Economics School of Louvain - ESL

Economics School of Namur - ESN

Does Mineral Exploration Attract Foreign Aid? Subnational Evidence from Africa

A Staggered Difference-in-Differences Analysis on a Continental ADM2 Panel, 1990–2020

Author : Ole Gottheil

Thesis Director : MAYSTADT, Jean-François

Thesis Reader : MARIANI, Fabio

Academic Year 2025–2026

Master in Economics – 120 credits – Focus : Professional focus

Université catholique de Louvain - Economics School of Louvain - Place Montesquieu 3, 1348 Louvain-La-Neuve
Université de Namur - Economics School of Namur - Rempart de la Vierge 8, 5000 Namur

Abstract

The energy transition has placed African critical minerals at the centre of contemporary diplomacy, driving strategic bilateral agreements between major donor governments and producer countries. While the agreements are recent, the resource-access motive they formalise is not. This thesis evaluates whether these national-level dynamics alter subnational allocation rules, testing whether African second-level administrative units (ADM2) receive systematically more foreign aid following the start of mineral exploration. To examine this question, I construct an annual continental panel of 6,279 African ADM2 polygons spanning 1990 to 2020. The dataset integrates geocoded aid disbursements from nineteen Western donors and the World Bank with a georeferenced mineral-exploration registry and global mining land-use polygons. Chinese development finance, for which only commitments are recorded, is analysed separately through its ODA-like component. Leveraging staggered difference-in-differences estimators robust to heterogeneous treatment effects, combined with within-country propensity-score matching, I find a null for total aid. The start of mineral exploration produces no detectable change in a district's subsequent aid receipts. Beneath the composite, the primary Poisson estimator attributes significant positive responses to World Bank aid and production-sector aid, but the secondary Callaway–Sant'Anna estimator finds no response for any composite and pre-trend diagnostics caution against a causal reading of the positive estimates. Chinese ODA on the post-2000 panel shows no response at the ADM2 level under either estimator. Supplementary descriptive country-and-year fixed-effects PPML regressions confirm that contemporaneous mineral status does not robustly predict aid allocation. Collectively, these findings indicate that while major donors increasingly negotiate high-level strategic agreements for resource access, there is no evidence in these data of a broad, localised targeting of recorded aid to mineral-bearing districts.

Contents

1 Introduction

Mineral endowments leave a strong imprint on the development trajectories of resource-rich countries. Resource abundance can raise growth where institutions and policy align, or lower it through Dutch-disease, weaker institutions and rent-seeking (Cust & Poelhekke, 2015; Ploeg, 2011; Sachs & Warner, 2001), and even the *announcement* of a discovery can distort fiscal and investment decisions before rents arrive (Cust & Mihalyi, 2017). At the subnational level, mineral activity is a documented driver of local conflict (Berman et al., 2017) and resource rents fuel observable regional favouritism in patronage where political institutions are weak (Hodler & Raschky, 2014).

These long-standing domestic and subnational vulnerabilities are increasingly intersecting with shifting international strategic priorities. Driven by the global energy transition, critical minerals have emerged as a central focus of contemporary diplomacy with African producer nations. Projections for electric vehicles, grid storage and renewable-energy infrastructure imply rapidly rising demand for cobalt, lithium, copper, nickel, manganese, graphite and rare-earth elements (International Energy Agency, 2025), and Africa holds a disproportionate share of known reserves and production. Major donor governments have responded with a series of explicit strategic agreements that bundle development finance with resource access, from the US-brokered Critical Minerals for Security and Peace agreement (the Washington Accords) and the Lobito Corridor compact to the European Union's Global Gateway raw-materials partnerships, while the DRC and China renegotiated the 2008 Sicomines minerals-for-infrastructure deal in March 2024. These moves sit against a heavily concentrated, largely Chinese-controlled supply side. Section 2 reviews these agreements, and the supply-chain concerns behind them, in detail. Arezki & Ploeg (2025) capture the dynamic as a contemporary super-power race for critical minerals echoing the 19th-century European scramble for African resources, with developing-country producers exposed to the same risks of capture and rent-seeking that produced the classical resource curse.

The cross-country evidence gives a clear starting point, and it shows that the behaviour the recent agreements formalise is not new. Arezki et al. (2024) show, on a 47-year donor–recipient panel, that recipient countries receive roughly 36% more bilateral aid after a major mineral discovery, consistent with donors competing for resource access. A sample split attributes the effect to the second half of their sample period, which overlaps the 1990–2020 window studied here. The resource-access motive that today's agreements state openly was thus already moving aid allocations, more quietly, during the study period. What this evidence cannot show is where the additional aid lands inside the recipient country. Mineral activity transforms specific districts, not the country as a whole, so the natural question is whether the aid response reaches the mineral-bearing districts themselves or is absorbed elsewhere in the country. That within-country allocation question is the gap this thesis tests. Section 2.3 develops the allocation channels that determine whether the country-level increment should, in theory, reach the producing district at all.

The research question of this thesis is whether African ADM2 units where mineral exploration begins receive systematically more foreign aid than otherwise-comparable never-treated units in the same country. I assemble an annual panel of African ADM2 polygons from 1990 to 2020, covering 50 countries and territories and 6,279 polygons, and combine four data sources: ADM2-level aid disbursements from the Geocoded Official Development

Assistance (ODA) Dataset [GODAD; Bomprezzi et al. (2025)]; the timing of first reported mineral exploration from the USGS Africa minerals compilation (Padilla et al., 2021); the global mining land-use polygons of Maus et al. (2022) for a sample-restriction bias check; and PRIO-GRID v.2.0 (Tollefsen et al., 2012) for geographic and political-economy controls. I use staggered difference-in-differences estimators that are robust to heterogeneous treatment effects (Callaway & Sant'Anna, 2021; Wooldridge, 2023), combined with within-country propensity-score matching to improve covariate balance between treated and never-treated units (Stuart, 2010).

The headline result is a null for total aid. The start of mineral exploration in a district produces no detectable change in the district's subsequent total aid flows. Beneath the composite, the primary Poisson estimator of Wooldridge (2023) attributes statistically significant positive responses to World Bank aid and production-sector aid, but the secondary Callaway & Sant'Anna (2021) estimator applied to the inverse-hyperbolic-sine (asinh) transformed outcome does not reproduce them. Since in addition the event-study leads and the joint pre-trends test flag deviations from parallel trends at longer horizons and the cell-level appendix estimates carry the opposite sign, I treat these positive estimates as suggestive rather than causal. The Chinese-ODA sub-analysis on the post-2000 panel shows no ADM2-level response under either estimator, and is the one specification where the formal pre-trends test comfortably passes. A descriptive country-and-year fixed-effects PPML on the same data (Section 5.3) finds that within-country aid tracks proximity to the capital and local economic size, while contemporaneous mining status adds no explanatory power.

The rest of the thesis proceeds as follows. Section 2 reviews the political-economy literature, develops a conceptual framework of within-country aid allocation, and states the three hypotheses. Section 3 describes the data and panel construction. Section 4 sets out the methodology and the balance comparison between the Total Sample and the PSM-Matched sample. Section 5 reports the empirical findings: the headline analysis, the donor and aid-category heterogeneity (Section 5.2) including the Chinese-ODA sub-analysis, and the descriptive country-and-year fixed-effects PPML correlate check (Section 5.3). Section 6 sets out the limitations that bound the inferential reach of the results. Section 7 concludes. The Annex reports the balance and overlap diagnostics, the Mining-Restricted bias check, additional Callaway–Sant'Anna results, the cell-level results, the Augmented Synthetic Control robustness check, and the estimation diagnostics.

2 Literature Review, Conceptual Framework and Hypotheses

2.1 Donor motivations and aid allocation

This section reviews two adjacent literatures (donor motivations behind aid allocation and the political economy of natural-resource endowments), develops a conceptual framework of who allocates aid within recipient countries, and uses it to motivate the three hypotheses that follow.

Altruistic motivations shaped aid volumes historically (Lumsdaine, 1993), and donors increasingly condition aid on recipient governance quality (Canavire-Bacarreza et al., 2005), but strategic motivations play a role as well. The political economy of foreign aid models the donor as a strategic actor whose allocation rules respond to its own foreign-policy and commercial interests. Alesina & Dollar (2000), using a panel of bilateral aid flows for the 1970–94 period, document that the direction of aid is dictated as much by political and strategic considerations as by recipient need and policy performance. Colonial history and political alliances, notably UN voting alignment, emerge as the dominant determinants. The commercial margin is documented by Berthélemy (2006), who shows that almost all bilateral donors direct more aid to their most significant trading partners, and by Hoeffler & Outram (2011), who show that bilateral donor–recipient trade flows correlate with higher aid disbursements. Both findings fit the broader pattern in which donor governments use aid as a complement to commercial relationships rather than as a substitute for them. Kuziemko & Werker (2006) sharpen the strategic-calculation reading. Temporary membership of the United Nations Security Council is associated with a sizeable, statistically significant increase in aid received from the United States and the United Nations system. The effect grows when the Council is voting on issues important to the United States, and disappears as soon as the rotating seat is given up. That on-off pattern is direct evidence that aid responds to short-run shifts in strategic value rather than only to long-run recipient need. Going further, donors use UN voting alignment as an ongoing screen on commitment (Adhikari, 2019; Kegley & Hook, 1991). The common thread is that donors weigh strategic, political and commercial considerations on top of any humanitarian rationale. The implication, which the thesis takes as its starting point, is that any change in the strategic or commercial value of a recipient unit should, *ceteris paribus*, move that unit's expected aid receipts.

2.2 Natural-resource endowments and donor engagement

The second strand examines how natural-resource endowments shape recipient outcomes and donor engagement. At the cross-country level, Sachs & Warner (2001) provide the canonical "resource curse" finding that resource-abundant economies grew significantly more slowly than resource-poor ones over the 1970–90 period, even after controlling for initial GDP and openness. Ploeg (2011)'s survey of the subsequent literature argues that the curse is conditional rather than mechanical. Institutional quality, fiscal-policy design and the degree of financial development can flip the sign of the resource effect on growth, so resource wealth is best read as a multiplier of pre-existing institutional conditions rather than as an independent driver. The subnational evidence sharpens the picture. Cust & Poelhekke (2015) show, on a panel of georeferenced oil and gas fields, that resource production delivers spatially concentrated benefits to the immediately surrounding districts rather than broad regional spillovers, and Cust & Mihalyi (2017) document a "presource curse" in which the mere announcement of a major discovery distorts macroeconomic outcomes years before any rents materialise. Berman et al. (2017) combine georeferenced African mining with conflict-event data and identify, off exogenous world-price variation, a robust positive effect of mining activity on local violence, with fighting-group control of mining areas scaling local conflict into country-wide armed conflict.

The energy transition has put critical minerals at the centre of donor–recipient engagement and turned it into a set of explicit strategic agreements. In June 2025, the United States brokered the Washington Accords between the DRC and Rwanda, which commit both parties to develop "transparent, formalised, end-to-end mineral supply chains" with American investors in exchange for (vaguely worded) U.S. regional-security guarantees (Amisi & Nkuba, 2026). Read together with the US-financed Lobito Corridor transport-and-investment compact for DRC and Zambian critical minerals, the agreement "goes far beyond what is usually meant by conventional development cooperation" (Baumgartner, 2026). It aligns supply chains for critical minerals with US industrial priorities and channels US private investment into the Congolese mining sector (Center for Strategic and International Studies, 2024). The European Union is a co-investor in the same Lobito Corridor under its Global Gateway framework (European Commission, 2023b), and has signed bilateral Strategic Partnership memoranda on sustainable raw-material value chains with Namibia (implemented by a 2024–25 roadmap under Global Gateway) (European Commission, 2023a), the DRC and Zambia (European Commission, 2023b) and Rwanda (European Commission, 2024). Africa is the largest regional target of Global Gateway, with 72 of the 138 flagship projects listed for 2024 (Eurodad and Counter Balance and Oxfam, 2024). A review of forty of these flagship projects by Eurodad and Counter Balance and Oxfam (2024) argues that the strategy prioritises EU commercial and geopolitical interests over traditional development objectives, a reading that the EU's own Joint Communication on Global Gateway makes hard to dispute: "In assisting others, the EU will also be contributing to the promotion of its own interests, to strengthening the resilience of its supply chains, and to opening up more trade opportunities for the EU economy" (Eurodad and Counter Balance and Oxfam, 2024). The EU–Rwanda raw-materials agreement in particular is criticised for risking an escalation of the M23 conflict in the DRC's mineral-rich east, where Rwanda is widely implicated The European Parliament has demanded the memorandum's suspension until Rwanda ceases all interference in the DRC, including the export of minerals mined in M23-controlled areas (European Parliament, 2025). The US, EU and Lobito-Corridor agreements above can be read jointly as a Western-donor diversification move on a heavily concentrated upstream. Chinese state-owned enterprises and policy banks are estimated to control around 80% of DRC cobalt output (Schoonover, 2025), and Chinese export restrictions since 2023 on graphite, gallium, germanium, antimony, tungsten and selected heavy rare earths have sharpened Western concerns over supply-chain resilience (International Energy Agency, 2025). The EU has encoded that diversification objective in law via a 65% single-third-country benchmark on strategic raw materials (European Union, 2024). Human-rights exposure in parts of the Chinese upstream is a second, narrower motivation behind US and EU due-diligence rules (Murphy & Elimä, 2021). The EU's own development-cooperation leadership has begun to frame the shift explicitly. The proposed EUR 200 bn "Global Europe" instrument for 2028–2034 is to be administered with a formal "European Preference" baked into project tenders and an explicit link from EU aid to critical-minerals offtake. In the words of Alexei Jones, who leads the EU foreign and development policy department at the European Centre for Development Policy Management (ECDPM) think tank quoted in the same coverage, the EU is transitioning "from a development instrument with geopolitical ambitions to a geopolitical instrument with development safeguards" (Becker, 2026).

Beijing on the other hand has acted to consolidate its supply chain control. In March 2024 the DRC and China renegotiated the 2008 Sicomines minerals-for-infrastructure agreement after EITI-DRC and the Inspectorate-General of Finance documented that the original deal delivered only around USD 822 million of the promised infrastructure against estimated Chinese profits of roughly USD 10 billion; the renegotiation raises the infrastructure commitment from USD 3 bn to USD 7 bn (Extractive Industries Transparency Initiative, 2024). The deal nonetheless retains commodity-price risk, lacks a fresh independent reserves study and, in expert assessments, remains largely tilted toward Chinese interests (Voice of America, 2024).

Arezki & Ploeg (2025) even argue that the contemporary super-power race for critical minerals in developing countries, together with absent governance reforms, displays the same dynamics that produced the classical resource curse, thereby risking a "critical minerals curse". On their reading, navigating the resource boom requires a deliberate balance between outward-facing institutions that govern the relationship with foreign buyers and inward-facing institutions that govern the domestic distribution of rents.

2.3 Conceptual framework: who allocates aid within recipient countries?

The two literatures reviewed above establish that donors respond to shifts in a recipient's strategic and commercial value, and that mineral activity transforms the local economies where it occurs. Neither, by itself, delivers a district-level prediction. Converting the country-level result of Arezki et al. (2024) into an ADM2-level hypothesis requires an account of who decides where aid projects land and what that actor wants. Formally, donors approve project locations, but in practice, project placement emerges from bargaining between donor agencies and the recipient government, which holds an informational advantage over local conditions (Jablonski, 2014). A pre-registered survey of World Bank task team leaders finds that they perceive recipient governments as most interested in targeting aid politically and in controlling project implementation (Briggs, 2021). Subnational aid allocation is therefore best read as jointly produced, and each side of the bargain supplies distinct channels through which the start of mineral exploration could move aid toward, or away from, the producing district.

Where the donor effectively controls location, two channels predict targeting of the exploring district itself. The first is a *commercial-foothold* channel. If aid complements future resource contracts, the donor has an incentive to place projects at the deposit. Transport, energy and water infrastructure makes later extraction and export by its own firms cheaper and less risky, and visible projects signal commitment to the host government during the licensing window that exploration opens. The subnational evidence is consistent with donors behaving this way. The World Bank disproportionately allocates infrastructure projects to districts where foreign investors stand to benefit (Nunnenkamp et al., 2017), and the EU's Global Gateway raw-materials partnerships explicitly bundle development finance with supply-chain objectives (Eurodad and Counter Balance and Oxfam, 2024). The second is a *social-licence* channel. Mining investment incites local opposition. The probability of protest or riot in a mining area roughly doubles once mining begins (Christensen, 2019); Chinese-owned critical-mineral mines raise the likelihood of local protest by around 15 percentage points relative to comparable non-Chinese operations, with negative environmental externalities as the key mechanism (Tsang, 2025); and georeferenced mining activity raises local violence (Berman et al., 2017). A donor whose supply-chain interests depend on uninterrupted extraction has a reason to buy local acceptance with visible social-sector projects in the affected district. Against these stands a third channel, *bargaining*. If aid is a side payment negotiated with the central government in exchange for access, nothing ties the payment's location to the deposit. The aid can land in the capital, along export corridors, or wherever the government directs it, in which case the country-level increment documented by Arezki et al. (2024) would be spatially decoupled from the producing district.

