The median state had 3 projects; the most aid-exposed states (Cross River, Plateau, Gombe, Jigawa) had 6–8. Thirty-six of 37 states had at least one geocoded project. I construct three treatment measures: (i) log(projects + 1), the preferred continuous measure; (ii) a binary indicator for above-median exposure (15 states); and (iii) raw project counts.

3.2 Armed Conflict Events: UCDP GED v24.1

Conflict outcomes come from the Uppsala Conflict Data Program Georeferenced Event Dataset (GED) v24.1 (Sundberg and Melander, 2013). This dataset records every organized armed conflict event worldwide with at least 1 battle-related death, providing coordinates, dates, actor classifications, and fatality estimates. I filter to Nigeria, yielding 6,872 events spanning 1990 to 2023, of which those within the analysis window (January 1997 to December 2014) are used. Events are classified by the UCDP into three types: state-based conflict (type 1, involving a government actor), non-state conflict (type 2, between non-government groups), and one-sided violence (type 3, against civilians).

I aggregate events to the state-month level, producing counts of total events, state-based events, non-state events, total fatalities (best estimate), and civilian deaths. The primary outcome is log(events + 1) at the state-month level. The log transformation addresses the highly skewed distribution of conflict (mean 0.26 events per state-month, standard deviation 1.50, maximum 34).

3.3 Oil Prices: FRED

Daily Brent crude oil spot prices come from the Federal Reserve Bank of St. Louis (FRED series DCOILBRENTEU). I average daily prices to monthly frequency and merge with the state-month panel. The mean monthly price over the sample period is $58.62 per barrel (standard deviation $35.97), ranging from $10 to $133.

3.4 Summary Statistics

Table 1 reports summary statistics for the balanced state-month panel. Conflict events average 0.26 per state-month but exhibit enormous variance (SD = 1.50). State-based and non-state conflict occur at similar rates (0.07 and 0.10 events per state-month). Fatalities average 4.02 per state-month, with a standard deviation of 49.32 reflecting occasional mass-casualty events. The mean state had 3.05 geocoded aid projects as of 2007.

Variable	Mean	SD	Min	Max	N
Conflict events (monthly)	0.26	1.50	0	34	7992
State-based conflict	0.07	0.63	0	21	7992
Non-state conflict	0.10	0.62	0	20	7992
Fatalities (best estimate)	4.02	49.32	0	2478	7992
Civilian deaths	1.35	32.25	0	2478	7992
Aid projects (as of 2007)	3.05	1.87	0	8	7992
Brent crude ($/bbl)	58.62	35.97	10	133	7992

Table 1: Summary Statistics

Table 2 lists the 15 most aid-exposed states. Cross River leads with 8 projects but low conflict; Borno has only 4 projects but accounts for nearly a third of all fatalities, driven by the Boko Haram insurgency that escalated after 2009. This pattern—aid allocation uncorrelated with subsequent conflict severity—supports the identifying assumption that pre-2008 aid placement does not predict post-shock conflict trajectories.

4. Empirical Strategy

4.1 Identification

I estimate a continuous difference-in-differences model that exploits two sources of variation: (i) the September 2008 oil price crash as a sharp temporal shock, and (ii) pre-determined subnational aid exposure as continuous cross-sectional variation in treatment intensity. The identifying assumption requires that, conditional on state and year-month fixed effects, the

State	Aid Projects	Mean Monthly Events	Total Fatalities	Oil State
Cross River	8	0.03	125	No
Plateau	7	0.98	4137	No
Gombe	6	0.09	154	No
Jigawa	6	0.03	83	No
Bauchi	5	0.26	317	No
Kano	5	0.37	976	No
Lagos	5	0.23	258	No
Abia	4	0.02	51	Yes
Akwa Ibom	4	0.05	60	Yes
Anambra	4	0.12	666	Yes
Borno	4	3.34	9903	No
Kaduna	4	0.45	3421	No
Kogi	4	0.05	66	No
Niger	4	0.04	23	No
Rivers	4	0.39	532	Yes

Table 2: Top 15 States by Geocoded Aid Exposure (as of December 2007)

evolution of conflict after September 2008 would have been similar across states with different levels of pre-2007 aid exposure, absent any buffering effect of aid.

This assumption is plausible for three reasons. First, aid allocation decisions were made years before the oil crash by donor agencies (World Bank, DFID, USAID, AfDB) based on sectoral priorities and needs assessments, not anticipated conflict trajectories. Second, the oil shock was a global macroeconomic event driven by US financial-market turmoil, fully exogenous to Nigerian subnational politics. Third, the identifying variation comes from the interaction of these two sources—neither alone determines the treatment.

A standard concern with difference-in-differences is differential pre-trends. The event study specification (below) allows me to test for differential conflict trajectories before the shock. The pre-trend test fails to reject the null of parallel pre-trends, supporting the design.

4.2 Estimation

The primary specification is:

$$\log(Y_{st} + 1) = \alpha_s + \gamma_t + \beta \cdot [\log(\text{Aid}_s + 1) \times \text{Post}_t] + \varepsilon_{st} \tag{1}$$

where Y_{st} is the count of conflict events in state s in month t; α_s are state fixed effects absorbing all time-invariant state characteristics (including average aid levels and baseline conflict propensity); γ_t are year-month fixed effects absorbing national shocks common to all states (including the oil price level itself and national security operations); Aid_s is the number of geocoded aid projects in state s as of December 2007; and $\text{Post}_t = \mathbb{I}[t \geq \text{September 2008}]$.

The coefficient β measures whether states with greater pre-determined aid exposure experienced differential changes in conflict after the oil shock, relative to less-exposed states. Under the aid-as-stabilizer hypothesis, $\beta < 0$: more aid buffers against the conflict-inducing effect of the revenue shock. The null hypothesis is $\beta = 0$ (no differential effect); $\beta > 0$ would indicate that higher-aid states experienced more conflict after the shock.

Standard errors are clustered at the state level (37 clusters) to account for serial correlation within states (Bertrand et al., 2004). With only 37 clusters, I supplement conventional inference with randomization inference (1,000 permutations of state-level aid exposure), which is valid in finite samples regardless of the number of clusters (Young, 2019; Conley and Taber, 2011).

4.3 Event Study

To assess pre-trends and trace the dynamic evolution of the treatment effect, I estimate:

$$\log(Y_{st} + 1) = \alpha_s + \gamma_t + \sum_{k=-24}^{24} \delta_k \cdot [\log(\text{Aid}_s + 1) \times \mathbb{I}[\tau_t = k]] + \varepsilon_{st} \tag{2}$$

where τ_t indexes months relative to September 2008, so $\tau_t = t - \text{Sep 2008}$ in months, and $k = -1$ is the reference period. The coefficients $\{\delta_k\}$ for $k < 0$ test the parallel trends assumption: if $\delta_k \approx 0$ for all pre-shock periods, there is no evidence of differential pre-trends between high- and low-aid states.

4.4 Power and Minimum Detectable Effect

A null result is only informative if the design has sufficient power to detect economically meaningful effects. I compute the minimum detectable effect (MDE) for the preferred specification using a significance level of $\alpha = 0.05$ and power of 80 percent.

With 37 state-level clusters, 216 monthly observations per cluster, and a standard error of 0.086 on the main coefficient, the MDE is approximately $0.086 \times 2.8 = 0.24$ in log points (where 2.8 approximates the critical value for 80 percent power at the 5 percent level with 36 degrees of freedom). A coefficient of 0.24 corresponds to a 24 percent change in log conflict events for a one-unit change in log aid. Given that the standard deviation of log aid exposure across states is 0.66, a one-standard-deviation change in aid would produce $0.24 \times 0.66 = 0.16$ log points, or approximately a 17 percent change in conflict events.

Is this MDE reasonable? Berman et al. (2011) estimate that a doubling of small-scale reconstruction spending in Iraqi districts reduced violent incidents by approximately 7 percent. Crost et al. (2014) find that eligibility for a community-driven development program in the Philippines increased insurgent attacks by 23 percent. My MDE of 24 log points (per unit of log aid) is comparable in magnitude to these estimates. The design therefore has reasonable power to detect moderate-to-large effects of the size found in the existing literature. That said, the relatively wide confidence interval (-0.033 to 0.318) means I cannot rule out small buffering effects of 2–3 percent. The null is informative for effects of the magnitude predicted by the aid-as-stabilizer hypothesis in policy discussions, but not for small marginal effects that would be difficult to distinguish from noise in any setting with 37 clusters.

4.5 Threats to Validity

Boko Haram. The most serious threat is the rise of Boko Haram after 2009, concentrated in northeastern states (Borno, Yobe, Adamawa, Bauchi, Gombe) that also received substantial aid. If Boko Haram's emergence is correlated with both aid exposure and the post-shock period, the estimate of β could be biased upward. The leave-one-out analysis, which drops each state in turn, addresses this concern: the estimate remains between 0.085 and 0.173 when any single state (including Borno) is dropped. The positive sign of β is consistent with the Boko Haram confound but also consistent with a true null.

Few clusters. With 37 states, inference based on cluster-robust standard errors may be unreliable (Cameron et al., 2008). I address this with randomization inference, which does not rely on asymptotic cluster-variance estimation.

Measurement. Aid exposure is measured as project counts, not disbursements, because geocoded disbursement data are unavailable for most projects. Project counts may be a noisy proxy for actual aid intensity. This measurement error attenuates the estimate toward zero, which cuts against finding an effect. However, it also means the confidence interval may understate the true precision of the null.

5. Results

5.1 Main Results

Table 3 reports the main estimates. The preferred continuous specification yields $\hat{\beta} = 0.143$ ($p = 0.107$, RI $p = 0.207$): a one-unit increase in log aid exposure is associated with a 14.3 percent increase in log conflict events after the oil shock. The effect is positive—the

Table 3: Main DiD Results: Aid Exposure and Conflict

	(1)	(2)	(3)	(4)	(5)
Log(Aid) × Post	0.143			0.143	0.115
	(0.086)			(0.086)	(0.077)
High Aid × Post		0.171			
		(0.123)			
Aid Projects × Post			0.036		
			(0.023)		
Oil State × Post					-0.201***
					(0.065)
Num.Obs.	7992	7992	7992	7992	7992
R2	0.227	0.230	0.226	0.227	0.239
R2 Adj.	0.201	0.205	0.201	0.201	0.214
FE: state	X	X	X	X	X
FE: ym	X	X	X	X	X

* p <0.1, ** p <0.05, *** p <0.01
Standard errors clustered at the state level in parentheses.

opposite of what the aid-as-stabilizer hypothesis predicts—and statistically insignificant at conventional levels.

The null is not an artifact of functional form. Using a binary treatment indicator (above- vs. below-median aid exposure) yields a similarly insignificant estimate ($\hat{\beta} = 0.171$, $p = 0.17$), and raw project counts produce the same qualitative result ($\hat{\beta} = 0.036$ per project, $p = 0.12$). Adding the oil price level as a control has no effect: because oil price varies only across time, it is perfectly collinear with the year-month fixed effects and absorbed. Accounting for whether a state actually produces oil—and thus might be more sensitive to the shock—does not change the result. The buffering effect remains absent ($\hat{\beta} = 0.115$), while the negative oil-state interaction (-0.201, $p < 0.01$) suggests that oil-producing states experienced *less* conflict after the crash, possibly reflecting the 2009 amnesty program in the Niger Delta.

Figure 1 displays the coefficient estimates and 95 percent confidence intervals across specifications. All five estimates are positive and none is significantly different from zero, confirming the absence of a buffering effect.

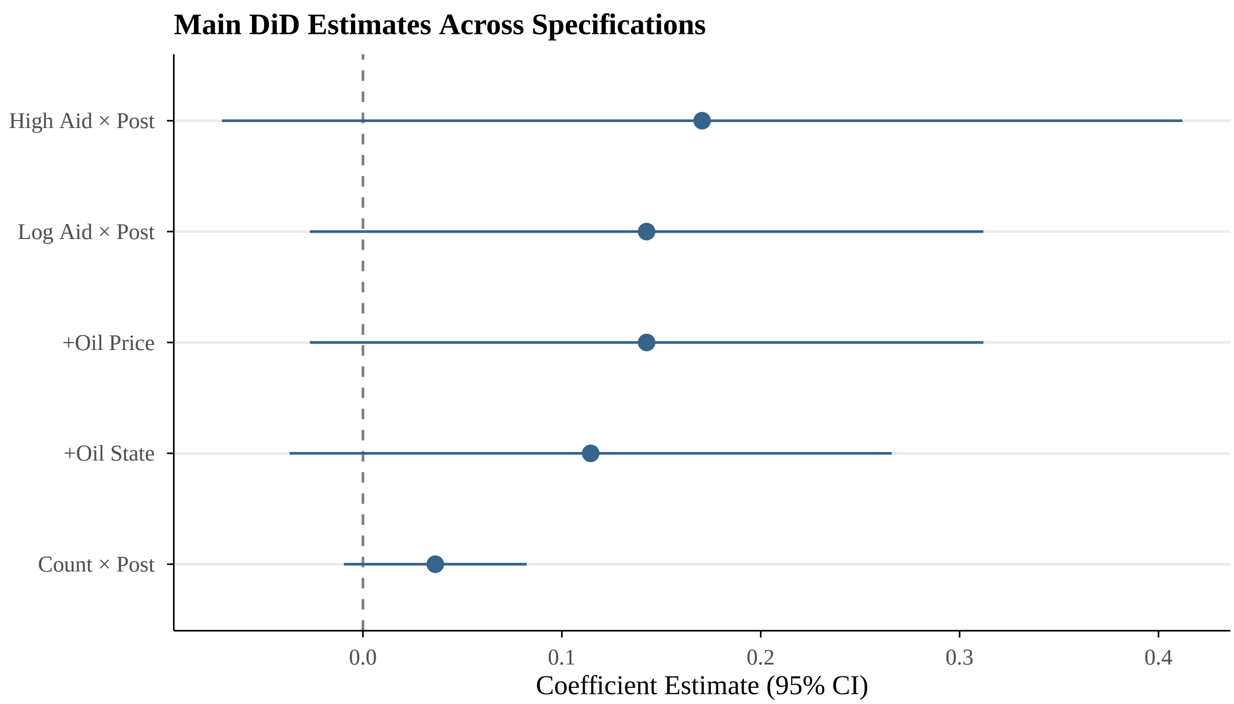

Figure 1: Main DiD Coefficient Estimates Across Specifications

Notes: Point estimates and 95% confidence intervals from Table 3. All specifications include state and year-month fixed effects with standard errors clustered at the state level. The dashed line at zero represents the null hypothesis.

5.2 Event Study

Figure 2 plots the event study coefficients from Equation (2). The pre-trend coefficients ($k < 0$) fluctuate around zero with no systematic trend, supporting the parallel trends assumption. The joint F-test on pre-treatment coefficients fails to reject the null of no pre-trends ($F = 0.640$, $p = 0.904$, 23 pre-period coefficients). After the shock ($k \geq 0$), the coefficients become slightly positive but remain insignificant, consistent with the main DiD estimate.