Where the central government effectively controls location, the literature identifies three further channels. The first is *patronage*. Leaders steer aid-financed projects toward their birth and co-ethnic regions (Dreher et al., 2019; Hodler & Raschky, 2014; Öhler & Nunnenkamp, 2014). Directly on point, Perrotta Berlin et al. (2023) show that when a recipient's geostrategic value to donors rises during a rotating UN Security Council seat, the country receives more World Bank projects overall, but the additional projects and loan commitments tilt disproportionately toward the leader's co-ethnic regions. A country-level strategic shock is thus channelled partly into subnational favouritism rather than toward the region that generates the strategic value. The second is an *electoral* channel. In competitive systems, governments use their influence over project placement to reward the voters most likely to keep them in power, in practice their core supporters, translating aid into votes (Anaxagorou et al., 2020; Jablonski, 2014). Under both channels, the exploring district receives the mineral-induced aid increment only where its political value happens to coincide with its geological value. The third, a *social-peace* channel, predicts genuine local targeting by contrast. Exploration and extraction generate grievance and protest in producing areas (Christensen, 2019; Tsang, 2025). A government that values calm, or whose donor partners do, therefore has a reason to direct visible projects to protest-prone producing districts, both to compensate affected populations and to secure the cooperation of local elites with extraction.

These channels should not operate uniformly across political regimes. Regional favouritism weakens where institutions constrain the executive (Hodler & Raschky, 2014), and democratic change disciplines ethnic favouritism in public investment (Burgess et al., 2015). The core-voter targeting of Chinese aid disappears under strong checks and balances (Anaxagorou et al., 2020), and the protest response to Chinese mining is weaker in stronger democracies (Tsang, 2025). In autocracies, the patronage and social-peace channels should therefore dominate the within-country allocation; in democracies, the electoral channel replaces them and multilateral safeguards bind more tightly. No regime type unambiguously predicts local targeting of mineral districts, which is itself a reason not to expect a large pooled effect.

The channels also might differ in which margin of aid they move first. A commitment records the allocation decision, the political act the framework describes; a disbursement records delivery. On average, commitments are largely disbursed within two years, but delivery lags are longest for infrastructure projects and disbursed volumes track commitments only loosely year to year (Bulíř & Hamann, 2008; Hudson, 2013). A strategic response to exploration should therefore appear first and most sharply in commitments, while disbursement responses arrive attenuated and delayed, most severely in the infrastructure-heavy categories that the commercial channels single out. The outcome measures in this thesis capture delivery for the Western donors (disbursements) and intent for China (ODA-like commitments), a distinction Section 3 details and one that the interpretation of the donor-specific results must respect.

In sum, the commercial-foothold, social-licence and social-peace channels predict that aid flows to the exploring district itself; the bargaining, patronage and electoral channels predict that the country-level increment is absorbed elsewhere in the recipient's administrative geography. The hypotheses that follow state the net prediction (H1) and use the donor and sectoral composition of aid (H2, H3) to probe the channels. The regime-type, political-alignment and protest moderators that the framework identifies lie beyond the design of this thesis and are taken up as room for further work in Section 7.

2.4 Hypotheses

2.4.1 Local aid response to mineral exploration (H1)

The framework in Section 2.3 yields a direct subnational prediction and its alternative. The local-targeting channels (commercial foothold, social licence, social peace) imply that part of any aid response to exploration is directed to the exploring district itself. The decoupling channels (country-level bargaining, patronage, electoral targeting) imply that the increment is absorbed elsewhere in the country. The country-level evidence fixes the size of the aggregate response to a major discovery, the strongest version of the mineral shock considered here. Arezki et al. (2024) show that recipient countries receive roughly 36% more bilateral aid following a major mineral discovery, consistent with donors competing for resource access. The treatment observable in this thesis, the start of exploration, is the earliest recorded stage of the same process and the stage at which the commercial-foothold channel predicts strategic donors act, during the licensing window and before access is contested. What the country-level evidence cannot show is whether any increment trickles down to the mineral-bearing area, and the within-country design of this thesis isolates exactly that question. Because treated districts are compared to never-treated districts in the same country, a positive average treatment effect on the treated (ATT) requires that donors or governments route resources to the exploring district specifically. H1 posits that the local-targeting channels dominate.

> **H1:** The start of mineral exploration in an African ADM2 unit causes a systematic *increase* in foreign aid received by that unit, all else equal.

H1 is tested as the headline ATT on total aid in Section 5, on the PSM-Matched ADM2 sample with a Mining-Restricted bias check and a supporting cell-level analysis reported in the Annex.

2.4.2 Cross-donor heterogeneity (H2)

The donor-as-strategic-actor framing of Alesina & Dollar (2000) and Kuziemko & Werker (2006) does not predict that all donors react identically to a given subnational shock. Donors differ in commercial exposure to the recipient's mineral sector, in geopolitical objectives, and in the institutional design of the aid programme, so a differential response across donors is a natural first conjecture. The framework in Section 2.3 gives it sharper content, since the commercial-foothold channel should bind most tightly for donors with direct mineral supply-chain exposure. Two strands of evidence push against that conjecture for the donor aggregates observable here. First, the World Bank's stated needs-based mandate makes it more selective on governance (Dollar & Levin, 2006; Morrison, 2011), but the subnational evidence on whether multilateral delivery differs from bilateral delivery is mixed. Briggs (2017) finds that multilateral aid within African recipients disproportionately reaches the regions where the richest quintiles live, and Dreher et al. (2019) show that recipient-leader birth-region targeting is detectable in Chinese but not in World Bank project allocation. Second, and decisive for the present design, the Chinese-aid evidence contradicts the popular portrayal of Chinese aid as "rogue aid" allocated by self-interest alone. Dreher & Fuchs (2015) test this charge and find that, while political considerations do shape Chinese allocation, China does not pay substantially more attention to politics than Western donors, and its aid allocation is largely independent of recipients' natural-resource endowments and institutional characteristics. Dreher et al. (2018) sharpen the picture by disaggregating

Chinese state financing into Official Development Assistance (ODA) and Other Official Flows (OOF). Resource access shapes the OOF component (non-concessional loans, export credits), while Chinese ODA is allocated primarily on foreign-policy criteria such as UN voting alignment and one-China recognition, and behaves much like Western aid. Since the Chinese outcome observable in this thesis is the ODA-like component (Section 3), the Dreher et al. (2018) result implies that none of the three observable donor aggregates should respond differentially to a subnational mineral shock:

> **H2:** Mineral exploration produces no systematically different aid response across the three donor aggregates (OECD-DAC bilateral, World Bank multilateral, and Chinese ODA).

H2 is tested across the three donor aggregates in Section 5.2, with the Chinese sub-analysis restricted to the post-2000 panel where GODAD records Chinese projects.

2.4.3 Aid-category heterogeneity (H3)

A third strand argues that, within an aid stream, the *sectoral composition* matters. The OECD-DAC classification distinguishes social-infrastructure aid (health, education, water, sanitation), economic-infrastructure aid (transport, energy, communications, banking), and production-sector aid (agriculture, industry, mining, trade). Donor motives differ systematically across these sectors. Dreher et al. (2024) synthesise the cross-donor evidence and conclude that, where donors have specific motives, they adjust the sectoral composition of aid accordingly. Commercial and export-competition incentives push allocations towards economic-infrastructure and production aid, while social-sector aid is more closely tied to recipient need and to governance considerations. The subnational evidence for a multilateral donor points in the same direction. Nunnenkamp et al. (2017)'s district-level analysis of World Bank projects in India finds that needs-based allocation emerges only for selected sectors and that the Bank disproportionately allocates infrastructure projects to districts where foreign direct investors stand to benefit. Bermeo (2017) completes the picture by arguing that donors further adjust the sectoral mix in response to recipient governance, so that humanitarian-need targeting and strategic targeting can co-exist within a country's overall aid, separated by sector. The Chinese-aid evidence shows the same pattern even more sharply. Guillon & Mathonnat (2020) show that the resource-acquisition signal in Chinese ODA loads on economic-infrastructure-and-services aid, with a weaker association in social-sector aid, rather than on production-sector aid. Taken together, the evidence locates strategic targeting somewhere in the commercially exposed categories, with the exact sectoral vehicle differing by donor and instrument.

> **H3:** The effect of mineral exploration on aid is heterogeneous across aid categories. Sectoral channels with closer commercial or resource exposure (production-sector aid and economic-infrastructure aid in particular) should respond more strongly than humanitarian-targeted channels (social-sector aid).

H3 is tested by running the same staggered-DiD design on total social-sector, total economic-sector, and total production-sector aid.

3 Data

3.1 Sources

I combine four geo-coded sources at annual frequency over the 1990–2020 window, plus conflict-event data used only in a descriptive sub-analysis:

- **Foreign aid (outcome).** Geocoded Official Development Assistance Dataset, version 1.0 [GODAD; Bomprezzi et al. (2025)]. GODAD provides geo-localised aid-project information for 18 European donors and the United States (1973–2020), based on the OECD Creditor Reporting System and Aid Information Management Systems data, together with World Bank (1995–2020), Chinese (2000–2021) and Indian (2007–2014) project records. The outcome measure is **disbursements**. GODAD's ADM2 release carries disbursement columns for all 19 Western donors and the World Bank, and disbursement coverage at the project level is substantially more complete than commitment coverage for these donors. Dollar-value disbursements are only systematically populated from 1990 onwards, which is the binding constraint that anchors the panel at 1990; coverage in the early 1990s is nonetheless thin, a point taken up in the Limitations. The ADM2 panel sums the donor-by-ADM2-year disbursement columns into six composites: total aid (19-donor DAC plus World Bank), total DAC, total World Bank, and the three sectoral composites (social, economic, production) built from the corresponding sectoral disbursement splits. Negative recorded disbursement values (accounting corrections and repayments) are truncated at zero, since the Poisson estimator requires a non-negative outcome and a negative entry records a repayment rather than an allocation decision; ADM2-years without any recorded project enter as structural zeros. Two structural coverage breaks matter for the design. World Bank disbursements begin in 1995, so all World-Bank-specific estimates are run on the 1995–2020 window, and the 1990–94 pre-treatment aid covariates are DAC-based by construction. Chinese development finance is recorded from 2000 and only as **commitments** (GODAD provides no disbursement values here); following Dreher et al. (2018), I use only the ODA-like component of Chinese commitments as the Chinese outcome, since the OOF-like component (non-concessional, resource-linked lending, which accounts for more than 80% of recorded Chinese commitment value in Africa) is a different economic object from the ODA flows of the Western donors. The Chinese sub-analysis is therefore a post-2000, ODA-commitments analysis and is labelled as such (Section 5.2). The cell-level appendix panel is built separately from the GODAD project-level release. Each project location is assigned to a PRIO-GRID 0.5° cell by point-in-polygon spatial join, filtered to precision codes 1 and 2 (exact coordinates and within 25 km) so the cell assignment is meaningful at the grid resolution; where a project's location-split disbursement is unavailable, the project total is divided evenly across its locations rather than assigned in full to each.
- **Mineral exploration (treatment).** USGS *Compilation of Geospatial Data (GIS) for the Mineral Industries and Related Infrastructure of Africa* (Padilla et al., 2021). The geodatabase records the georeferenced locations of mineral commodity production and processing facilities, mineral exploration and development sites, and mineral commodity exporting ports in Africa. For each ADM2 unit I record the year of the first reported mineral exploration or development event in that unit; this defines the treatment cohort variable

$G_i \in \{1995, 1996, \ldots, 2018, 0\}$ ($G_i = 0$ for never-treated units within the analytic sample).[1] The USGS geodatabase is partially taken from foreign governmental and open-source sources and therefore is not a systematic census of exploration activity, so coverage is likely denser for larger operations, for countries with stronger geological-survey infrastructure, and for more recent periods.

- **Mining polygons (sample restriction).** Global mining land-use polygons from Maus et al. (2022), distributed via PANGAEA at https://doi.pangaea.de/10.1594/PANGAEA.942325. The dataset contains 44,929 polygon features covering 101,583 km^2 of large-scale and artisanal mining. Because the polygons are derived from satellite imagery taken around 2019, they are measured *after* most treatment dates; they are therefore used only to define the Mining-Restricted bias-check sample (Section 3.2) and are deliberately excluded from the propensity-score model of the headline matched sample, where conditioning on a post-treatment-measured variable could introduce selection on the outcome of exploration itself.
- **Pre-treatment covariates and cell structure.** PRIO-GRID v.2.0 (Tollefsen et al., 2012). PRIO-GRID supplies the $0.5° \times 0.5°$ cell layer used in the cell-level appendix analysis as well as harmonised population, gross cell product, distance-to-capital, nighttime-lights, and urban-share variables. Pre-treatment covariates are constructed as 1990–94 unit averages, which precede the earliest treated cohort (G = 1995). Cell-level PRIO variables are mapped onto the GADM ADM2 polygons by centroid containment for intensive quantities and by area-weighted intersection for extensive quantities (population, gross cell product), with a nearest-cell fallback for polygons too small to contain any centroid. This addresses the change-of-support problem of mapping gridded covariates onto administrative units (Gotway & Young, 2002). The grid and the ADM2 polygons partition space differently, so a polygon may cover several cells or only a fraction of one, and cell values must be transferred between the two supports without inventing or losing mass. Assigning intensities by cell centroid and dividing totals according to overlap area achieves this, following standard practice in the subnational-development literature (Donaldson & Storeygard, 2016; Michalopoulos & Papaioannou, 2014).
- **Conflict-event data (descriptive sub-analysis only).** Armed Conflict Location & Event Data Project (ACLED) (Raleigh et al., 2010). ACLED records georeferenced political-violence events with date, latitude, longitude, event type and fatalities for African countries from 1997 onwards. I retain Battles, Violence against civilians and Explosions/Remote violence events (excluding Protests and Riots), point-in-polygon-join them to ADM2 polygons and aggregate to ADM2-year counts and fatalities. ACLED is used only in the descriptive PPML sub-analysis in Section 5.3; the 1997 coverage start, which post-dates the earliest treated cohort (G = 1995), rules out a clean pre-treatment conflict baseline for the headline staggered-DiD.

Assigning point and polygon layers to ADM2 units. Three input layers are mapped onto the GADM ADM2 polygons, each posing a different overlap question. The USGS exploration

[1] Throughout the thesis, "exploration" is shorthand for the first recorded activity at a mineral exploration or development site: for 76 percent of the treated districts in the headline sample, the treatment-defining first site carries exploration status (71 percent of sites in the USGS layer overall), and the remainder enter through sites that had advanced to later stages, mostly feasibility studies and reserves development, by the time the USGS compiled the data.

records are points, so each is assigned to the unique ADM2 polygon that contains it, and a unit's treatment cohort is the earliest exploration year among its interior points. The Maus et al. (2022) mining footprints are areal and can straddle a boundary; rather than split a straddling footprint between units, I flag every ADM2 polygon it intersects as mining-feasible. This any-touch rule is deliberately conservative, retaining a unit whenever mining is physically present in any part of it, which suits its only role in the design as a binary feasibility flag for the Mining-Restricted control set. The PRIO-GRID covariate cells, at 0.5°, are larger than many ADM2 units, so intensive cell variables are assigned by centroid containment and extensive variables by area-weighted intersection as described above.