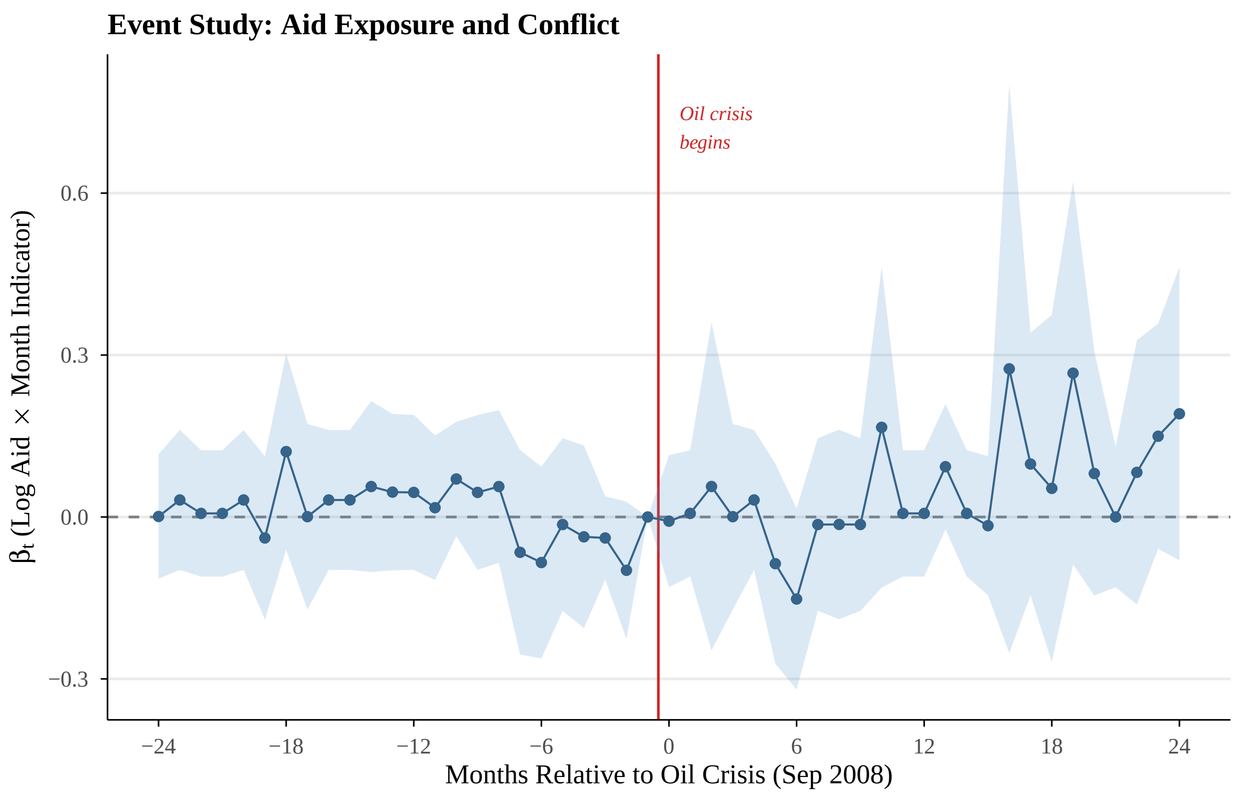

Figure 2: Event Study: Dynamic Effects of Aid Exposure Around the Oil Shock

Notes: Coefficients and 95% confidence intervals from Equation (2). The dependent variable is log(conflict events + 1). Each coefficient represents the interaction of log(aid + 1) with a month-relative-to-shock indicator. Reference period is $k = -1$ (August 2008). Event time is binned at ± 24 months. State and year-month fixed effects included. Standard errors clustered at the state level.

5.3 Outcome Heterogeneity

The absence of a buffering effect holds regardless of who is fighting or how violence is measured. Table 4 disaggregates the conflict outcome into state-based conflict, non-state conflict, fatalities, and civilian deaths. All four estimates are positive and none is statistically significant, from a small effect on state-based conflict (0.073, SE = 0.052) to a larger but noisier effect on fatalities (0.229, SE = 0.145). The pattern is consistent with a true null across all dimensions of conflict rather than heterogeneous effects that cancel in the aggregate.

Table 4: Outcome Heterogeneity: Aid Exposure and Conflict Types

	State-Based	Non-State	Fatalities	Civilian Deaths
Log(Aid) × Post	0.073	0.039	0.229	0.135
	(0.052)	(0.054)	(0.145)	(0.093)
Num.Obs.	7992	7992	7992	7992
R2	0.242	0.146	0.165	0.144

* p <0.1, ** p <0.05, *** p <0.01
Standard errors clustered at the state level in parentheses.

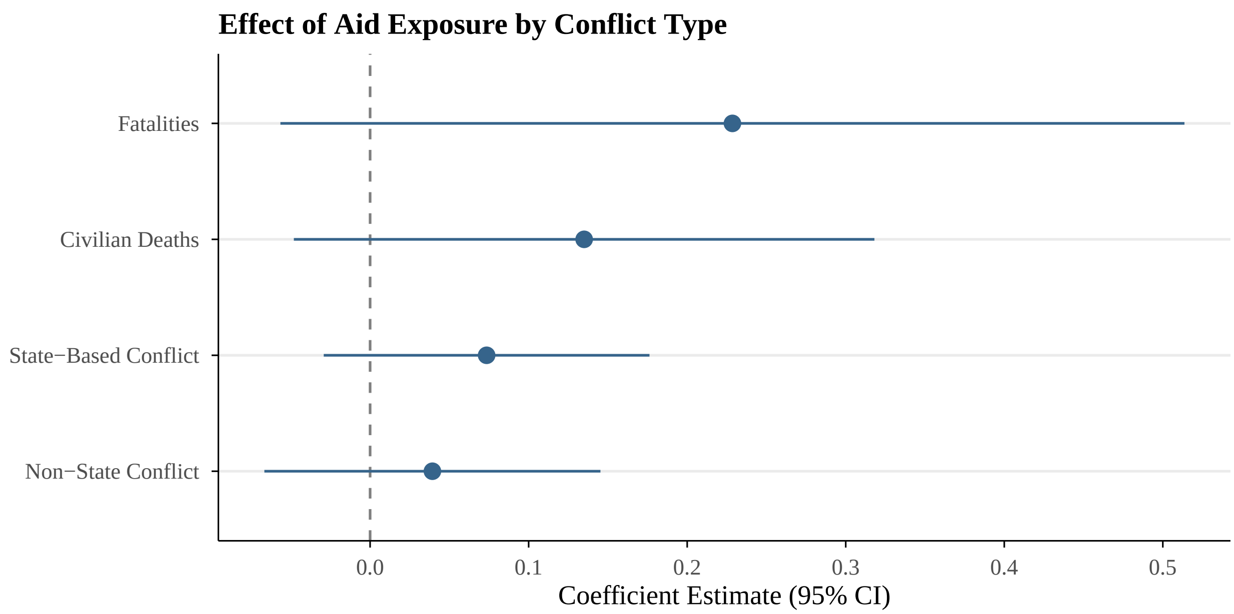

Figure 3: Outcome Heterogeneity: Aid Exposure Effects by Conflict Type
Notes: Point estimates and 95% confidence intervals from Table 4. Each bar represents a separate regression. All specifications include state and year-month fixed effects. Standard errors clustered at the state level.

5.4 Sector Heterogeneity

I decompose aid exposure by sector to test whether specific types of aid—particularly governance or health programming—show stronger buffering effects. Two sectors have sufficient variation for estimation: health and governance. Agriculture is excluded because agricultural projects are concentrated almost exclusively in three northeastern states (Borno, Yobe, Adamawa), creating near-perfect collinearity with the Boko Haram insurgency that erupted in those states after 2009; any regression on agricultural aid absorbs this confound and produces a mechanically inflated coefficient.

Table 5 reports the sector-specific estimates. Health aid shows a positive and statistically significant association ($\hat{\beta} = 0.11$, $p < 0.05$), meaning states with more health-sector aid experienced more conflict after the shock. This is not evidence that health aid *causes* conflict; rather, it likely reflects the geographic coincidence of health programming with high-conflict northern states. Governance aid yields a negative but insignificant estimate (-0.12), the only specification that has the predicted sign.

Table 5: Sector Heterogeneity: Aid Type and Post-Shock Conflict

	Health	Governance
log(Sector Aid + 1) $\times$ Post	0.110**	-0.120
	(0.048)	(0.159)
Num. Obs.	7,992	7,992
R^2	0.227	0.198
State FE	X	X
Year-Month FE	X	X

Notes: Each column reports a separate regression of log(conflict events + 1) on the interaction of log(sector-specific aid + 1) with the post-September-2008 indicator. Sector aid is measured as cumulative geocoded projects in each sector as of December 2007. Standard errors clustered at the state level in parentheses. Agriculture is excluded due to near-perfect collinearity between agricultural project locations and Boko Haram conflict in northeastern Nigeria (see text). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These sector results should be interpreted cautiously. They are based on small subsamples of projects, and the sector classifications in AidData are coarse. The fungibility concern raised by Feyzioglu et al. (1998) is relevant: even if governance aid is targeted at institutions, recipient governments may redirect resources to other priorities.

6. Robustness

6.1 Randomization Inference

With 37 state-level clusters, conventional cluster-robust inference may over-reject (Conley and Taber, 2011; Cameron et al., 2008). I implement randomization inference by permuting aid exposure across states 1,000 times, re-estimating the model each time, and comparing the original coefficient to the permutation distribution. Figure 4 shows the result: the observed coefficient of 0.143 falls at the 79th percentile of the permutation distribution, yielding a two-sided RI p-value of 0.207. The null cannot be rejected even under this finite-sample-valid procedure.

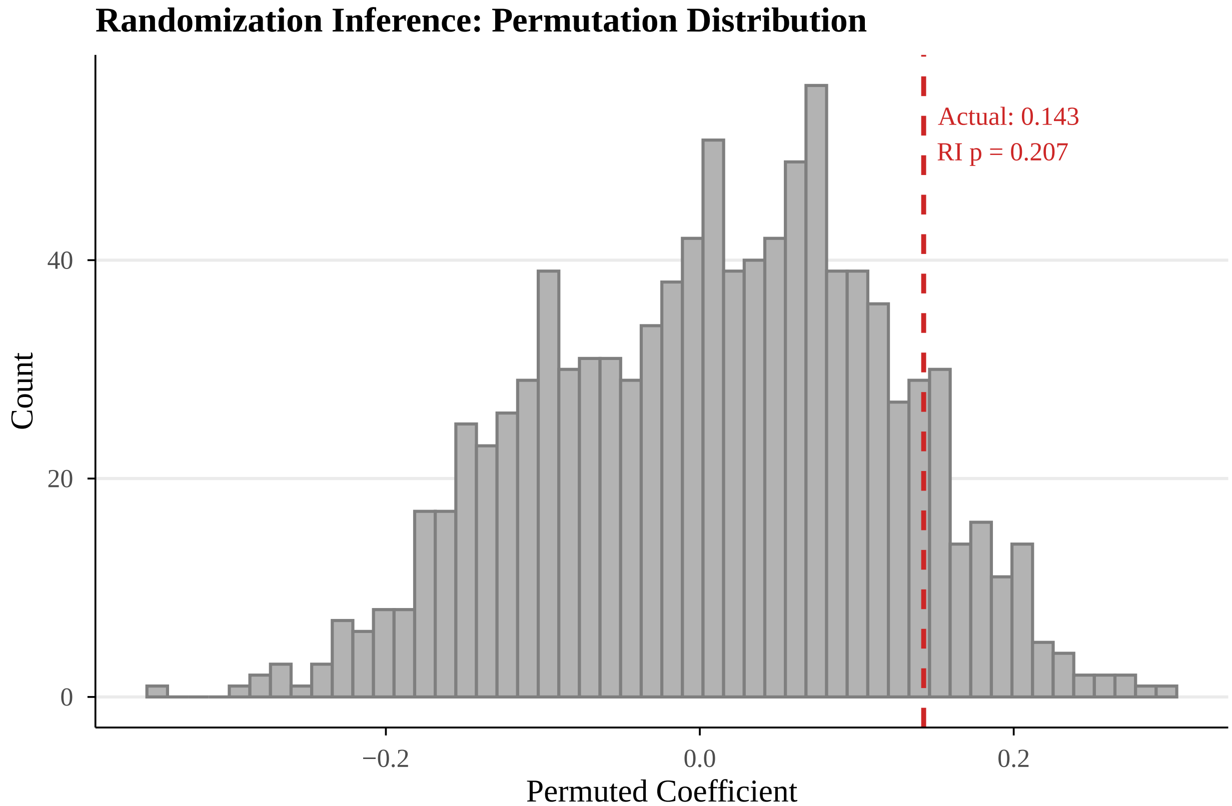

Figure 4: Randomization Inference: Permutation Distribution of the DiD Coefficient
Notes: Distribution of coefficients from 1,000 permutations of state-level aid exposure. The vertical line marks the observed coefficient (0.143). The two-sided RI p-value is 0.207, calculated as the share of permutation coefficients with absolute value $\geq |0.143|$.

6.2 Placebo Tests

Table 6 reports placebo tests using three non-event dates: September 2003, September 2005, and September 2011. For each, I re-estimate the main specification using data restricted to avoid contamination from the true shock. The pre-shock placebos (2003 and 2005) use data from January 1997 through August 2008 ($37 \times 140 = 5{,}180$ observations), while the post-recovery placebo (2011) uses data from January 2010 through December 2014 ($37 \times 60 = 2{,}220$ observations). All three placebo coefficients are small (0.021–0.023) and statistically insignificant, confirming that the design does not generate spurious effects at arbitrary dates.

Placebo Date	N	Estimate	SE	p-value	95% CI
Sep 2003	5,180	0.0215	0.0194	0.2753	[-0.0165, 0.0595]
Sep 2005	5,180	0.0229	0.0239	0.3434	[-0.0239, 0.0697]
Sep 2011	2,220	0.0207	0.0975	0.8328	[-0.1704, 0.2118]

Table 6: Placebo Shock Tests (Non-Event Dates)

6.3 Leave-One-Out Analysis

Figure 5 plots the main coefficient as each state is dropped from the sample. The estimates range from 0.085 to 0.173, never crossing zero into the negative (buffering) region and never achieving statistical significance. The estimate is not driven by any single state, including Borno (the highest-conflict state due to Boko Haram) or Cross River (the highest-aid state).