3.2 Sample construction

The full continental panel covers 50 African countries and territories and 6,279 ADM2 polygons.[2] I work with three samples:

1. **Total Sample**: the full continental panel before any sample restriction. Used both for the descriptive balance check in Section 4 that motivates the choice of analytic sample and as the basis for the alternative-explanations PPML in Section 5.3.
2. **Mining-Restricted sample**: the panel is restricted to polygons that overlap a Maus et al. (2022) mining land-use polygon. Substantively, these are districts of demonstrated mineral feasibility, places where the underlying geology has already supported active extraction, so I read them as a proxy for geological potential and mineral wealth rather than for mining activity itself. Restricting the comparison to this set means treated districts are evaluated only against other mineral-endowed districts, which strips out the most basic confounder of the continental comparison, geology. Because the footprints are measured around 2019, conditioning on them is conditioning on a mostly post-treatment variable; the sample is therefore used only as a *bias check* against the headline, not as the headline itself.
3. **PSM-Matched sample**: each treated polygon ($G_i > 0$) is matched 1:1 without replacement to its nearest never-treated control **within the same country** (exact matching on the country code, nearest-neighbour logit propensity score within country). The propensity score is estimated on pre-treatment aid activity (the mean of total disbursements and the number of recorded aid projects) and on geography and economic development (population, gross cell product, distance to capital, urban share, nighttime lights), all constructed from the 1990–94 pre-treatment window (Caliendo & Kopeinig, 2008; Stuart, 2010). Missing values in the geography and development covariates are filled before matching, with zero where zero is a valid level of the variable (nighttime lights, urban share) and with the sample median for population, gross cell product and distance to capital, where a zero would not be a meaningful level; the balance tables in Section 4 evaluate the covariates as they enter the matching. The within-country restriction aligns the matched comparison with the research question (treated versus comparable units *in the same country*) and absorbs country-level differences in total aid volumes and aid cycles that unit and year fixed effects alone do not remove; treated units in countries without any never-treated candidate

[2] The panel covers 48 of the continent's 54 sovereign states plus the territories of La Réunion and Saint Helena; Cabo Verde, the Comoros, Lesotho, Libya, Mauritius, the Seychelles and Western Sahara are not covered by the assembled GODAD-based panel. Exact within-country matching further restricts the PSM-Matched sample to 36 countries that contain both treated units and never-treated candidate controls.

are dropped, and their number is reported in the sample table. **The PSM-Matched sample is the main analytic sample for the headline ATT.**

3.3 Treatment map

Figure 1: Treatment status by ADM2 polygon, PSM-Matched sample

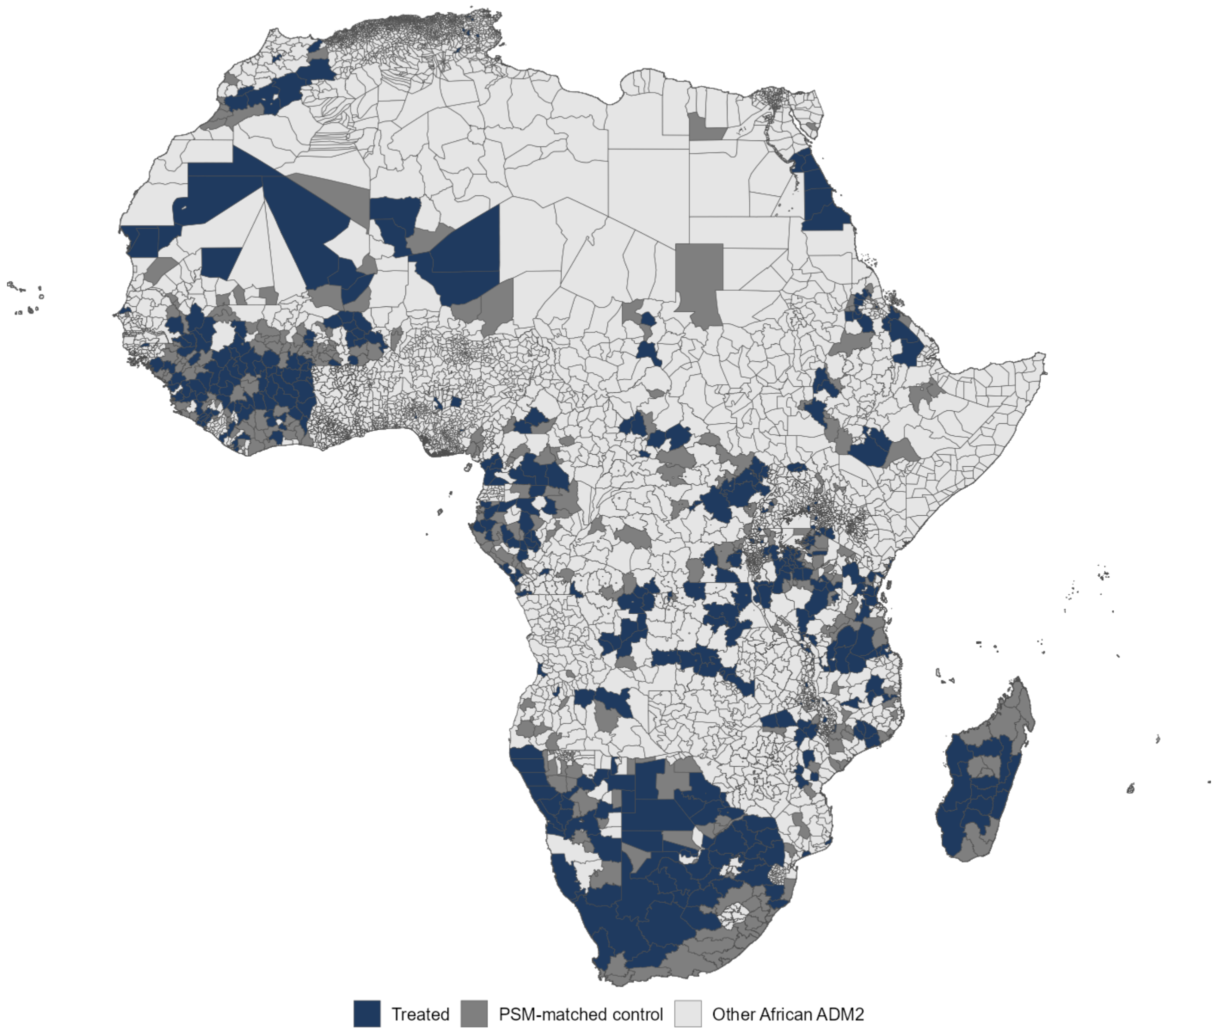

Navy: treated polygons, where mineral exploration or development began at some point during the 1995 to 2018 treatment window. Grey: PSM-matched never-treated controls (same-country matches). Light grey: all other African ADM2 polygons in the continental cross-section, neither treated nor matched.

3.4 Treatment rollout and cohort timing

Figure 2 shows the distribution of treatment-cohort years for the ADM2 PSM-Matched sample. The treatment cohorts span 1995 to 2018, with cohort sizes peaking in the mid-2000s.

Figure 2: Treatment cohort sizes, ADM2 PSM-Matched sample

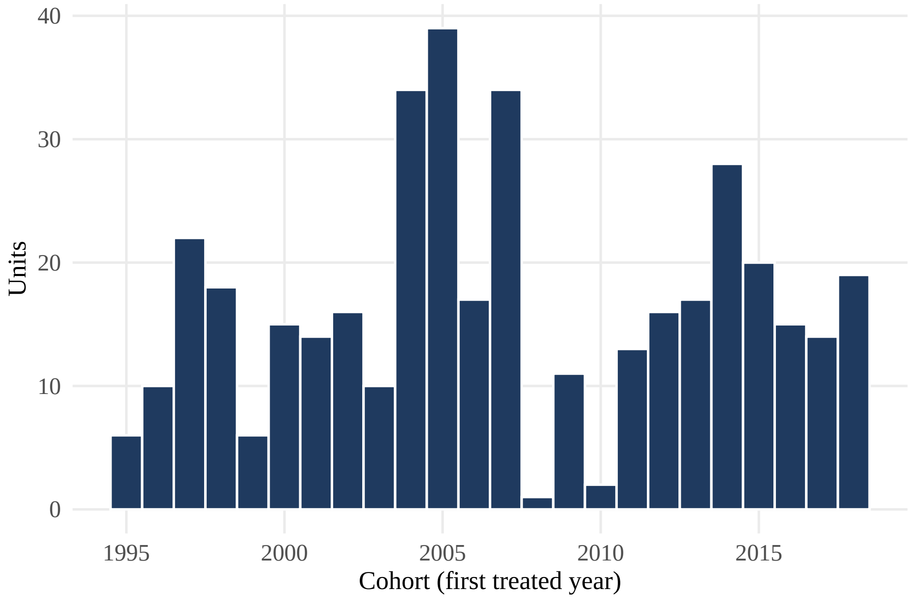

Each bar counts the matched treated ADM2 polygons whose first recorded mineral exploration falls in the given year.

3.5 Outcome trajectories by cohort group

Figure 3 traces the mean inverse-hyperbolic-sine of total aid over the panel window for each cohort group in the PSM-Matched sample. The visual comparison of pre-treatment trajectories complements the formal event-study evidence in Section 5.

Figure 3: Mean asinh(Total Aid) over time by cohort group, ADM2 PSM-Matched sample

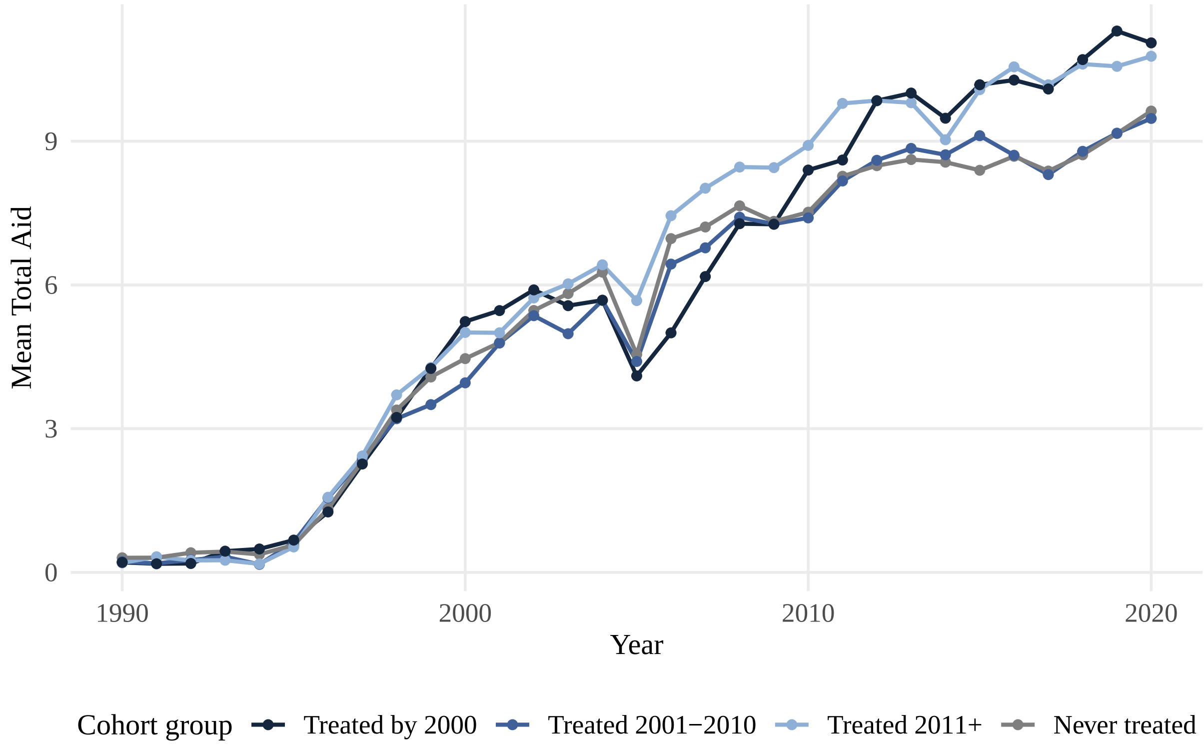

Total aid is the 19-donor DAC plus World Bank disbursement composite; the level shift around 1995 reflects the start of World Bank disbursement coverage and is common to all groups (absorbed by year fixed effects in the estimators).

3.6 Summary statistics

Table 1: Sample sizes by ADM2 sample. Treated counts the unique ADM2 polygons with $G_i > 0$ (any cohort year in the panel). Exact within-country matching drops the 15 treated polygons whose country contains no never-treated candidate control.

Sample	ADM2 polygons	ADM2 treated
Total Sample	6279	412
Mining-Restricted	468	238
PSM-Matched (headline)	794	397

4 Methodology

4.1 Identification setup

Let i index ADM2 polygons and t index years. The treatment timing $G_i \in \{0\} \cup \{g_{\min}, \ldots, g_{\max}\}$ records the year of first exploration ($G_i = 0$ for never-treated). Define $D_{it} = \mathbf{1}\{G_i > 0 \text{ and } t \geq G_i\}$ as the post-treatment indicator. The estimand of interest is the average treatment effect on the treated (ATT), aggregated across cohorts and post-treatment years.

Identification rests on the two standard difference-in-differences assumptions. The first is conditional parallel trends. Absent exploration, treated and never-treated districts with the same pre-treatment covariates would have followed the same average aid path. It is supported by four pieces of evidence, reported in Section 4, Section 5 and the Annex: within-country PSM balance on the pre-treatment covariates, the pre-treatment coefficients of a heterogeneity-robust event study, HonestDiD sensitivity bounds that quantify how large a violation would have to be to overturn the conclusions, and an Augmented Synthetic Control robustness check whose validity rests on reproducing pre-treatment outcome paths rather than on parallel trends. The second assumption is no anticipation. Aid does not respond to exploration before it begins. The headline specification sets the anticipation window to zero, and a sensitivity grid over $\{0, 1, 2\}$ years of anticipation is reported in Annex Table 19 and discussed in Section 5.1.5.

4.2 Estimators

The standard two-way fixed-effects (TWFE) regression is biased when treatment is staggered and treatment effects differ across cohorts. Already-treated units enter the control set for later cohorts, and the TWFE coefficient becomes a weighted average of cohort-specific ATTs in which some weights can be negative (Borusyak et al., 2024; Goodman-Bacon, 2021); the same contamination affects its event-study leads and lags (Sun & Abraham, 2021). I therefore use two heterogeneity-robust estimators (a primary and a secondary specification), a heterogeneity-robust event study for the dynamic evidence, and two complementary checks (an Augmented Synthetic Control robustness check and HonestDiD sensitivity bounds).

Primary: Poisson extended two-way fixed effects (Wooldridge, 2023). I model the conditional mean of raw aid as

$$\mathbb{E}[Y_{it} \mid G_i = g, t] \;=\; \exp(\lambda_g + \tau_t + \delta_{g,t}\, D_{it}),$$

where λ_g are cohort fixed effects (one effect per treatment cohort rather than one per district; Wooldridge (2023) shows this is sufficient to recover the ATTs in this design), τ_t are year effects, and $\delta_{g,t}$ is the treatment effect for cohort g in year t, which the post-treatment indicator D_{it} switches on only once treatment has started. The cohort-by-year effects are aggregated to a single ATT, interpretable as a semi-elasticity (log-percentage change in raw aid). The Poisson link suits this outcome for two reasons. First, well over half of the ADM2-year observations record zero aid (the % Zero column of Annex Table 28), which rules out OLS on log-transformed aid without arbitrary, unit-dependent adjustments. Second, Poisson quasi-maximum-likelihood is consistent for the conditional mean even when the conditional variance is misspecified (Wooldridge, 1999), and it requires parallel trends only in the ratio of conditional means, a weaker condition than the additive parallel trends of log-linearised OLS. The covariates from the matching stage do not re-enter the estimating equation; conditioning on X_i operates through the construction of the matched sample. Standard errors are cluster-robust at the province level throughout, one level above the unit of observation, to absorb within-province correlation in aid disbursements and in the timing of mineral exploration across districts that share administrative infrastructure.

Event study (ETWFE with never-treated controls). The R implementation of the ETWFE estimator includes treatment-effect terms only for post-treatment cohort-by-year cells; the pre-treatment cells form the estimation baseline, so the headline regression produces no lead coefficients to display. Setting the control group to the never-treated units resolves this, because the never-treated units then identify the time effects in every period and the pre-treatment cells become estimable. The event study is therefore the event-time aggregation of the same Poisson regression with never-treated controls, with the cohort's final pre-treatment year ($t = -1$) as the reference period. Because each event-time coefficient is a cohort-share weighted average of the cohort-by-year effects, the leads and lags are free of the contamination that affects a naive dynamic TWFE event study under heterogeneous effects (Sun & Abraham, 2021), so the pre-treatment coefficients are a clean diagnostic for parallel trends. Event-time effects are reported on the link scale (comparable with the ATT semi-elasticities) with delta-method standard errors from the province-clustered covariance matrix, and the display window spans event times -5 to $+10$.

Secondary: Callaway & Sant'Anna (2021) estimator on the asinh-transformed outcome. The asinh transformation behaves like the logarithm for large values but is defined at zero, which suits an outcome where most district-years record no aid (Bellemare & Wichman, 2020). The Callaway & Sant'Anna (2021) framework constructs $\text{ATT}(g, t)$ cell-by-cell via a doubly-robust combination of regression adjustment and propensity weighting; I aggregate to a simple overall ATT. The specification conditions on three time-invariant pre-treatment covariates (log population, log GCP PPP, log distance to capital), with multiplier-bootstrap standard errors clustered at the province level. The control group is the not-yet-treated set, which is far larger than the never-treated set on this panel. The China sub-analysis (Section 5.2) reverts to never-treated controls because its latest cohorts have too few not-yet-treated comparison units. Because asinh-scale magnitudes are sensitive to the units of the outcome, in the same way that OLS on $\log(1 + y)$ is (Chen & Roth, 2024), the Poisson ATT carries the headline magnitude and this estimator adds evidence on the sign of the effect.

The Callaway–Sant'Anna framework also provides a *joint pre-trends Wald test*. Under conditional parallel trends the statistic is asymptotically chi-squared, so a small p-value flags evidence against pre-treatment parallel trends. On this panel the test should be read with care. With thousands of pre-treatment cohort-time cells entering jointly it rejects for economically small deviations, and the computed statistics are often inflated by a near-singular pre-period covariance matrix (specifications where the statistic is not computable at all are marked in the tables). A rejection is therefore treated as a flag pointing to the event-study and HonestDiD evidence, not as a verdict by itself.

Robustness: Augmented Synthetic Control Method (Ben-Michael et al., 2021). ASCM constructs, for each treated unit (or, where that problem is too large or ill-conditioned, for each treatment cohort, that is, all treated units sharing the same first-exploration year, fitted jointly with a common synthetic control), a weighted combination of never-treated units that reproduces the treated unit's own pre-treatment outcome trajectory, with ridge regularisation where the fit is otherwise unstable. ASCM replaces the parallel-trends assumption with a different requirement, namely that the weighted combination reproduces the treated units' pre-treatment aid path. Its validity therefore rests on the quality of that pre-treatment fit rather than on the assumption under scrutiny in Section 5.1.5, which makes it a useful cross-check where the pre-trend evidence is imperfect. Like the secondary estimator, ASCM is fitted to the asinh-transformed aid outcome, so the reported ATTs are on the asinh scale even though the outcome names in the tables are shown untransformed. ASCM ATTs for the Mining-Restricted and PSM-Matched samples are reported in Annex Tables 25 and 26.

Sensitivity: HonestDiD (Rambachan & Roth, 2023). HonestDiD provides bounds on the ATT under bounded violations of parallel trends, governed by the relative-magnitude bound $\bar{M}$ (treated-versus-control deviation) and the smoothness bound M. Its based on the dynamic Callaway–Sant'Anna event-study coefficients and their influence-function covariance matrix, following the implementation recommended by the HonestDiD authors for staggered designs, since HonestDiD restricts additive deviations from parallel trends on the scale at which the estimator is defined, which is what the asinh Callaway–Sant'Anna specification provides. The Poisson estimator instead identifies effects from parallel trends in the ratio of conditional means, where the same bounds would not carry the same meaning. I report the HonestDiD plots for the headline analysis (PSM-Matched ADM2, total aid) in Section 5.1.4.

4.3 Estimation windows

The donor-specific coverage breaks in GODAD give each analysis its own estimation window. Table 2 collects the estimation windows, treatment-cohort ranges and covariate vintages in one place; the donor-specific choices are motivated where the respective results are introduced (Section 5.2).

Table 2: Estimation windows and timing conventions by analysis

Analysis	Sample	Window	Cohorts	Covariates
Headline	PSM-Matched ADM2	1990–2020	1995–2018	1990–94
World Bank	Headline matched pairs	1995–2020	1995–2018	1990–94
Chinese ODA	China-adapted PSM	2000–2020	2000–2018	1990–99 / 1990–94
Descriptive PPML	PSM, Total, Mining	1995–2005	–	contemporaneous

Note: Headline covers both the ETWFE Poisson and the Callaway–Sant'Anna estimator; the World Bank and Chinese rows apply the same estimator pair on their restricted windows. The covariate window precedes the earliest treated cohort (G = 1995); the Chinese entry reads matching window / doubly-robust adjustment window. The World Bank rows reuse the headline matched pairs on the truncated window; the six units of the 1995 cohort have no pre-treatment year inside that window and are dropped automatically by the Callaway–Sant'Anna estimator, while the Poisson estimator retains them with their first treated year absorbed as the cohort baseline; the same applies to the 2000 cohort of the China-adapted panel. Event studies display event times -5 to $+10$ with $t = -1$ as the reference period. The PPML window is the effective window with measured PRIO-GRID covariates (contemporaneous regressors); its ACLED variant runs on the 1997+ truncation of each sample, effective 2000–05.

4.4 Sample comparability and choice of analytic sample

The first step in defending the parallel-trends assumption is to demonstrate that, after matching, treated and control polygons are comparable on the dimensions that plausibly drive the outcome. Table 3 reports balance on the *Total Sample* (the full continental panel of African ADM2 polygons, treated $G > 0$ versus all $G = 0$). Table 4 reports balance for the PSM-Matched sample. The Mining-Restricted comparison is reported as Annex Table 14. Throughout, the *normalised difference* in column 4 is defined as $(\bar{X}_T - \bar{X}_C)/\sqrt{(s_T^2 + s_C^2)/2}$, where $\bar{X}_T$ and $\bar{X}_C$ are the treated and control sample means and s_T^2 and s_C^2 the corresponding sample variances. Following the practice of Imbens & Wooldridge (2009), an absolute value above 0.25 is read as flagging substantial imbalance. Because World Bank disbursements begin in 1995, the 1990–94 pre-treatment aid rows measure DAC disbursements; the pre-window aid measures are also very sparse (most units record no disbursement in 1990–94), so the matching is in practice anchored by the geography and development covariates.

Table 3: Covariate balance, Total Sample (all African ADM2)

Variable	T mean	C mean	Norm. diff.	p-value	Group
Pre-treatment Total Aid	8.49e+04	3.61e+04	0.043	0.457	Pre-treatment Aid
Pre-treatment Social Aid	4.11e+03	4.97e+03	-0.006	0.789	
Pre-treatment Economic Aid	8.06e+04	3.02e+04	0.047	0.440	
Pre-treatment Production Aid	58.3	232	-0.043	0.065	
Donors Active (1990-94)	0.0233	0.0117	0.102	0.072	
P(Aid > 0)	0.0218	0.0107	0.113	0.049	
Nighttime Lights (1990-94)	0.0222	0.0653	-0.486	0.000	Geography & Economic Development
Population (1990-94)	506	619	-0.114	0.002	
GCP PPP (1990-94)	0.619	1	-0.190	0.001	
Distance to Capital (km)	494	348	0.409	0.000	
Urban Intensity (1990-94)	0.155	0.657	-0.373	0.000	

Note: N treated = 412, N control = 5867. Pre-treatment covariates and pre-aid means are computed over 1990–1994, the panel's first five years, prior to the earliest treated cohort (G = 1995); the same window feeds the PSM matching and the Callaway–Sant'Anna conditional specification. Aid rows are disbursements; the pre-window contains DAC flows only because World Bank disbursement coverage starts in 1995.

Table 4: Covariate balance, PSM-Matched sample (ADM2)

Variable	T mean	C mean	Norm. diff.	p-value	Group
Pre-treatment Total Aid	8.65e+04	2.43e+04	0.065	0.357	Pre-treatment Aid
Pre-treatment Social Aid	2.58e+03	7.69e+03	-0.097	0.173	
Pre-treatment Economic Aid	8.37e+04	1.63e+04	0.071	0.317	
Pre-treatment Production Aid	60.5	0	0.071	0.318	
Donors Active (1990-94)	0.0196	0.0287	-0.070	0.326	
P(Aid > 0)	0.0181	0.0272	-0.078	0.270	
Nighttime Lights (1990-94)	0.0169	0.0196	-0.110	0.120	Geography & Economic Development
Population (1990-94)	500	406	0.142	0.046	
GCP PPP (1990-94)	0.535	0.423	0.063	0.376	
Distance to Capital (km)	508	508	0.001	0.988	
Urban Intensity (1990-94)	0.0931	0.15	-0.132	0.063	

Note: Exact within-country matching. N treated = 397, N control = 397.

The Total Sample shows substantial imbalance on the geography and political-economy covariates. Treated polygons are on average farther from the capital (normalised difference 0.41), less urbanised (-0.37), and have lower pre-treatment nighttime lights (-0.49) than the never-treated set, consistent with mineral exploration occurring disproportionately in the more peripheral, less-developed parts of recipient countries. Several normalised differences approach or exceed the 0.25 threshold of Imbens & Wooldridge (2009); treated and control sets are not directly comparable on these dimensions, so the Total Sample is unsuitable for the headline ATT. The Mining-Restricted sample limits the control set to ADM2 polygons that overlap a Maus et al. (2022) mining land-use polygon, ensuring that the never-treated units are at least geologically plausible counterfactuals; its balance is reported in Annex Table 14. Its control-group composition is, however, determined by a footprint measured around 2019, i.e. partly after treatment, which is why it serves only as a bias check.

The PSM-Matched sample is built by estimating a logit propensity score for treatment assignment on the pre-treatment covariates and selecting nearest-neighbour matched controls within the same country (exact matching on the country code). I follow the standard implementation guidance of Caliendo & Kopeinig (2008) (logit propensity score, nearest-neighbour matching, post-match balance diagnostics). The within-country design means every treated-control contrast is taken inside a single country's aid allocation. Table 4 shows the resulting balance; the largest absolute normalised difference in the matched sample is 0.14, and capital distance falls from 0.41 in the Total Sample to 0.00 after matching. The covariate-overlap density curves in Annex Figures 6 and 7 complement the table. I therefore adopt the PSM-Matched sample as the main analytic sample because it provides the tightest defence of the conditional parallel-trends assumption, which requires that treated and control units be comparable on the dimensions that drive the outcome. The Mining-Restricted sample is retained as a bias check. If the null replicates there despite the different control-selection logic, it is unlikely to be an artefact of the matching. Its results are reported in the Annex.

5 Empirical Findings

5.1 Headline analysis

5.1.1 ETWFE Poisson event study

The pre-treatment coefficients of the headline event study are the cleanest non-parametric test of the conditional parallel-trends assumption. Figure 4 plots the ETWFE Poisson event study with never-treated controls (which makes the leads estimable, see Section 4) on the PSM-Matched ADM2 panel, with the omitted reference period at $t = -1$ and standard errors clustered at the province level. Among the four pre-treatment leads, 3 are statistically distinguishable from zero at the 5% level, and among the eleven post-treatment lags, 9. The interpretation of this pattern, and its sensitivity under bounded violations of parallel trends, is taken up in Sections 5.1.4 and 5.1.5.

Figure 4: ETWFE Poisson event study: ADM2 PSM-Matched, Total Aid

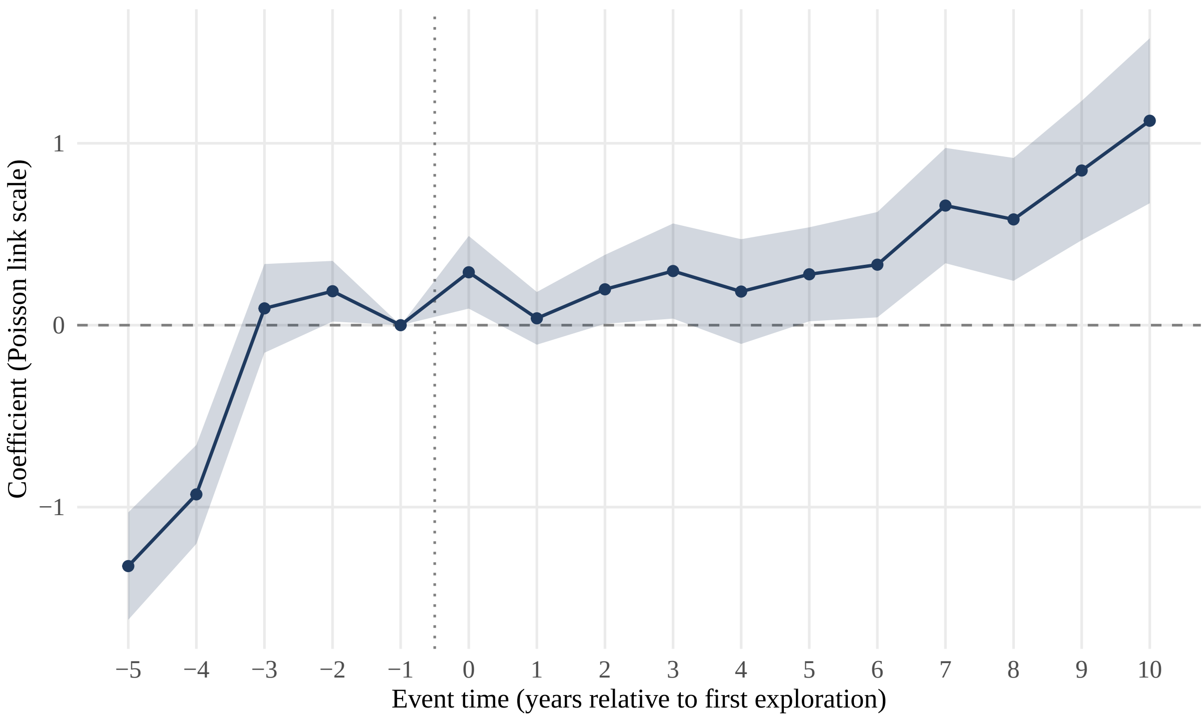

Event-time aggregation of the Wooldridge ETWFE cohort-by-time effects with never-treated controls; coefficients estimate the year-by-year ATT on the link scale relative to the year before treatment ($t = -1$, normalised to zero); shaded band: 95% pointwise confidence intervals, cluster-robust at the province level.

5.1.2 Headline ATT: ETWFE Poisson

Table 5: Primary headline ATT on total aid (ETWFE Poisson, ADM2 PSM-Matched)

Level	Outcome	ATT	SE	p-value
Adm2	Total Aid	0.0192	(0.0274)	0.483

Note: Coefficients are link-scale semi-elasticities (log-percentage changes in raw aid); cluster-robust SEs at province level. *p<0.1, **p<0.05, ***p<0.01.

Table 5 shows the headline ETWFE Poisson ATT on total aid. The point estimate is 0.019 (cluster-robust SE 0.027), statistically indistinguishable from zero, so the null of no aid response to the start of mineral exploration cannot be rejected on the headline sample. The 95% confidence interval spans -0.035 to 0.073 on the link scale, i.e. roughly -3% to +8% of aid, far inside even the 36% country-level benchmark that Arezki et al. (2024) estimate for major discoveries, a stronger shock than exploration onset. The donor-specific and sector-specific composites are reported separately in Tables 8 and 10 (Section 5.2) The headline null does not extend uniformly to them, and the heterogeneity beneath the composite is the subject of that section.

5.1.3 Callaway and Sant'Anna doubly-robust check

Table 6: Secondary headline ATT on total aid (Callaway–Sant'Anna conditional, ADM2 PSM-Matched)

Level	Outcome	ATT	SE	p-value	PT Wald p
Adm2	Total Aid	-0.1212	(0.3252)	0.709	<0.001

Note: Doubly-robust estimates on the asinh of aid. *p<0.1, **p<0.05, ***p<0.01. PT Wald p: joint test of the pre-treatment ATT(g,t)

Table 6 shows the Callaway–Sant'Anna conditional estimate of -0.121 (SE 0.325) on the asinh scale, agreeing with the ETWFE verdict of no detectable response. The joint pre-trends Wald test rejects at the 5% level, so the formal pre-test flags evidence against conditional parallel trends; the HonestDiD analysis below quantifies how much the conclusions depend on this.

5.1.4 HonestDiD sensitivity

The HonestDiD framework of Rambachan & Roth (2023) formalises sensitivity to bounded violations of parallel trends. I apply it to the dynamic Callaway–Sant'Anna event-study coefficients with their influence-function covariance matrix, following the implementation the HonestDiD authors recommend for staggered designs. Figure 5 displays both the relative-magnitudes restriction (which bounds the post-period violation by a multiple of the largest observed pre-period violation) and the smoothness restriction (which bounds the second derivative of pre-period trends).

For a null headline, the relevant robustness question is whether the null depends on assuming exact parallel trends. Figure 5 shows that it does not. Under both restrictions every robust confidence set continues to include zero, up to relative-magnitude bounds of $\bar{M} = 2$ (twice the largest observed pre-period deviation). The sets widen as the allowance grows, which prices in the imperfect pre-period evidence documented in Section 5.1.5, but the conclusion does not change at any considered bound.