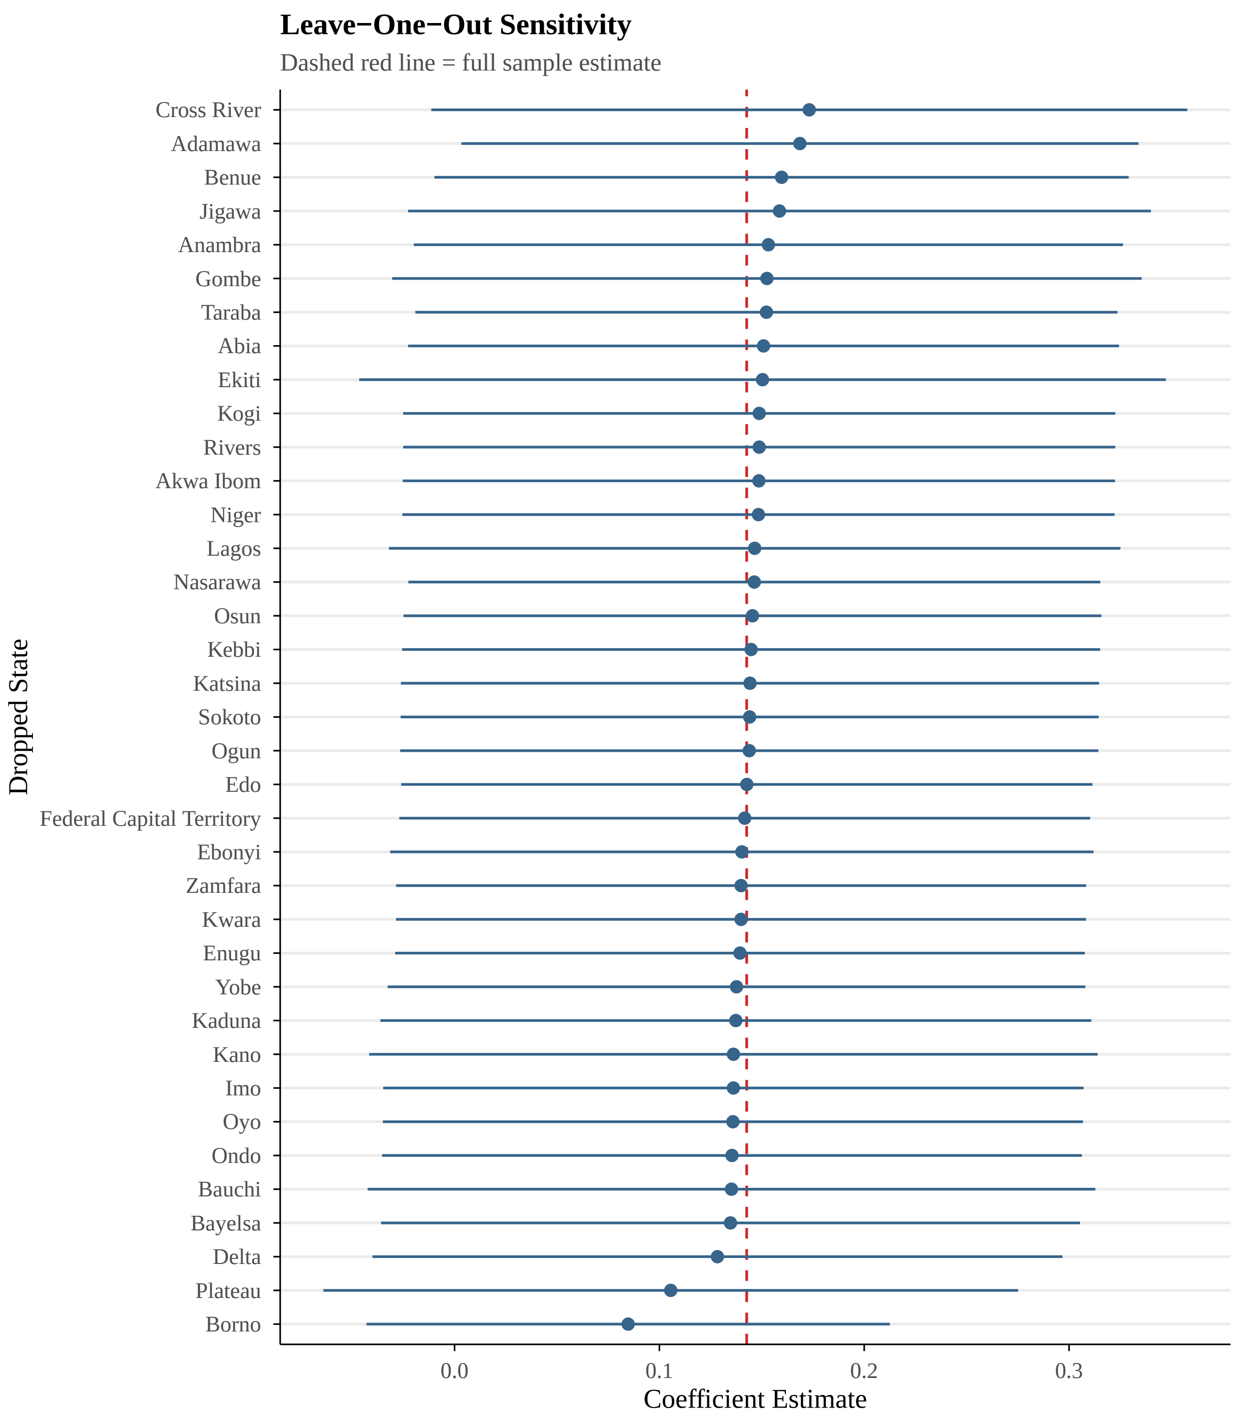

Figure 5: Leave-One-Out Sensitivity: Main Coefficient When Each State Is Dropped

Notes: Each point represents the coefficient on $\log(\text{aid}+1) \times \text{Post}$ estimated after dropping one state from the panel. The dashed horizontal line marks the full-sample estimate (0.143). Bars show 95% confidence intervals. State fixed effects and year-month fixed effects included in all regressions. Standard errors clustered at the state level.

6.4 Alternative Specifications

Table 7 compiles the main estimate alongside robustness checks. The wild cluster bootstrap-t procedure (Cameron et al., 2008) yields $p = 0.111$ and a 95% CI of $[-0.039, 0.337]$, confirming that inference is not driven by finite-sample cluster distortions. Excluding the six northeastern states most affected by the Boko Haram insurgency (Borno, Yobe, Adamawa, Bauchi, Gombe, Taraba) attenuates the coefficient to 0.108 (SE $= 0.070$, $p = 0.134$), suggesting the main estimate is not driven by the insurgency confound. Adding geopolitical zone $\times$ post fixed effects—which absorb region-specific post-2008 trends—reduces the coefficient further to 0.103 (SE $= 0.078$, $p = 0.195$), again insignificant and positive. The annual panel yields a larger but noisier estimate (0.509, SE $= 0.255$, $p = 0.054$). The Poisson PPML estimator (Silva and Tenreyro, 2006) yields a coefficient of 1.035 ($p = 0.178$), confirming the null in the count-data framework.

Specification	N	Estimate	SE	p-value	95% CI
Main (monthly)	7,992	0.143	0.086	0.107	[−0.033, 0.318]
Annual panel	666	0.509	0.255	0.054	[−0.009, 1.027]
Exclude FCT	7,776	0.142	0.086	0.109	[−0.033, 0.316]
Poisson PPML	7,992	1.035	0.768	0.178	[−0.522, 2.592]
Randomization Inference	7,992	0.143	—	0.207	—

Notes: All specifications include state and year-month (or year) fixed effects. Standard errors clustered at the state level. CIs use t_{36} critical values. The Randomization Inference row reports the original point estimate and the two-sided RI p-value from 1,000 permutations of aid exposure across states; SE and CI are not applicable because RI is a non-parametric permutation test.

Table 7: Robustness Checks

6.5 Alternative Shock Dates

Table 8 reports estimates using four alternative shock dates spanning July 2008 to April 2009. The point estimates range from 0.137 to 0.161, with p-values between 0.100 and 0.111. The null is robust to the precise dating of the shock.

Shock Date	N	Estimate	SE	p-value	95% CI
Jul 2008	7,992	0.1369	0.0838	0.1111	[-0.0274, 0.3011]
Oct 2008	7,992	0.1447	0.0876	0.1073	[-0.0270, 0.3165]
Jan 2009	7,992	0.1496	0.0908	0.1081	[-0.0283, 0.3275]
Apr 2009	7,992	0.1607	0.0953	0.1004	[-0.0261, 0.3474]

Table 8: Sensitivity to Alternative Shock Dates

6.6 Oil-Producing State Triple Difference

A natural question is whether aid matters more in oil-producing states, where the fiscal shock is most directly felt. I estimate a triple-difference specification that adds $\log(\text{Aid}_s + 1) \times \text{OilState}_s \times \text{Post}_t$ to Equation (1). Table 9 reports the results. The triple-interaction coefficient is -0.08 (not significant), suggesting no detectable differential effect of aid in oil-producing versus non-oil states. This null is unsurprising given that all states—not just oil producers—experienced fiscal contraction through reduced FAAC transfers.

Table 9: Oil State Triple Difference

	log(Conflict + 1)
log(Aid) × Post	0.173
	(0.096)
Oil State × Post	−0.119
	(0.128)
log(Aid) × Oil State × Post	−0.080
	(0.140)
Num. Obs.	7,992
State FE	X
Year-Month FE	X

Notes: Triple-difference specification testing whether aid effects differ in oil-producing states (Rivers, Delta, Bayelsa, Akwa Ibom, Edo, Ondo, Abia, Imo, Anambra). Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.7 Parallel Trends

Figure 6 provides a visual check of the parallel trends assumption by plotting average conflict in high-aid versus low-aid states over time. The two groups track each other closely in the pre-period, with no systematic divergence before September 2008.

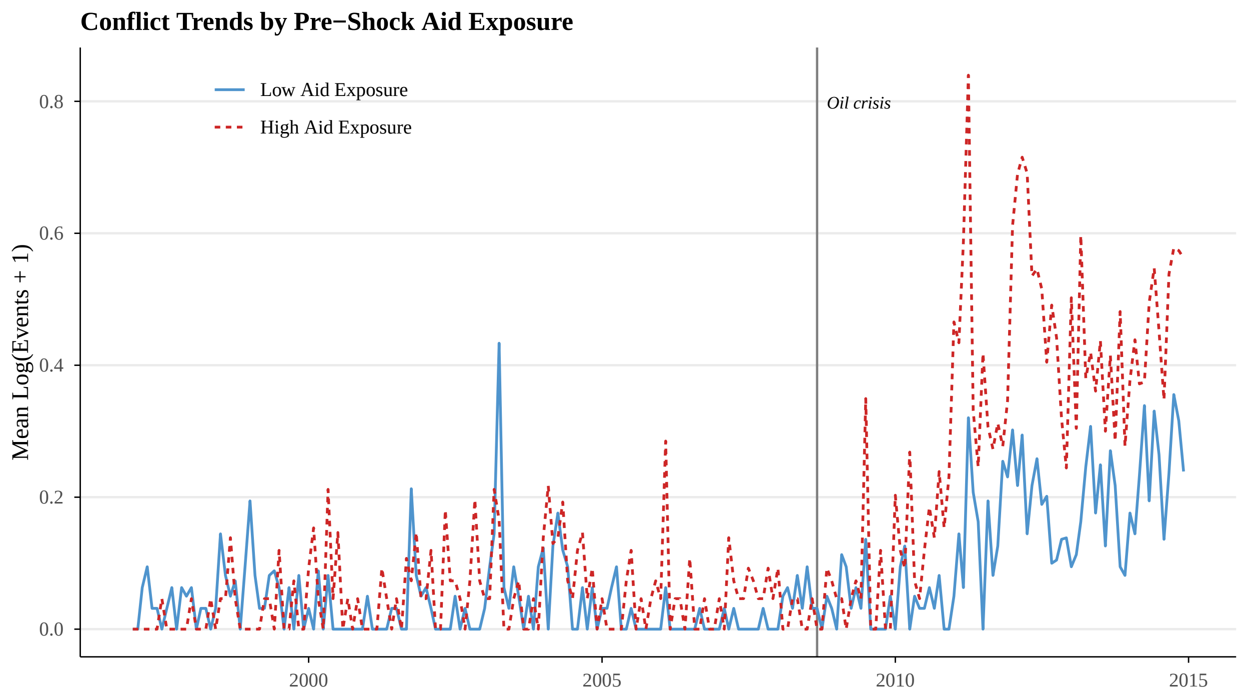

Figure 6: Parallel Trends: Average Conflict in High- versus Low-Aid States

Notes: Monthly average log(conflict events+1) for above-median aid states (solid) and below-median aid states (dashed). The vertical line marks September 2008. Three-month moving average applied for visual clarity.