Figure 5: HonestDiD sensitivity analysis for the headline (ADM2 PSM-Matched, asinh of Total Aid)

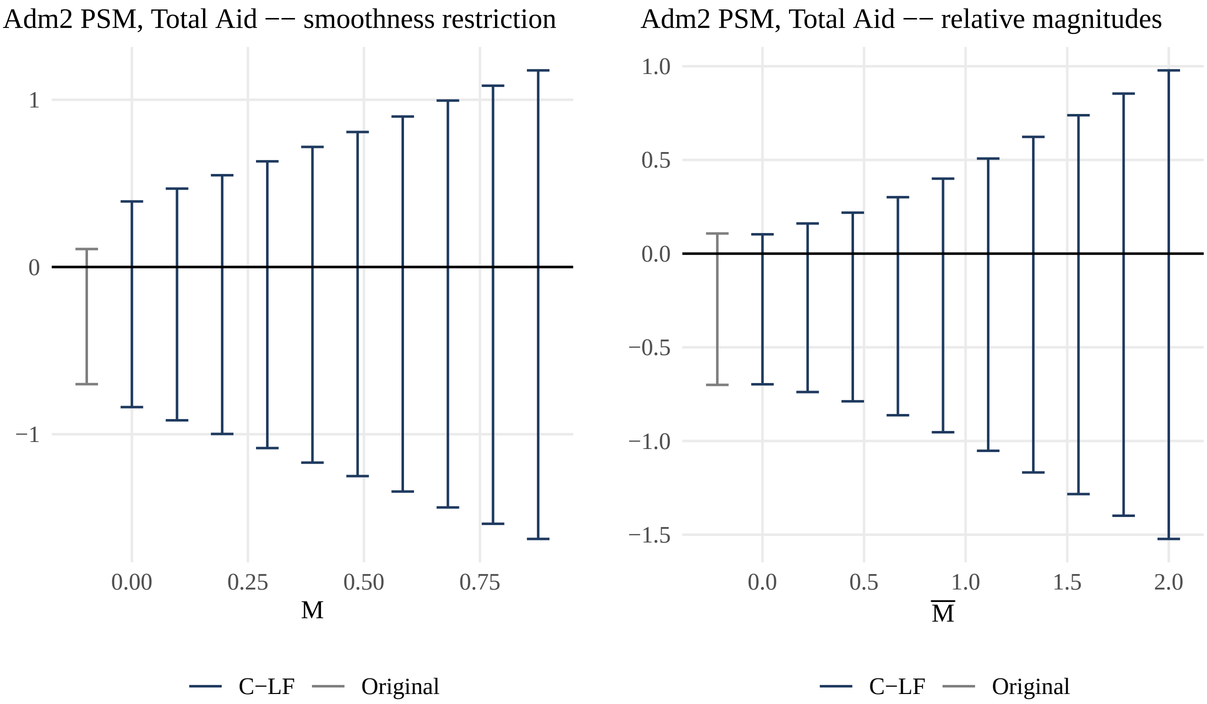

Inputs are the Callaway–Sant'Anna dynamic event-study coefficients. Left panel: smoothness restriction. Right panel: relative-magnitudes restriction. The grey segment is the original 95% confidence interval for the first post-treatment period; navy segments are robust confidence sets under increasingly generous bounds on the parallel-trends violation.

5.1.5 Headline interpretation

For the total-aid composite I fail to reject the null that the start of mineral exploration has no average effect on aid received by an ADM2 unit. The 95% confidence interval stops at roughly +8% of aid, so true responses larger than that are incompatible with the data; a district-level counterpart of the 36% country-level increase in Arezki et al. (2024) can be ruled out, noting that their benchmark is estimated for major discoveries, a stronger shock than exploration onset, and therefore bounds the expected response from above. The secondary Callaway–Sant'Anna estimate agrees on the verdict, though at much lower precision on its asinh scale.

Two qualifications apply. First, the event-study path is not flat. The leads and lags listed in Section 5.1.1 include significant departures, the corresponding Mining-Restricted event study (Annex Figure 8) shows a comparable pattern of significant distant leads and positive lags, and the joint pre-trends Wald test rejects on the headline panel. The violations are concentrated in the distant leads. Treated districts show significantly lower relative aid five and four years before exploration (-1.32 and -0.93 respectively), while the recent leads flip sign, with -3 statistically indistinguishable from zero and -2 modestly positive and significant at the 5% level, so relative aid was rising into treatment. Part of the distant-lead violation is mechanical rather than behavioural. For the 1995–99 cohorts those leads fall in calendar years 1990–94, where disbursement coverage is structurally thin. Re-estimating the event study on the 2000+ cohorts alone, whose leads lie entirely in full-coverage years, roughly halves the distant-lead coefficients (to -0.74 and -0.38) and renders the leads at -3 and -2 statistically indistinguishable from zero (Annex Figure 11). The residual run-up matters mainly for the *direction* of any bias. If the rising pre-treatment path continued mechanically after treatment, it would inflate positive post-treatment estimates. A violation of this form cannot manufacture the headline null; it works against it. The null on total aid is therefore conservative with respect to the observed pre-trend pattern, while the positive sub-composite estimates are discounted further by it. The HonestDiD bounds in Figure 5 take the full event-study path as input and quantify how much of the conclusion survives bounded violations.

The event-study figure raises the natural question of how a null pooled ATT can coexist with individually positive post-treatment coefficients. Three mechanical differences account for most of the gap: the event study benchmarks against the never-treated controls alone (the construction that makes the leads estimable) while the pooled ATT draws on the not-yet-treated comparison; the event-study path weights all horizons equally while the pooled ATT weights cohort-by-year cells by size; and the reference period $t = -1$ sits at the end of the documented pre-treatment run-up, which every lag inherits. A gradually rising lag path is, moreover, exactly what the framework's delivery-lag argument predicts for a disbursement outcome (Section 2.3), so the shape alone cannot separate a delayed disbursement response from a control-group artefact. The supporting evidence favours the artefact reading. The Callaway–Sant'Anna dynamic event study, which keeps the not-yet-treated comparison (on the asinh scale), is statistically indistinguishable from zero at every horizon (Annex Figure 9), and the ASCM trajectory, which rests on matched pre-treatment paths rather than on parallel trends, is likewise flat (Annex Figure 10). I therefore read the positive lag path as a feature of the never-treated benchmark.

Second, the cell-level appendix estimates (Annex Tables 20 and 21) carry significant negative Poisson coefficients, the opposite sign from the ADM2 sub-composites, on a precision-filtered sample whose composition differs from the ADM2 universe; the sign instability across aggregation levels reinforces a cautious reading. The honest summary is that the data show no detectable shift in total aid towards exploring districts; the donor and sector decompositions beneath the composite are taken up in Section 5.2.

The headline null replicates on the Mining-Restricted ADM2 sample under the ETWFE Poisson estimator (Annex Table 15) and under both the conditional and unconditional Callaway–Sant'Anna estimators (Annex Tables 16 and 17). The Ben-Michael et al. (2021) Augmented Synthetic Control Method deserves particular weight in light of the pre-trend discussion above. It matches each treated unit's (or treatment cohort's) own pre-treatment aid trajectory directly, so its validity rests on the quality of that fit rather than on parallel trends and its verdict does not inherit the distant-lead problem. It delivers null ATTs on the converged outcomes of both the Mining-Restricted and the PSM-Matched panels (Annex Tables 25 and 26). The Callaway–Sant'Anna anticipation sensitivity grid (Annex Table 19) reports the estimates when permitting one- and two-year treatment anticipation. Across the annex ADM2 Callaway–Sant'Anna tables the one starred entry is a marginally significant positive DAC coefficient under the conditional estimator on the Mining-Restricted sample (Annex Table 16); it does not reappear in the unconditional specification and is not read as evidence. The Pearson dispersion diagnostic (Annex Table 27) shows that the Poisson outcome variance is large but bounded, and the SE diagnostics (Annex Table 28) report the outcome zero shares alongside the cluster-robust headline estimates.

Read through the conceptual framework (Section 2.3), the pattern of a precise district-level null alongside the country-level increase documented by Arezki et al. (2024) is exactly what the decoupling channels predict. Aid negotiated as a country-level side payment, or routed through patronage and electoral logic once inside the country, raises the country's total aid receipts without moving the allocation rule facing the producing district. The design cannot distinguish a donor that does not respond to exploration from a donor whose response is absorbed elsewhere in the recipient's geography, but both readings run through the same within-country allocation politics.

5.2 Heterogeneity across donors and aid categories

The headline finding in the previous sub-section aggregates across donors and across aid categories. If the strategic-and-commercial-value mechanism behind H1 operates anywhere, it would most plausibly show up in donor-specific responses or in particular sectoral channels rather than in the pooled total-aid composite.

5.2.1 Cross-donor comparison

H2 is tested across three observable donor aggregates whose strategic interests and mandates differ: the DAC donors, the World Bank, and Chinese ODA.

Two donor aggregates require restricted windows. World Bank disbursements are recorded from 1995, so the World Bank rows are estimated on the 1995–2020 window, reusing the headline matched pairs; Table 2 summarises the estimation windows, including the handling of the 1995 cohort, which has no pre-treatment year inside the World Bank window. Chinese development finance is recorded from 2000 onwards and only as commitments; mixing pre-2000 zeros (structural absence of coverage) with post-2000 flows would inflate the never-treated control pool and bias the ATT toward zero by construction. I therefore restrict the Chinese sub-analysis to a 2000+ panel and redefine the treatment cohort as $G_i^{\mathrm{CHN}} = G_i$ for $G_i \geq 2000$ and $G_i^{\mathrm{CHN}} = \mathrm{NA}$ for $0 < G_i < 2000$ (units treated before 2000 are dropped because they have no pre-period under Chinese-aid coverage). These rows are estimated on a separately matched sample rather than on the headline pairs, built as a fresh 1:1 within-country PSM of the 2000+ treated cohorts against never-treated controls, with the matching covariates computed from a 1990–99 pre-treatment window (Table 7). The pre-treatment aid covariates measure Western (DAC and World Bank) disbursements, the only aid recorded before 2000, and proxy a district's general aid-attractiveness rather than exposure to Chinese aid specifically. The Chinese outcome is the ODA-like component of commitments. GODAD also records the OOF-like component (which dominates Chinese state financing in Africa by value), but OOF is a categorically different instrument (non-concessional, often resource-backed lending) and Dreher et al. (2018) show its allocation logic differs sharply from ODA; pooling the two would blur exactly the distinction H2 is about. The OOF channel, where the resource-seeking motive is documented to operate, is left for future work. The DAC ATTs are read off the full 1990–2020 panel.

Table 7: China-adapted ADM2 panel summary

Sample	ADM2 polygons	ADM2 treated

China-adapted (2000+)	686	343

Note: Within-country PSM on the 2000+ treated cohort against never-treated controls drawn from the full continental cross-section, with the same covariate set as the headline PSM computed from a 1990–1999 pre-treatment window. Outcome: Chinese ODA-like commitments. Cohort range: 2000–2018.

Table 8 reports the ETWFE Poisson ATTs by donor on the PSM-Matched ADM2 panel (total aid, total DAC, total WB) alongside the China-restricted ATT. All rows are link-scale semi-elasticities from the same estimator and runner, so the columns are directly comparable across donors. Table 9 reports the same comparison under the Callaway–Sant'Anna conditional estimator on asinh of aid.

Table 8: Donor heterogeneity: ETWFE Poisson ATTs by donor (ADM2 PSM-Matched)

Level	Outcome	ATT	SE	p-value
Adm2	Total Aid	0.0192	(0.0274)	0.483
	DAC Aid	-0.0467	(0.0368)	0.204
	World Bank Aid	0.1698***	(0.0435)	<0.001
	Chinese ODA	-1.2410	(5.5369)	0.823

Note: All rows are link-scale semi-elasticities. The World Bank row is estimated on the 1995+ window (start of WB disbursement coverage); the Chinese row is the ODA-like commitment component on the 2000+ subpanel where GODAD records Chinese projects. *p<0.1, **p<0.05, ***p<0.01.

Table 9: Donor heterogeneity: Callaway–Sant'Anna conditional ATTs by donor (ADM2 PSM-Matched)

Level	Outcome	ATT	SE	p-value	PT Wald p
Adm2	Total Aid	-0.1212	(0.3252)	0.709	<0.001
	DAC Aid	-0.0803	(0.2969)	0.787	<0.001
	World Bank Aid	-0.1209	(0.3475)	0.728	<0.001
	Chinese ODA	-0.0326	(0.1495)	0.827	0.830

Note: Estimates on the asinh of aid. *p<0.1, **p<0.05, ***p<0.01. PT Wald p: joint test of the pre-treatment ATT(g,t)

The donor decomposition in Table 8 is where the headline null differentiates. The DAC row is a clear null under both estimators. The World Bank row, by contrast, attracts a positive and statistically significant Poisson ATT of 0.170 (SE 0.043), roughly a 19% increase in World Bank disbursements, though the Callaway–Sant'Anna estimate for the same outcome in Table 9 is -0.121 (SE 0.348), a null. Furthermore, three caveats prevent a causal reading. The pre-trends evidence on the headline panel is imperfect (Section 5.1.5); the cell-level World Bank Poisson estimates carry the opposite (negative) sign on the precision-filtered cell sample (Annex Tables 20 and 21); and the ASCM robustness check finds no World Bank effect (Annex Table 25). The

Chinese-ODA ADM2 rows are null under both estimators (the Poisson row imprecisely so, see Section 6; the precise Chinese null is the Callaway–Sant'Anna row), and the ADM2 sub-analysis is the one specification whose joint pre-trends Wald test clearly passes (p = 0.83), making it the cleanest null in the thesis; the doubly-robust adjustment reuses the same three 1990–94 pre-treatment covariates as the headline specification, while the matching stage uses the longer 1990–99 window. The China-adapted cell-level rows, reported in Annex Table 24, repeat the divergence pattern of the other cell-level results. The Poisson coefficient is significantly negative on the precision-filtered cell sample while the Callaway–Sant'Anna coefficient is marginally positive, so the cell level offers no consistent evidence of a response in either direction and the aligned ADM2 design carries the Chinese verdict. That ADM2 null is consistent with Dreher et al. (2018). Chinese ODA is allocated primarily on foreign-policy criteria and behaves like Western aid, while the resource-pursuit motive is documented in the OOF component, which is deliberately excluded here and remains an open question for future work. The commitments margin adds a second reading. Commitments record the allocation decision itself and should carry a strategic response first and most sharply (Section 2.3), so a null on Chinese commitments is harder to attribute to delivery frictions than the disbursement-based nulls of the Western donors.

On balance, H2's no-difference prediction is not rejected by the Callaway–Sant'Anna estimator, but the Poisson estimator points to the World Bank as the one donor aggregate with a (suggestive) differential response.

5.2.2 Aid-category heterogeneity

The headline ETWFE Poisson and Callaway–Sant'Anna conditional fits in Section 5 were estimated on the pooled total-aid composite. Tables 10 and 11 report the same estimators run separately on the three OECD-DAC sectoral composites (social-sector, economic-sector, production-sector), with the total-aid row repeated as a comparison anchor.

Table 10: Sector heterogeneity: ETWFE Poisson ATTs by aid category, ADM2 PSM-Matched

Level	Outcome	ATT	SE	p-value
Adm2	Total Aid	0.0192	(0.0274)	0.483
	Social Aid	-0.0196	(0.0289)	0.498
	Economic Aid	0.0283	(0.0339)	0.405
	Production Aid	0.5340***	(0.0486)	<0.001

Note: *p<0.1, **p<0.05, ***p<0.01.

Table 11: Sector heterogeneity: Callaway–Sant'Anna conditional ATTs (asinh of aid) by aid category, ADM2 PSM-Matched

Level	Outcome	ATT	SE	p-value	PT Wald p
	Total Aid	-0.1212	(0.3252)	0.709	<0.001
	Social Aid	-0.2041	(0.3180)	0.521	<0.001
	Economic Aid	-0.5579	(0.3505)	0.111	<0.001

Adm2				
Production Aid	-0.2301	(0.2657)	0.386	<0.001

Note: *p<0.1, **p<0.05, ***p<0.01. PT Wald p: joint test of the pre-treatment ATT(g,t)

The sectoral decomposition in Table 10 shows the same two-estimator pattern as the donor split, and here it lines up with H3's prediction. Under the Poisson estimator, the social-sector and economic-sector composites are nulls while the production-sector composite, the channel with the most direct commercial and extractive exposure (it includes mining-sector aid), attracts a large positive ATT of 0.534 (SE 0.049). The ordering predicted by Dreher et al. (2024), Bermeo (2017) and Guillon & Mathonnat (2020), commercially exposed channels above humanitarian-targeted ones, is exactly what the Poisson point estimates display. The Callaway–Sant'Anna conditional estimates in Table 11, however, find no significant response for any sectoral composite (production: -0.230, SE 0.266), so the sectoral signal shares the structure of the World Bank result, positive under the Poisson estimator and absent under the Callaway–Sant'Anna estimator. The same three caveats apply, and the cell-level production-aid coefficient is again negative on the precision-filtered sample.

The verdict on H3 is therefore split by estimator. The Poisson evidence is directionally consistent with the hypothesised sectoral ordering, concentrated in production-sector aid, but the effect does not survive the Callaway–Sant'Anna specification or the pre-trends scrutiny, so H3 receives at most suggestive support.

5.3 Descriptive country-and-year FE check

As an additional descriptive check, this section reviews the geographic correlates of subnational aid receipt that the literature has identified and then estimates a country-and-year fixed-effects PPML on the same panel.

5.3.1 Competing geographic correlates from the literature

The conceptual framework (Section 2.3) implies a concrete alternative to H1. If the decoupling channels dominate, subnational aid follows the recipient's administrative, political and logistical geography rather than its mineral geography. The literature identifies three variants of that alternative geography, each with an observable ADM2-level handle.

A long tradition in the subnational aid literature emphasises *administrative geography.* Briggs (2017) shows for African subnational data that aid disproportionately reaches the regions where more of the richest wealth quintiles live rather than the poorest sub-units, and Nunnenkamp et al. (2017) confirm a similar bias toward better-endowed districts in the World Bank's allocation of project aid in India. Both readings imply that the practical anchors of subnational aid disbursement are administrative reach and logistical accessibility, captured at the ADM2 level by distance to the capital, urbanisation, and the lights/population/gross-cell-product cluster of development covariates. Briggs (2021) traces the mechanism to implementation incentives, as World Bank staff expect projects in remote areas to be harder to deliver. If administrative geography drives aid, the start of mineral exploration in peripheral districts will not move aid receipts because the donor logistics are anchored elsewhere.

A second strand emphasises *humanitarian and conflict targeting.* Findley et al. (2011) geocode approximately 65,000 foreign-aid projects across African recipients between 1989 and 2008 and use case studies of Sierra Leone, Angola and Mozambique to show that aid and conflict co-locate at fine spatial scales in ways that country-level totals cannot reveal. At the subnational scale, aid is frequently delivered to or near conflict-affected districts, a pattern country-level regressions misread as aid either causing or reducing violence. The empirical handle on this strand at the ADM2 level is geocoded conflict-event data from Raleigh et al. (2010), available from 1997 onwards.

A third strand returns to poverty- and humanitarian-need targeting. The pro-rich pattern documented by Briggs (2017) suggests that need-based allocation is at best partially implemented within recipients. The strand has no single covariate handle in this panel and is the residual hypothesis if the administrative-geography and conflict-targeting strands fail to attract a robust coefficient.

These strands point to ADM2-level dimensions that should plausibly correlate with aid receipt: distance to the capital, population, gross cell product, urbanisation, nighttime lights, and in a 1997+ sub-analysis ACLED conflict counts and fatalities. The political channels of the framework, patronage and electoral targeting (Dreher et al., 2019; Hodler & Raschky, 2014; Jablonski, 2014), have no covariate handle in this panel, since the PPML includes no leader-region or electoral variables.

5.3.2 Descriptive PPML specification

To document which of these dimensions systematically correlate with aid in the data, I estimate a country-and-year fixed-effects PPML on the panel:

$$\mathbb{E}[Y_{it} \mid X_{it}, \alpha_c, \tau_t] \;=\; \exp\Big(X_{it}'\beta + \alpha_c + \tau_t\Big),$$

where Y_{it} is total aid in ADM2 unit i in year t, X_{it} contains the asinh of contemporaneous distance to capital, nighttime lights, population and gross cell product, plus an urban-share dummy and an indicator for the presence of a mining polygon. α_c is a country fixed effect and τ_t a year fixed effect. The country fixed effect absorbs donor-recipient political relationships; the year fixed effect absorbs global aid cycles. Standard errors are cluster-robust at the province level. The model is estimated on three samples (PSM-Matched, Total Sample, and Mining-Restricted) to gauge sensitivity. In the Mining-Restricted sample every unit has a mining polygon by construction, so the indicator is dropped as collinear.

The regression is *descriptive, not causal.* There is no exogenous variation in the covariates, so the coefficients are read as conditional correlations that indicate which within-country dimensions move with aid once country and year effects are removed. One measurement constraint applies. The PRIO-GRID source series end before the panel does (calibrated nighttime lights in 2012, gross cell product in 2005), so observation-years without measured covariates are excluded rather than imputed, the effective estimation window is reported beneath each table, and the urban share, measured only every ten years, is interpolated linearly within units. A sub-specification refits the same PPML on the post-1997 truncation of each sample, adding the asinh of ADM2-year conflict-event counts and fatalities from ACLED (Raleigh et al., 2010).

Table 12: Descriptive PPML: within-country correlates of total aid, three samples

Variable	PSM	Total	Mining
Distance to capital	-0.500*** (0.188)	-0.383*** (0.087)	-0.202 (0.248)
Nighttime lights	0.818 (4.875)	-1.036 (1.124)	11.056** (5.479)
Population	-0.099 (0.094)	-0.107*** (0.034)	-0.164 (0.110)
GCP PPP	1.995*** (0.385)	0.887*** (0.245)	1.448*** (0.256)
Urban share	0.012 (0.303)	0.079** (0.035)	-0.643* (0.364)
Has mining polygon	0.414 (0.270)	0.288 (0.178)	– (collinear)

Note: Country and year fixed effects; each cell reports the PPML coefficient with cluster-robust (province-level) SE in parentheses. Distance to capital, nighttime lights, population and gross cell product enter the model in asinh; urban share and the mining-polygon indicator are dummies; row labels are shown untransformed. Observation-years without measured covariates are excluded (no imputation). Effective samples (measured covariates only): Mining: n = 1,317, 1995–2005; Total: n = 17,043, 1995–2005; PSM: n = 2,224, 1995–2005. *p<0.1, **p<0.05, ***p<0.01.

5.3.3 Results

Two covariates carry a stable pattern across the three samples in Table 12, and both belong to the administrative-geography strand. Distance to capital enters negatively throughout, with a coefficient of -0.383 in the Total Sample and -0.500 in the PSM-Matched headline sample (both significant ($p < 0.01$)), so aid disbursements decline with distance from the political centre. Gross cell product enters positively and significantly in all three samples (Total 0.887, Mining-Restricted 1.448, PSM 1.995; stars in the table). Within countries, economically larger districts receive systematically more aid, the within-country analogue of the pro-rich targeting documented by Briggs (2017) and Briggs (2021). The remaining covariates show no robust pattern. Population is significantly negative only in the Total Sample, the urban-share dummy is unstable across samples, and nighttime lights is insignificant in the Total Sample (coefficient -1.036) but implausibly large in the Mining-Restricted sample (11.06), a magnitude that points to near-collinearity with the development block in the small mining-feasible cross-section rather than to an economic effect. The mining-polygon indicator is insignificant in the PSM sample and insignificant in the Total Sample.

Table 13 reports the ACLED-augmented sub-specification on the 1997+ truncation of each sample; because the PRIO covariate series bind, the effective estimation window is 2000–2005. The conflict-event count is significant ($p < 0.01$) in the Total Sample (coefficient 0.588) and marginally significant ($p < 0.10$) in the PSM-Matched sample (coefficient 0.330); conflict fatalities likewise change sign and significance across the three samples (Total: -0.092; Mining: 0.266; PSM: -0.124). The conflict-targeting strand therefore does not survive the covariate matching that anchors the headline DiD. In the PSM-Matched sample, which is the headline analytic sample, neither conflict regressor attracts a robust coefficient (and the mining-polygon indicator is only marginally significant there). The conflict correlates thus share the instability of the night-lights pattern in Table 12, in contrast to the distance-to-capital and gross-cell-product pattern, which carries over unchanged in the ACLED specification.

Table 13: Descriptive PPML with ACLED conflict covariates (1997+ samples)

Variable	PSM	Total	Mining
Distance to capital	-0.505*** (0.192)	-0.378*** (0.074)	-0.257 (0.232)
Nighttime lights	4.027 (5.342)	0.921 (0.869)	14.094** (5.562)
Population	-0.112 (0.088)	-0.096*** (0.034)	-0.188* (0.106)
GCP PPP	2.063*** (0.306)	0.820*** (0.224)	1.418*** (0.376)
Urban share	-0.153 (0.313)	0.059 (0.037)	-0.770** (0.385)
Conflict events	0.330* (0.188)	0.588*** (0.158)	0.282 (0.227)
Conflict fatalities	-0.124 (0.124)	-0.092 (0.116)	0.266* (0.159)
Has mining polygon	0.457* (0.278)	0.263 (0.178)	– (collinear)

Note: Specification mirrors Table 12 with two extra regressors: the conflict-event count and conflict fatalities from a point-in-polygon spatial join of ACLED events to the ADM2 grid. All continuous covariates (including the two conflict regressors) enter in asinh, urban share and the mining-polygon indicator are dummies, and row labels are shown untransformed. Observation-years without measured covariates are excluded (no imputation). Effective samples (measured covariates only): Mining: n = 878, 2000–2005; Total: n = 11,394, 2000–2005; PSM: n = 1,496, 2000–2005. Country and year fixed effects, cluster-robust SEs at province. *p<0.1, **p<0.05, ***p<0.01.

The PPML thus identifies two robust within-country correlates of aid receipt, proximity to the capital and local economic size, both consistent with the administrative-geography reading in which donor logistics and implementation feasibility anchor subnational disbursement (Briggs, 2017, 2021). Two readings follow. First, the PPML supports the DiD null on the narrow point that contemporaneous mineral status does not add explanatory power to aid prediction once these administrative-geographic controls are absorbed. Second, the poverty- and humanitarian-need-targeting strand has no single covariate handle in this panel and cannot be tested directly by the PPML.

6 Limitations

Five limitations bound what can be inferred from the results.

(i) *Statistical power, unevenly distributed.* Not every null in the thesis is equally informative. The headline confidence interval is tight (Section 5.1.5), but the Callaway–Sant'Anna confidence intervals on the asinh scale are an order of magnitude wider, and the China-window ADM2 Poisson estimate is essentially uninformative (its standard error exceeds 5 log points; the Chinese verdict rests on the far more precise Callaway–Sant'Anna row). Each null in the text is therefore read against its confidence interval rather than treated as evidence of absence by default.

(ii) *Standard-error precision under heavy outcome dispersion.* The Poisson PML estimator is consistent for the conditional mean function regardless of the variance, but the cluster-robust SEs scale with the square root of the dispersion. Annex Table 27 reports the Pearson dispersion across the headline Poisson fits, and Annex Table 28 reports the outcome zero shares alongside the cluster-robust estimates. In this panel, where mineral-exploration timing and aid disbursements are spatially clustered within provinces, clustering one level above the unit of observation is the appropriate inference.

(iii) *Data-quality features of the underlying GODAD and the pre-trends evidence.* GODAD project-level coverage starts in 1973 but disbursement dollar values are only systematic from 1990 onwards, which anchors the panel at 1990; disbursement coverage in the early 1990s is nonetheless thin, so the 1990–94 pre-treatment aid covariates are sparse and the matching is anchored mainly by the geography covariates. World Bank disbursements begin in 1995 and Chinese records (commitments only) in 2000, which is why the WB and Chinese estimates use restricted windows. The pre-trends evidence on the headline panel is imperfect. The event study shows significant negative leads at distant horizons and the joint Wald test rejects, while the Chinese sub-analysis passes the same test cleanly. The cohort-restriction diagnostic in Section 5.1.5 attributes roughly half of the distant-lead violation to the thin pre-1995 coverage window, and the direction of the residual violation (a rising pre-treatment path) biases toward finding positive effects, not toward the null. The headline conclusions are therefore anchored on the HonestDiD sensitivity bounds, which remain valid under bounded violations, and the positive sub-composite estimates are flagged as suggestive rather than causal precisely because of this limitation.

(iv) *Treatment measurement error.* As discussed, the USGS compilation records exploration events from heterogeneous sources and is not a complete census. Related USGS data products on African mineral activity have been characterised as carrying limited and patchy coverage (Cust & Poelhekke, 2015), and similar concerns may apply here. Under-reporting would attenuate the ATT toward zero, so the null finding is bounded by the quality of the treatment proxy. The treatment moreover captures exploration onset rather than confirmed discoveries: a donor response concentrated at the discovery stage would appear in the design only with a lag, where a discovery follows the first recorded exploration inside the event window, and would be attenuated where it does not. I contacted USGS for further information on the data collection process but did not receive an answer.

(v) *Covariate staleness and measurement.* The pre-treatment covariates are constructed from 1990–94 unit averages; for cohorts treated late in the panel (e.g. 2018) this baseline is more than two decades stale, which limits the precision of the doubly-robust adjustment. The contemporaneous covariates used in the descriptive PPML inherit the PRIO-GRID source coverage, which ends before the panel does, so the PPML is estimated on the shorter window reported beneath its tables.

7 Conclusion

This thesis tests three hypotheses about the relationship between mineral exploration and subnational foreign aid in Africa. H1 predicts a positive aid shift in a district following the start of mineral exploration. H2 predicts no systematic difference in that response across OECD-DAC, World Bank and Chinese-ODA aid. H3 predicts heterogeneity across aid categories, with economic-infrastructure and production-sector aid responding more strongly than humanitarian-targeted social-sector aid. Using a 1990–2020 panel of 50 African countries and territories combining GODAD aid disbursements (Chinese ODA commitments for the post-2000 sub-analysis), the USGS Africa-minerals compilation, the Maus et al. (2022) mining polygons and PRIO-GRID covariates, and applying two staggered-DiD estimators with within-country propensity-score matching, I reach a differentiated verdict. H1 is not supported for total aid. The start of mineral exploration produces no detectable shift in a district's subsequent total aid flows. Beneath the composite, the Poisson estimator attributes significant positive responses to World Bank aid and production-sector aid that the Callaway–Sant'Anna estimator does not reproduce; because the pre-trends evidence on the headline panel is imperfect and the cell-level estimates flip sign, I read these positive estimates as suggestive rather than causal. H2's no-difference prediction survives under the Callaway–Sant'Anna estimator, with the World Bank as the one donor aggregate flagged by the Poisson estimator; the Chinese-ODA null at the ADM2 level is the cleanest result in the thesis, passing the formal pre-trends test. H3 receives at most suggestive support, concentrated in production-sector aid under the Poisson estimator only.

The descriptive country-and-year fixed-effects PPML in Section 5.3 supports the DiD null on the narrow point that contemporaneous mineral status does not predict aid once standard administrative-geographic controls are absorbed, and identifies proximity to the capital and local economic size as the two robust within-country correlates of aid receipt, while the remaining geographic controls vary across samples.

Given the recent strategic agreements, in which donor governments bundle aid into broader resource-access arrangements with African producer countries, these findings might change with more current data, a more recent timeframe, and outcome measures that capture the non-concessional financing instruments (such as Chinese OOF) through which the resource-access motive is documented to operate. In the framework's terms (Section 2.3), those agreements operate through the country-level bargaining channel, which predicts responses in the country's total aid rather than in the producing district, so the district-level null found here does not contradict their prominence per se.

The conceptual framework also identifies heterogeneity that this thesis does not test and that is left as room for further work: splitting treated cohorts by the recipient's regime type at treatment, interacting treatment with the leader's birth and co-ethnic regions, conditioning on the producing district's pre-treatment protest exposure, and re-running the design on commitments rather than disbursements to capture the allocation decision before delivery frictions. All four are feasible on the panel assembled here.

A natural extension is to move from the allocation question to the implications question: do the spatial patterns of aid in mineral-active districts have downstream consequences? Gehring et al. (2022) and Moscona (2025) show, for example, that aid can both raise and lower local conflict probability depending on donor type and project-management quality, but multiple other potential consequences leave room for further research.

Finally, as mentioned earlier Arezki & Ploeg (2025) argue that the super-power race for critical minerals risks a "critical minerals curse". One channel runs through national politics, where geopolitical rents for leaders aligning with the competing powers are back, slowing or reversing democratisation and dimming the prospects for the redistributive institutions that would turn resource proceeds into benefits for citizens. On that reading, aid that follows resource-rich countries (Arezki et al., 2024) rather than need is troubling no matter where it lands inside the country. It is no more reassuring that this aid does not reach the mineral-bearing districts within those countries, since they carry the environmental, social and conflict costs of extraction.

Even if the strategic motive is taken as given, donors retain discretion over how they engage. Arezki & Ploeg (2025) call on advanced economies to build an international governance model for critical minerals, and on producer governments to ensure that investment delivers local content, environmental protection and jobs in extracting communities. A donor persuaded by that logic would direct compensating support to the districts bearing the costs of extraction. The evidence here indicates that, over this period, none was.