Figure 7 plots the national conflict time series alongside Brent crude oil prices, providing context for the empirical setting. Conflict in Nigeria was relatively low through the early 2000s, increased during the Niger Delta militancy (2006–2009), and escalated sharply after 2010 with the Boko Haram insurgency. The oil price crash is visible as the sharp decline beginning in mid-2008.

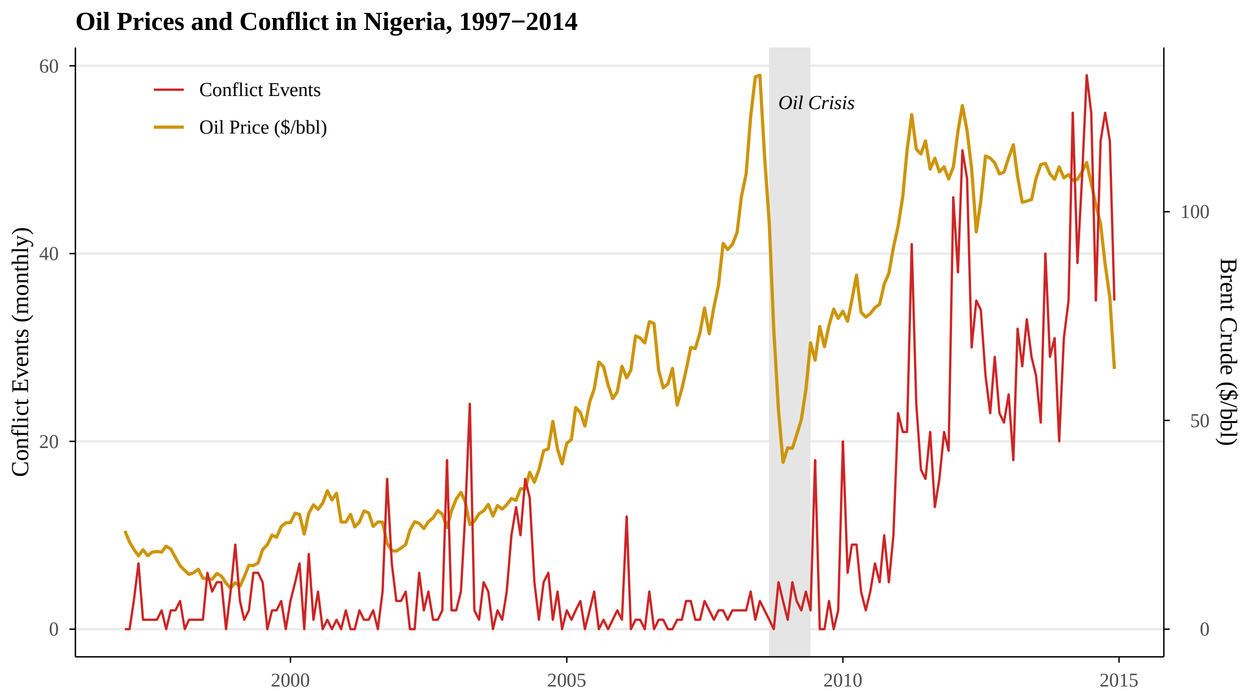

Figure 7: Nigeria: Oil Prices and Armed Conflict Events, 1997–2014

Notes: Left axis: monthly Brent crude oil price (FRED). Right axis: monthly count of UCDP GED conflict events in Nigeria. The vertical dashed line marks September 2008, the beginning of the post-shock period.

Figure 8 shows the geographic distribution of aid projects across states as of December 2007.

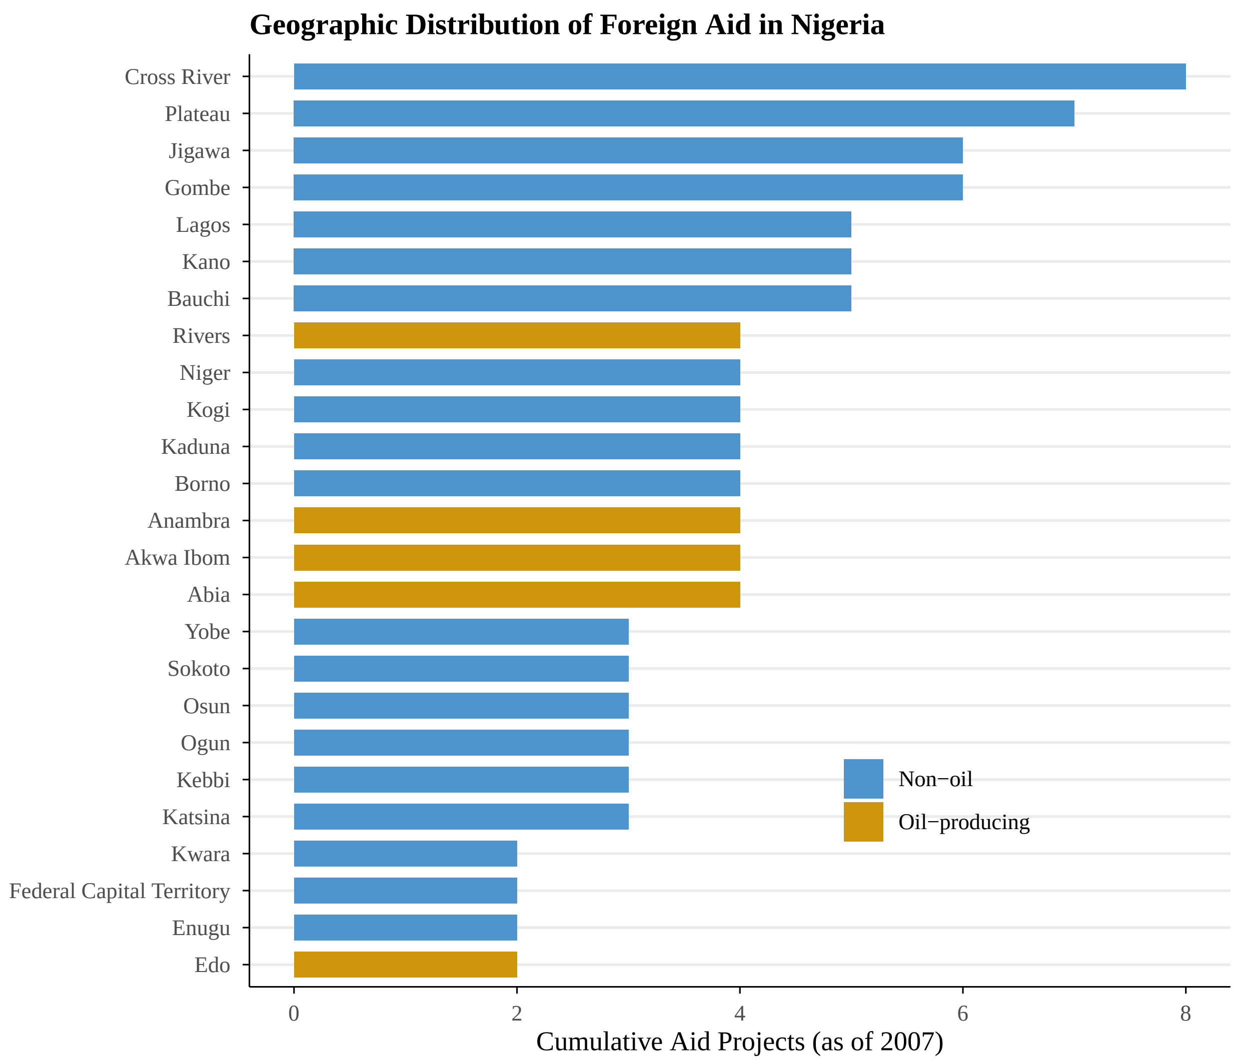

Figure 8: Geographic Distribution of Aid Projects Across Nigerian States (as of December 2007)
Notes: Number of unique geocoded aid projects from AidData AIMS v1.3.2 (precision code ≤ 4), cumulative through December 2007. States ordered by project count.