References

Adhikari, B. (2019). United nations general assembly voting and foreign aid bypass. *International Politics*, *56*(4), 514–535. https://doi.org/10.1057/s41311-018-0151-3

Alesina, A., & Dollar, D. (2000). Who gives foreign aid to whom and why? *Journal of Economic Growth*, *5*(1), 33–63. https://doi.org/10.1023/A:1009874203400

Amisi, E. M., & Nkuba, B. (2026). *Minerals for peace – or war over minerals in the DR Congo?* Development and Cooperation (D+C). https://www.dandc.eu/en/article/critical-minerals-security-and-peace-deal-drc

Anaxagorou, C., Efthyvoulou, G., & Sarantides, V. (2020). Electoral motives and the subnational allocation of foreign aid in sub-Saharan Africa. *European Economic Review*, *127*, 103430. https://doi.org/10.1016/j.euroecorev.2020.103430

Arezki, R., Camara, Y., Ploeg, F. van der, & Rota-Graziosi, G. (2024). *Bad samaritans in foreign aid: Evidence from major mineral discoveries* (Working Paper 11427). CESifo. https://hdl.handle.net/10419/307357

Arezki, R., & Ploeg, F. van der. (2025). *The Critical Minerals Curse* (CESifo Working Paper 11966). CESifo. https://hdl.handle.net/10419/322528

Baumgartner, A. (2026). *The price of peace: US strategy and the DRC's critical minerals*. Review of African Political Economy (ROAPE) Blog. https://roape.net/2026/02/04/the-price-of-peace-us-strategy-and-the-drcs-critical-minerals/

Becker, V. (2026). *Can Brussels make "unsentimental" foreign aid a reality?* Euractiv. https://www.euractiv.com/section/economy/news/can-brussels-make-unsentimental-foreign-aid-a-reality

Bellemare, M. F., & Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, *82*(1), 50–61. https://doi.org/10.1111/obes.12325

Ben-Michael, E., Feller, A., & Rothstein, J. (2021). The augmented synthetic control method. *Journal of the American Statistical Association*, *116*(536), 1789–1803. https://doi.org/10.1080/01621459.2021.1929245

Berman, N., Couttenier, M., Rohner, D., & Thoenig, M. (2017). This mine is mine! How minerals fuel conflicts in Africa. *American Economic Review*, *107*(6), 1564–1610. https://doi.org/10.1257/aer.20150774

Bermeo, S. B. (2017). Aid allocation and targeted development in an increasingly connected world. *International Organization*, *71*(4), 735–766. https://doi.org/10.1017/S0020818317000315

Berthélemy, J.-C. (2006). Bilateral donors' interest vs. Recipients' development motives in aid allocation: Do all donors behave the same? *Review of Development Economics*, *10*(2), 179–194. https://doi.org/10.1111/j.1467-9361.2006.00311.x

Bomprezzi, P., Dreher, A., Fuchs, A., Hailer, T., Kammerlander, A., Kaplan, L. C., Marchesi, S., Masi, T., Perlik, K., & Robert, C. (2025). *Wedded to prosperity? Spousal favoritism in foreign aid and regional development* [Discussion Paper]. CEPR Geopolitics Research Programme. https://grp.cepr.org/publications/discussion-paper/wedded-prosperity-informal-influence-and-regional-favoritism

Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, *91*(6), 3253–3285. https://doi.org/10.109

3/restud/rdae007
Briggs, R. C. (2017). Does foreign aid target the poorest? *International Organization*, *71*(1), 187–206. https://doi.org/10.1017/S0020818316000345
Briggs, R. C. (2021). Why does aid not target the poorest? *International Studies Quarterly*, *65*(3), 739–752. https://doi.org/10.1093/isq/sqab035
Bulíř, A., & Hamann, A. J. (2008). Volatility of development aid: From the frying pan into the fire? *World Development*, *36*(10), 2048–2066. https://doi.org/10.1016/j.worlddev.2007.02.019
Burgess, R., Jedwab, R., Miguel, E., Morjaria, A., & Padró i Miquel, G. (2015). The value of democracy: Evidence from road building in Kenya. *American Economic Review*, *105*(6), 1817–1851. https://doi.org/10.1257/aer.20131031
Caliendo, M., & Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, *22*(1), 31–72. https://doi.org/10.1111/j.1467-6419.2007.00527.x
Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001
Canavire-Bacarreza, G., Nunnenkamp, P., Thiele, R., & Triveño, L. (2005). *Assessing the allocation of aid: Developmental concerns and the self-interest of donors* (Working Paper 1253). Kiel Institute for the World Economy.
Center for Strategic and International Studies. (2024). *A window of opportunity to build critical mineral security in Africa*. CSIS Analysis. https://www.csis.org/analysis/window-opportunity-build-critical-mineral-security-africa
Chen, J., & Roth, J. (2024). Logs with zeros? Some problems and solutions. *Quarterly Journal of Economics*, *139*(2), 891–936. https://doi.org/10.1093/qje/qjad054
Christensen, D. (2019). Concession stands: How mining investments incite protest in Africa. *International Organization*, *73*(1), 65–101. https://doi.org/10.1017/S0020818318000413
Cust, J., & Mihalyi, D. (2017). The presource curse. *Finance & Development*, *54*(4), 36–40.
Cust, J., & Poelhekke, S. (2015). The local economic impacts of natural resource extraction. *Annual Review of Resource Economics*, *7*, 251–268. https://doi.org/10.1146/annurev-resource-100814-125106
Dollar, D., & Levin, V. (2006). The increasing selectivity of foreign aid, 1984–2003. *World Development*, *34*(12), 2034–2046. https://doi.org/10.1016/j.worlddev.2006.06.002
Donaldson, D., & Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, *30*(4), 171–198. https://doi.org/10.1257/jep.30.4.171
Dreher, A., & Fuchs, A. (2015). Rogue aid? An empirical analysis of China's aid allocation. *Canadian Journal of Economics*, *48*(3), 988–1023. https://doi.org/10.1111/caje.12166
Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., & Tierney, M. J. (2019). African leaders and the geography of China's foreign assistance. *Journal of Development Economics*, *140*, 44–71. https://doi.org/10.1016/j.jdeveco.2019.04.003
Dreher, A., Fuchs, A., Parks, B., Strange, A. M., & Tierney, M. J. (2018). Apples and dragon fruits: The determinants of aid and other forms of state financing from China to Africa. *International Studies Quarterly*, *62*(1), 182–194. https://doi.org/10.1093/isq/sqx052
Dreher, A., Lang, V., & Reinsberg, B. (2024). Aid effectiveness and donor motives. *World Development*, *176*, 106501. https://doi.org/10.1016/j.worlddev.2023.106501
Eurodad and Counter Balance and Oxfam. (2024). *Who profits from the Global Gateway?*

The EU's new strategy for development cooperation. Eurodad / Counter Balance / Oxfam, Brussels. https://assets.nationbuilder.com/eurodad/pages/4589/attachments/original/1728405785/01_EU-GG-report-2024-FINAL.pdf

European Commission. (2023a). *EU and Namibia adopt first roadmap (2024–2025) of the strategic partnership on sustainable raw materials value chains and renewable hydrogen.* Press Release IP/23/5263, Brussels. https://ec.europa.eu/commission/presscorner/detail/en/ip_23_5263

European Commission. (2023b). *EU signs strategic partnerships on critical raw materials value chains with the Democratic Republic of the Congo and Zambia.* Press Release IP/23/5303, Brussels. https://ec.europa.eu/commission/presscorner/detail/en/ip_23_5303

European Commission. (2024). *EU and Rwanda sign a Memorandum of Understanding on Sustainable Raw Materials Value Chains.* Press Release IP/24/822, Brussels. https://ec.europa.eu/commission/presscorner/detail/en/ip_24_822

European Parliament. (2025). *MEPs want to suspend EU–Rwanda deal on sustainable value chains for critical raw materials.* Press release on the resolution of 13 February 2025 on the escalation of violence in the eastern DRC, ref. 20250206IPR26752. https://www.europarl.europa.eu/news/en/press-room/20250206IPR26752/meps-want-to-suspend-eu-rwanda-deal-on-critical-raw-materials

European Union. (2024). *Regulation (EU) 2024/1252 of the European Parliament and of the Council of 11 April 2024 establishing a framework for ensuring a secure and sustainable supply of critical raw materials (critical raw materials act).* Official Journal of the European Union, L series. https://eur-lex.europa.eu/eli/reg/2024/1252/oj

Extractive Industries Transparency Initiative. (2024). *Sicomines: How the EITI in DRC helped secure USD 4 billion in additional revenue.* EITI Blog Post, Oslo. https://eiti.org/blog-post/sicomines-how-eiti-drc-helped-secure-4-billion-additional-revenue

Findley, M. G., Powell, J., Strandow, D., & Tanner, J. (2011). The localized geography of foreign aid: A new dataset and application to violent armed conflict. *World Development*, *39*(11), 1995–2009. https://doi.org/10.1016/j.worlddev.2011.07.022

Gehring, K., Kaplan, L. C., & Wong, M. H. L. (2022). China and the World Bank — How contrasting development approaches affect the stability of African states. *Journal of Development Economics*, *158*, 102902. https://doi.org/10.1016/j.jdeveco.2022.102902

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014

Gotway, C. A., & Young, L. J. (2002). Combining incompatible spatial data. *Journal of the American Statistical Association*, *97*(458), 632–648. https://doi.org/10.1198/016214502760047140

Guillon, M., & Mathonnat, J. (2020). What can we learn on Chinese aid allocation motivations from available data? A sectorial analysis of Chinese aid to African countries. *China Economic Review*, *60*, 101265. https://doi.org/10.1016/j.chieco.2019.01.004

Hodler, R., & Raschky, P. A. (2014). Regional favoritism. *Quarterly Journal of Economics*, *129*(2), 995–1033. https://doi.org/10.1093/qje/qju004

Hoeffler, A., & Outram, V. (2011). Need, merit, or self-interest — what determines the allocation of aid? *Review of Development Economics*, *15*(2), 237–250. https://doi.org/10.1111/j.1467-9361.2011.00605.x

Hudson, J. (2013). Promises kept, promises broken? The relationship between aid commitments and disbursements. *Review of Development Finance*, *3*(3), 109–120. https://doi.org/10.1016/j.rdf.2013.08.001

Imbens, G. W., & Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, *47*(1), 5–86. https://doi.org/10.1257/jel.47.1.5

International Energy Agency. (2025). *Global critical minerals outlook 2025*. IEA Report. https://www.iea.org/reports/global-critical-minerals-outlook-2025

Jablonski, R. S. (2014). How aid targets votes: The impact of electoral incentives on foreign aid distribution. *World Politics*, *66*(2), 293–330. https://doi.org/10.1017/S0043887114000045

Kegley, C. W., Jr., & Hook, S. W. (1991). U.s. Foreign aid and U.N. Voting: Did Reagan's linkage strategy buy deference or defiance? *International Studies Quarterly*, *35*(3), 295–312. https://doi.org/10.2307/2600702

Kuziemko, I., & Werker, E. (2006). How much is a seat on the Security Council worth? Foreign aid and bribery at the United Nations. *Journal of Political Economy*, *114*(5), 905–930. https://doi.org/10.1086/507155

Lumsdaine, D. H. (1993). *Moral vision in international politics: The foreign aid regime, 1949–1989*. Princeton University Press.

Maus, V., Giljum, S., Silva, D. M. da, Gutschlhofer, J., Rosa, R. P. da, Luckeneder, S., Gass, S. L. B., Lieber, M., & McCallum, I. (2022). An update on global mining land use. *Scientific Data*, *9*, 433. https://doi.org/10.1038/s41597-022-01547-4

Michalopoulos, S., & Papaioannou, E. (2014). National institutions and subnational development in Africa. *Quarterly Journal of Economics*, *129*(1), 151–213. https://doi.org/10.1093/qje/qjt029

Morrison, K. M. (2011). As the world bank turns: Determinants of IDA lending in the cold war and after. *Business and Politics*, *13*(2), Article 2. https://doi.org/10.2202/1469-3569.1365

Moscona, J. (2025). The management of aid and conflict in Africa. *American Economic Journal: Economic Policy*, *17*(4), 228–259. https://doi.org/10.1257/pol.20220219

Murphy, L. T., & Elimä, N. (2021). *In broad daylight: Uyghur forced labour and global solar supply chains*. Sheffield Hallam University Helena Kennedy Centre for International Justice.

Nunnenkamp, P., Öhler, H., & Sosa Andrés, M. (2017). Need, merit and politics in multilateral aid allocation: A district-level analysis of World Bank projects in India. *Review of Development Economics*, *21*(1), 126–156. https://doi.org/10.1111/rode.12259

Öhler, H., & Nunnenkamp, P. (2014). Needs-based targeting or favoritism? The regional allocation of multilateral aid within recipient countries. *Kyklos*, *67*(3), 420–446. https://doi.org/10.1111/kykl.12061

Padilla, A. D., Otarod, D., Deloach-Overton, S. W., Kemna, R. F., Freeman, P. A., Wolfe, E. R., Bird, L. R., Gulley, A. L., Trippi, M. H., Dicken, C., Hammarstrom, J. M., & Brioche, A. S. (2021). *Compilation of geospatial data (GIS) for the mineral industries and related infrastructure of Africa*. U.S. Geological Survey data release. https://doi.org/10.5066/P97EQWXP

Perrotta Berlin, M., Desai, R. M., & Olofsgård, A. (2023). Trading favors? UN Security Council membership and subnational favoritism in aid recipients. *The Review of International Organizations*, *18*(2). https://doi.org/10.1007/s11558-022-09467-z

Ploeg, F. van der. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature*, *49*(2), 366–420. https://doi.org/10.1257/jel.49.2.366

Raleigh, C., Linke, A., Hegre, H., & Karlsen, J. (2010). Introducing ACLED: An armed conflict location and event dataset. *Journal of Peace Research*, *47*(5), 651–660. https://doi.org/10.1177/0022343310378914

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, *90*(5), 2555–2591. https://doi.org/10.1093/restud/rdad018

Sachs, J. D., & Warner, A. M. (2001). The curse of natural resources. *European Economic Review*, *45*(4-6), 827–838. https://doi.org/10.1016/S0014-2921(01)00125-8

Schoonover, N. (2025). *China in Africa: March 2025*. Council on Foreign Relations. https://www.cfr.org/articles/china-africa-march-2025

Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, *25*(1), 1–21. https://doi.org/10.1214/09-STS313

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

Tollefsen, A. F., Strand, H., & Buhaug, H. (2012). PRIO-GRID: A unified spatial data structure. *Journal of Peace Research*, *49*(2), 363–374. https://doi.org/10.1177/0022343311431287

Tsang, K. (2025). *Mining the energy transition: Chinese critical mineral investment and protest in sub-Saharan Africa* [Job Market Paper]. KU Leuven; UCLouvain.

Voice of America. (2024). *Uncertainties remain with renegotiated Chinese mining deal in DRC*. VOA News. https://www.voanews.com/a/uncertainties-remain-with-renegotiated-chinese-mining-deal-in-drc-/7458908.html

Wooldridge, J. M. (1999). Quasi-likelihood methods for count data. In M. H. Pesaran & P. Schmidt (Eds.), *Handbook of applied econometrics, volume II: Microeconomics* (pp. 321–368). Blackwell. https://doi.org/10.1111/b.9780631216339.1999.00009.x

Wooldridge, J. M. (2023). Simple approaches to nonlinear difference-in-differences with panel data. *Econometrics Journal*, *26*(3), C31–C66. https://doi.org/10.1093/ectj/utad016

Appendix

A. Balance and covariate overlap

Table 14: Covariate balance, Mining-Restricted sample (ADM2)

Variable	T mean	C mean	Norm. diff.	p-value	Group
Pre-treatment Total Aid	1.27e+05	1.04e+04	0.096	0.294	Pre-treatment Aid
Pre-treatment Social Aid	4.34e+03	334	0.167	0.070	
Pre-treatment Economic Aid	1.23e+05	1e+04	0.093	0.311	
Pre-treatment Production Aid	101	0	0.092	0.318	
Donors Active (1990-94)	0.0227	0.0157	0.075	0.415	
P(Aid > 0)	0.0227	0.0157	0.075	0.415	
Nighttime Lights (1990-94)	0.0269	0.0279	-0.017	0.855	Geography & Economic Development
Population (1990-94)	544	541	0.005	0.960	
GCP PPP (1990-94)	0.919	0.647	0.113	0.220	
Distance to Capital (km)	487	400	0.242	0.009	
Urban Intensity (1990-94)	0.195	0.248	-0.076	0.411	

Note: N treated = 238, N control = 230. Pre-window 1990–1994 (disbursements; DAC flows only, since WB disbursement coverage starts in 1995).