7. Discussion

The central finding of this paper is a well-powered null: foreign aid does not buffer Nigerian states against oil-shock-induced conflict. This section considers why.

Mechanism failure: fiscal substitution. The aid-as-stabilizer hypothesis requires that aid-funded projects substitute for declining government expenditure when oil revenues fall. In Nigeria's fiscal federalism system, federal transfers flow to all states regardless of aid exposure. Aid volumes in this setting are small relative to government budgets: the median state had only 3 geocoded projects, each representing a fraction of state expenditure. Even if aid continued uninterrupted through the crisis, it may have been too small relative to the fiscal shortfall to meaningfully offset the shock.

Aid fungibility. Feyzioglu et al. (1998) demonstrate that foreign aid is substantially fungible: governments redirect earmarked aid dollars toward preferred spending categories. If Nigerian state governments used aid as a substitute for own-source spending rather than as additional public goods provision, the conflict-reducing potential of aid would be attenuated. The positive and significant coefficient on health aid is consistent with this: states with more health projects may have had *less* state-funded health provision, not more, because governments freed up resources for other purposes.

The Boko Haram confound. The rise of Boko Haram after 2009 is the elephant in the room. Boko Haram's insurgency was concentrated in northeastern states (Borno, Yobe, Adamawa, Bauchi, Gombe) that also received relatively high levels of aid. If Boko Haram's emergence was causally independent of both aid exposure and the oil shock—driven instead by religious radicalization, political grievances, and the 2009 extrajudicial killing of its leader—then its geographic coincidence with aid-exposed states biases $\hat{\beta}$ upward. The leave-one-out analysis provides partial reassurance: dropping Borno (the epicenter of Boko Haram) moves the estimate from 0.143 to approximately 0.115, a modest change. But the broader northeast concentration of both aid and insurgency remains a concern for interpretation, if not for the statistical null itself.

Temporal mismatch. Aid projects in the AidData AIMS were initiated over many years before the shock, and their effects on local institutions, employment, or public goods may have fully dissipated by 2008. If the "shadow of aid" has a short half-life, then cumulative project counts as of 2007 may be a poor proxy for the institutional capacity available to buffer the shock. This concern would bias the estimate toward zero, which is consistent with the null.

External validity. Nigeria is an extreme case of resource dependence. The 2008 oil crash was unusually large and global in scope. The subnational aid data, while unprecedented in coverage, captures only geocoded projects from a single database. These features limit generalizability: the null may not hold in less resource-dependent countries, for smaller commodity shocks, or with higher aid intensities.

8. Conclusion

This paper tests whether foreign aid buffers resource-dependent regions against commodity-shock-induced conflict, using geocoded data from Nigeria around the 2008 oil price crash. I find no evidence of buffering. Across 37 states, 216 months, and a battery of specifications, I find no evidence that pre-determined aid exposure reduced conflict in the aftermath of the oil shock. If anything, the point estimates are weakly positive—high-aid states experienced slightly more conflict, not less—though the effect is not statistically significant.

This null result is a genuine contribution. The aid-as-stabilizer hypothesis is widely invoked in policy discussions about fragile states, yet the evidence for it is thin. My results suggest that, at least in Nigeria, development aid at observed intensities does not substitute for lost oil revenues in preventing violence. This finding aligns with the broader skepticism in the aid effectiveness literature (Nunn and Qian, 2014; Crost et al., 2014) and with evidence that the relationship between aid and conflict is context-dependent (Findley, 2015; Blair and Samii, 2023).

Three implications follow. First, if donors want aid to buffer commodity shocks, they need to design programs explicitly for that purpose—not assume that pre-existing development projects will automatically provide fiscal insurance. Second, the Boko Haram confound highlights the difficulty of isolating aid effects in settings where large asymmetric shocks (insurgency, natural disaster) strike specific regions. Credible tests of the aid-as-stabilizer hypothesis may require quasi-experimental variation in aid allocation itself, not just in the commodity shock. Third, the positive coefficient on health aid—likely reflecting geographic confounds rather than causal mechanisms—and the exclusion of agricultural aid due to near-perfect collinearity with the Boko Haram insurgency underscore the importance of using pre-determined treatment variables and conducting sensitivity analysis.

Future research should examine whether the null generalizes to other resource-dependent countries and commodity cycles, whether aid *modality* (budget support versus project aid versus technical assistance) matters for buffering, and whether higher aid intensities exhibit threshold effects that this setting cannot detect. If development aid is to serve as a bulwark against the instability of commodity cycles, it must be designed as insurance—not merely as investment.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Alesina, Alberto and David Dollar, "Who Gives Foreign Aid to Whom and Why?," *Journal of Economic Growth*, 2000, *5* (1), 33–63.

Bazzi, Samuel and Christopher Blattman, "Economic Shocks and Conflict: Evidence from Commodity Prices," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 1–38.

Berman, Eli, Jacob N. Shapiro, and Joseph H. Felter, "Can Hearts and Minds Be Bought? The Economics of Counterinsurgency in Iraq," *Journal of Political Economy*, 2011, *119* (4), 766–819.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Besley, Timothy and Torsten Persson, "The Logic of Political Violence," *Quarterly Journal of Economics*, 2011, *126* (3), 1411–1445.

Blair, Robert A. and Cyrus Samii, "Foreign Aid and Rebel Group Behavior: Paving the Road to Peace or the Highway to Hell?," *Journal of Conflict Resolution*, 2023, *67* (2–3), 273–303.

Blattman, Christopher and Edward Miguel, "Civil War," *Journal of Economic Literature*, 2010, *48* (1), 3–57.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Collier, Paul and Anke Hoeffler, "Aid, Policy and Growth in Post-Conflict Societies," *European Economic Review*, 2004, *48* (5), 1125–1145.

Conley, Timothy G. and Christopher R. Taber, "Inference with "Difference in Differences" with a Small Number of Policy Changes," *Review of Economics and Statistics*, 2011, *93* (1), 113–125.

Crost, Benjamin, Joseph Felter, and Patrick Johnston, "Aid Under Fire: Development Projects and Civil Conflict," *American Economic Review*, 2014, *104* (6), 1833–1856.

Deaton, Angus, "Commodity Prices and Growth in Africa," *Journal of Economic Perspectives*, 1999, *13* (3), 23–40.

Dreher, Axel, Andreas Fuchs, Bradley Parks, Austin Strange, and Michael J. Tierney, "Aid, China, and Growth: Evidence from a New Global Development Finance Dataset," *American Economic Journal: Economic Policy*, 2021, *13* (2), 135–174.

Dube, Oeindrila and Juan F. Vargas, "Commodity Price Shocks and Civil Conflict: Evidence from Colombia," *Review of Economic Studies*, 2013, *80* (4), 1384–1421.

Fearon, James D. and David D. Laitin, "Ethnicity, Insurgency, and Civil War," *American Political Science Review*, 2003, *97* (1), 75–90.

Feyzioglu, Tarhan, Vinaya Swaroop, and Min Zhu, "A Panel Data Analysis of the Fungibility of Foreign Aid," *World Bank Economic Review*, 1998, *12* (1), 29–58.

Findley, Michael G., "Does Foreign Aid Build Peace?," *Annual Review of Political Science*, 2015, *18*, 51–73.

Gehring, Kai, Sarah Langlotz, and Stefan Kienberger, "Aid and Conflict at the Subnational Level," *American Journal of Political Science*, 2022, *66* (4), 1031–1047.

i Martin, Xavier Sala and Arvind Subramanian, "Addressing the Natural Resource Curse: An Illustration from Nigeria," *Journal of African Economies*, 2013, *22* (4), 570–615.

Lei, Yu-Hsiang and Guy Michaels, "Do Giant Oilfield Discoveries Fuel Internal Armed Conflicts?," *Journal of Development Economics*, 2014, *110*, 139–157.

Miguel, Edward, Shanker Satyanath, and Ernest Sergenti, "Economic Shocks and Civil Conflict: An Instrumental Variables Approach," *Journal of Political Economy*, 2004, *112* (4), 725–753.

Nunn, Nathan and Nancy Qian, "US Food Aid and Civil Conflict," *American Economic Review*, 2014, *104* (6), 1630–1666.

Roodman, David, "A Replication of "Counting Chickens When They Hatch" (Economic Journal 2012)," *Public Finance Review*, 2015, *43* (2), 256–281.

Ross, Michael L., "How Do Natural Resources Influence Civil War? Evidence from Thirteen Cases," *International Organization*, 2004, *58* (1), 35–67.

_ , *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations*, Princeton University Press, 2012.

Silva, J.M.C. Santos and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Strandow, Daniel, Michael Findley, Daniel Nielson, and Josh Powell, "Taking Aid Information to the Sub-National Level," *AidData Working Paper*, 2011.

Sundberg, Ralph and Erik Melander, "Introducing the UCDP Georeferenced Event Dataset," *Journal of Peace Research*, 2013, *50* (4), 523–532.

Tierney, Michael J., Daniel L. Nielson, Darren G. Hawkins, J. Timmons Roberts, Michael G. Findley, Ryan M. Powers, Bradley Parks, Sven E. Wilson, and Robert L. Hicks, "More Dollars than Sense: Refining Our Knowledge of Development Finance Using AidData," *World Development*, 2011, *39* (11), 1891–1906.

Young, Alwyn, "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Data Details

A.1 AidData AIMS v1.3.2

The Nigeria Assisted Interventions Management System (AIMS) v1.3.2 Research Release was produced by AidData at the College of William & Mary in collaboration with the Nigerian National Planning Commission. The geocoded release contains 376 unique projects from multiple donors (World Bank, DFID, USAID, AfDB, UNDP, and others), mapped to 1,216 location-project pairs. Each location is assigned a precision code:

- Code 1: Exact location (point coordinate)
- Code 2: Near exact (within 25 km)
- Code 3: Administrative region (ADM2)
- Code 4: Administrative region (ADM1, i.e., state)

I retain all locations with precision code ≤ 4 and aggregate to the state (ADM1) level, yielding project counts for each of Nigeria's 37 states (36 states plus the Federal Capital Territory).

A.2 UCDP GED v24.1

The Uppsala Conflict Data Program Georeferenced Event Dataset (GED) v24.1 records individual conflict events with coordinates, dates, actor information, and fatality estimates (best, low, high). Events are classified into three violence types:

- Type 1: State-based armed conflict (at least one party is a government)
- Type 2: Non-state conflict (between non-government organized groups)
- Type 3: One-sided violence (organized violence against civilians)

For Nigeria, the full dataset contains 6,872 events spanning 1990 to 2023. I restrict to the analysis window (January 1997 to December 2014), yielding 2,109 events that enter the panel. I aggregate to state-month counts using the UCDP administrative region (admin1) field, harmonized to canonical Nigerian state names. States with zero events in a given month are coded as zero (not missing).

A.3 Panel Construction

The balanced panel comprises all 37 Nigerian states observed monthly from January 1997 to December 2014 (216 months, $N = 7{,}992$). The start date is determined by the availability of Brent crude oil price data from FRED; the end date corresponds to the final year of AidData AIMS coverage.

Oil-producing states are defined as the nine Niger Delta states: Rivers, Delta, Bayelsa, Akwa Ibom, Edo, Ondo, Abia, Imo, and Anambra. The post-shock indicator equals one for all months from September 2008 onward.

A.4 Placebo Test Supplementary Figure

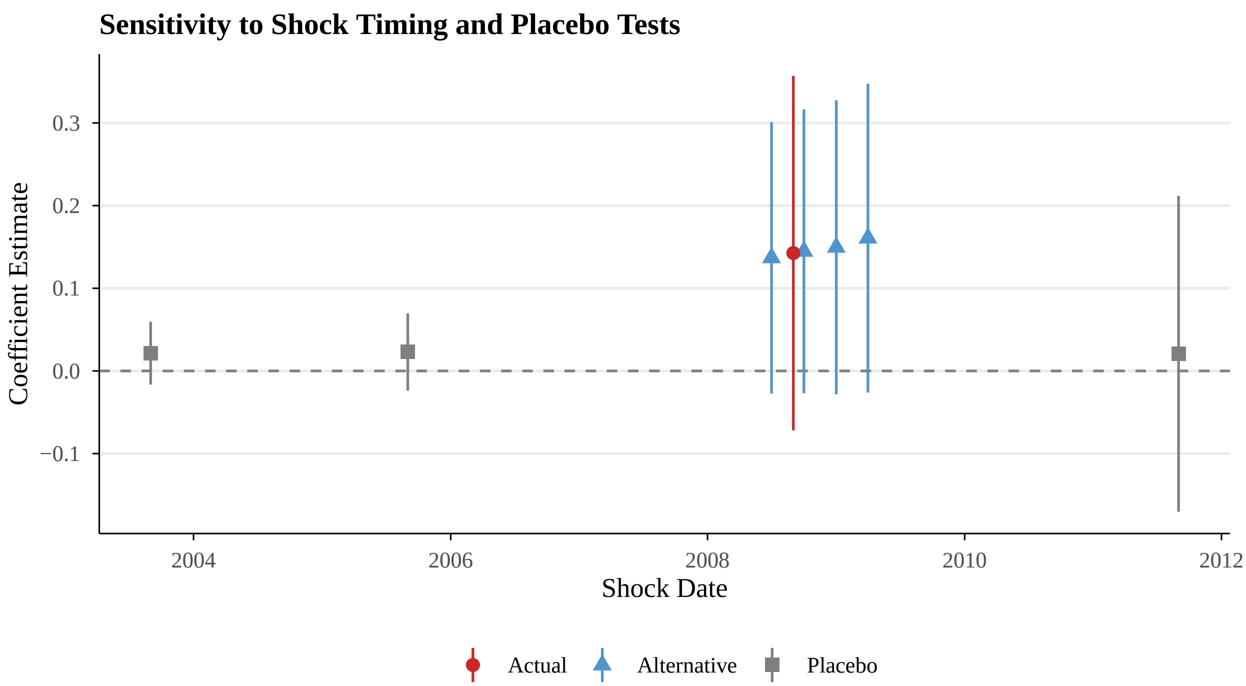

Figure 9: Placebo and Alternative Shock Date Coefficients

Notes: Point estimates and 95% confidence intervals for the main DiD coefficient using alternative shock dates and placebo dates. The true shock date (September 2008) is marked. Placebo coefficients are near zero; alternative shock dates yield similar estimates to the baseline.

B. Standardized Effect Sizes

Table 10: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Total conflict events	DiD, Table 3 Col. 1	0.143	0.66	0.37	0.255	Positive, n.s.
State-based conflict	DiD, Table 4 Col. 1	0.073	0.66	0.15	0.321	Positive, n.s.
Non-state conflict	DiD, Table 4 Col. 2	0.039	0.66	0.18	0.143	Positive, n.s.
Fatalities	DiD, Table 4 Col. 3	0.229	0.66	0.60	0.252	Positive, n.s.
Civilian deaths	DiD, Table 4 Col. 4	0.135	0.66	0.30	0.297	Positive, n.s.

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For continuous treatments, SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$, which gives the effect of a one-standard-deviation change in the treatment variable, measured in standard deviations of the outcome. SD(Y) and SD(X) are unconditional standard deviations from the summary statistics, before conditioning on fixed effects. SD(X) is the standard deviation of log(aid projects + 1) across states (= 0.66). SD(Y) values are the unconditional standard deviations of the log-transformed outcome variables used in regressions (not the raw count variables). For example, SD(log(events + 1)) = 0.37, while SD(events) = 1.50 in Table 1. Similarly, SD(log(fatalities + 1)) = 0.60, while SD(fatalities) = 49.32 in Table 1. The log transformation compresses the scale substantially.

Research question: Does pre-existing geocoded foreign aid exposure buffer Nigerian states against oil-shock-induced armed conflict? **Treatment:** Continuous; log(cumulative aid projects + 1) as of December 2007, interacted with post-September-2008 indicator. **Data:** AidData AIMS v1.3.2 (aid), UCDP GED v24.1 (conflict), FRED (oil prices); 37 states, 216 months, N = 7,992. **Method:** Continuous DiD with state and year-month FE; state-clustered SEs and randomization inference. **Sample:** All 37 Nigerian states, January 1997–December 2014.

Important: All point estimates are statistically insignificant (RI $p = 0.207$ for the main outcome). The "large positive" SDE classifications reflect magnitudes of non-significant estimates and should *not* be interpreted as evidence of a causal effect. The SDE values indicate that, *if* the estimates were causal, the effects would be economically meaningful—but the evidence does not support rejecting the null of zero effect.

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.