Figure 6: Covariate overlap: ADM2 Mining-Restricted

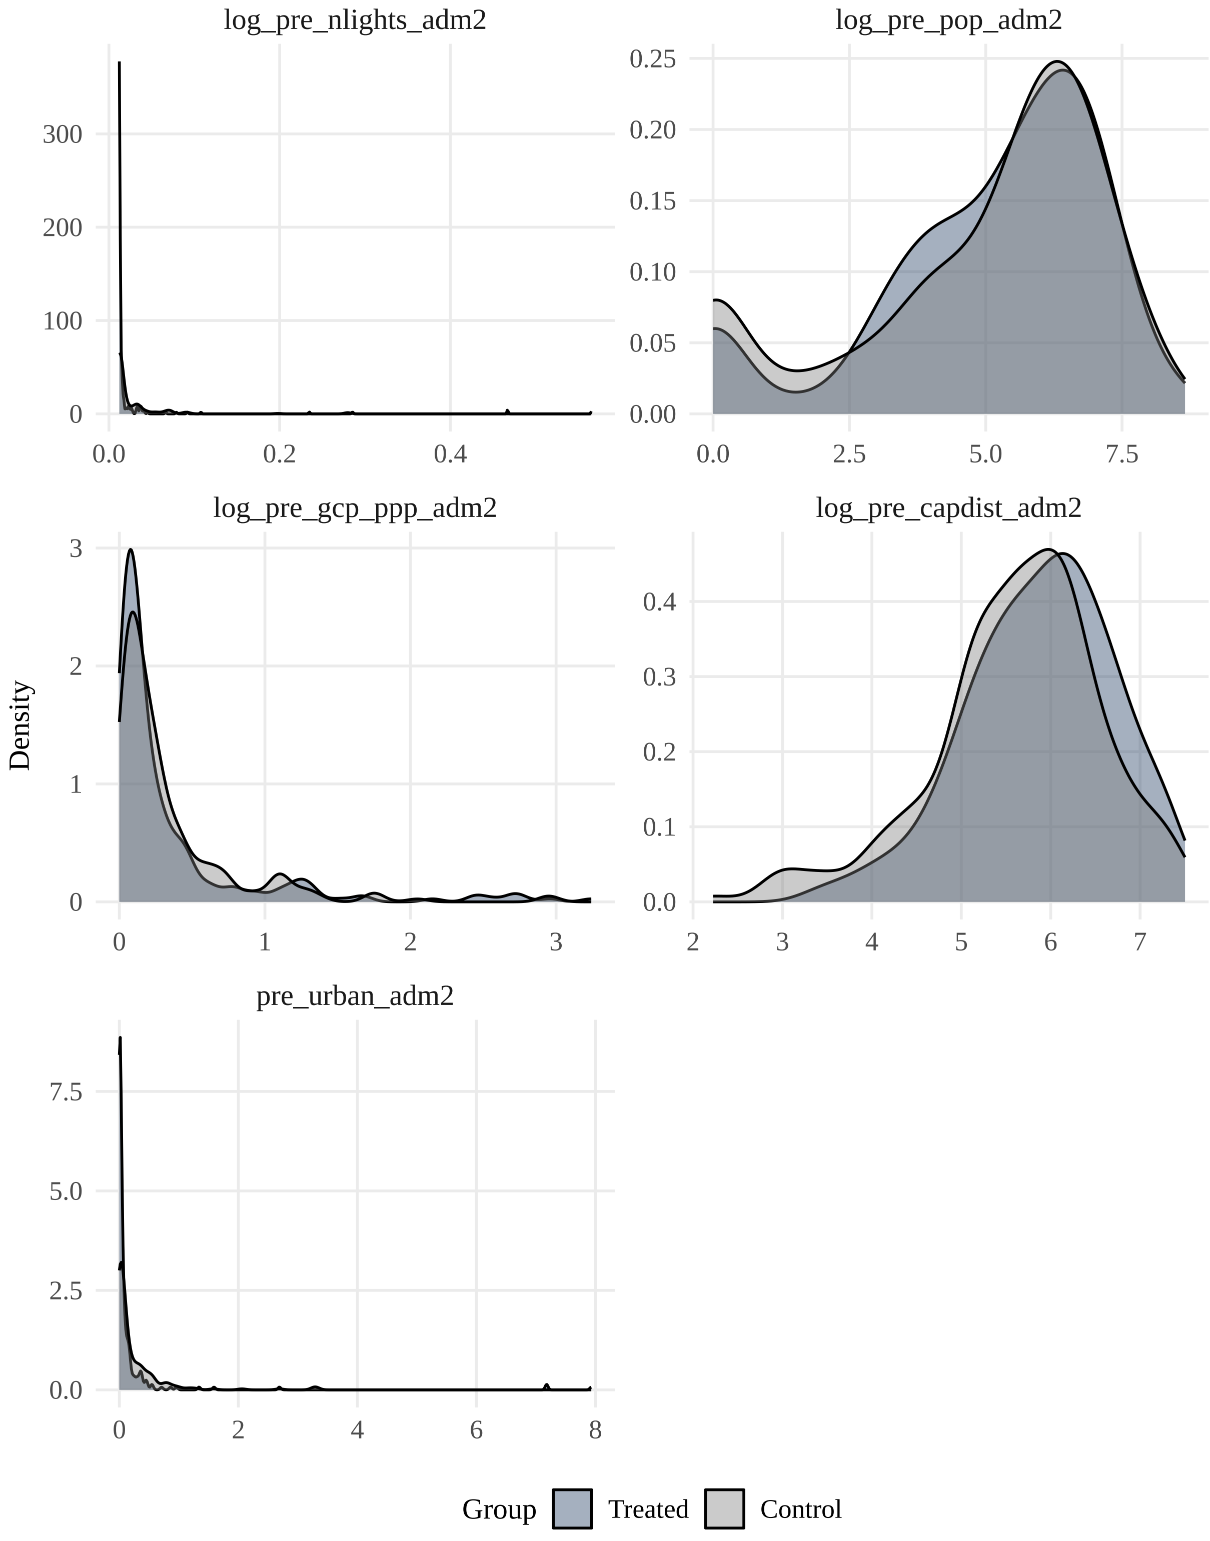

Density curves of the pre-treatment covariates by treated/control status.

Figure 7: Covariate overlap: ADM2 PSM-Matched

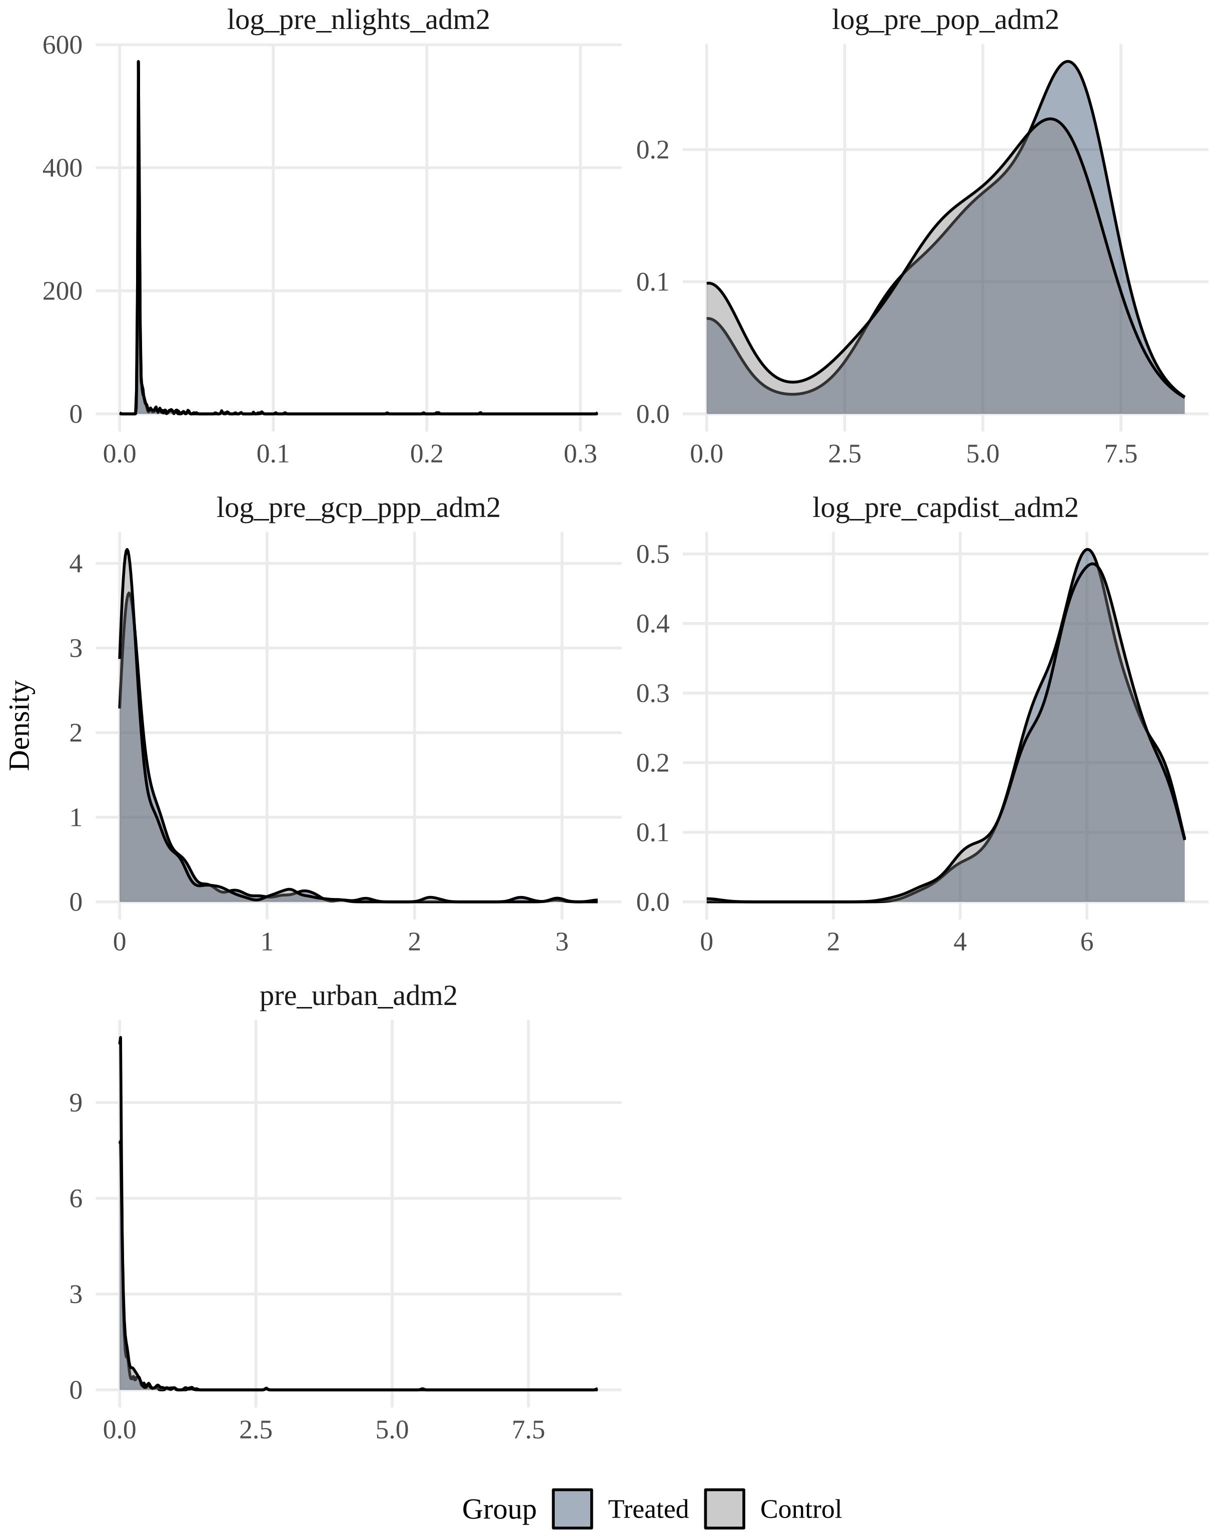

Within-country matching brings the treated and control densities into tighter alignment than the Mining-Restricted figure above.

B. Mining-Restricted bias check

Table 15: ETWFE Poisson ATT, ADM2 Mining-Restricted

Level	Outcome	ATT	SE	p-value
Adm2	Total Aid	0.0458	(0.0417)	0.272
	DAC Aid	0.0590	(0.0500)	0.238
	World Bank Aid	0.2724***	(0.0643)	<0.001
	Social Aid	-0.0331	(0.0479)	0.489
	Economic Aid	0.0904**	(0.0435)	0.038
	Production Aid	0.5310***	(0.0942)	<0.001

Note: The World Bank row is estimated on the 1995+ window.
*p<0.1, **p<0.05, ***p<0.01.

Figure 8: ETWFE Poisson event study: ADM2 Mining-Restricted, Total Aid

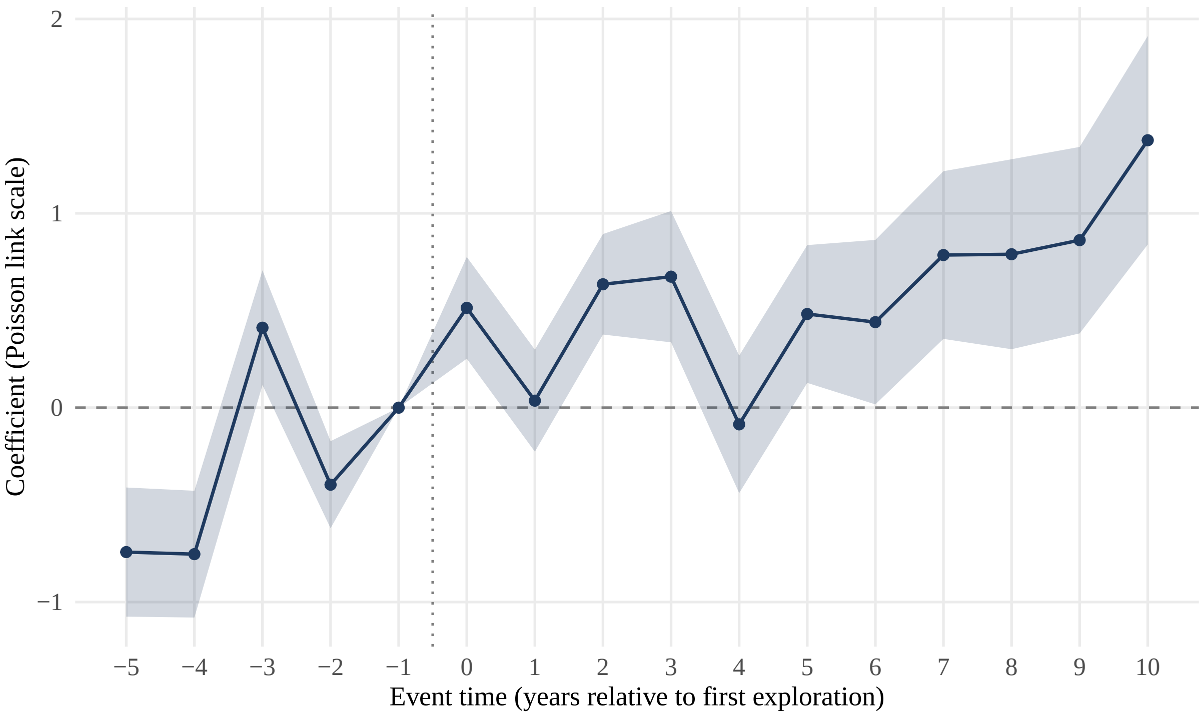

Never-treated controls; event-time coefficients on the link scale relative to $t = -1$. Referenced in Section 5.1.5.

Table 16: Callaway–Sant'Anna conditional, ADM2 Mining-Restricted

Level	Outcome	ATT	SE	p-value
	Total Aid	0.3435	(0.4322)	0.427
	DAC Aid	0.7079*	(0.3854)	0.066
	World Bank Aid	-0.1365	(0.4519)	0.763
	Social Aid	-0.0017	(0.4066)	0.997
	Economic Aid	0.3702	(0.4312)	0.391

Level	Outcome	ATT	SE	p-value
Adm2				
	Production Aid	-0.0718	(0.4026)	0.858

Note: *p<0.1, **p<0.05, ***p<0.01.

Table 17: Callaway–Sant'Anna unconditional, ADM2 Mining-Restricted

Level	Outcome	ATT	SE	p-value
Adm2	Total Aid	0.2685	(0.4206)	0.523
	DAC Aid	0.5943	(0.3776)	0.116
	World Bank Aid	-0.1270	(0.4145)	0.759
	Social Aid	0.0640	(0.4114)	0.876
	Economic Aid	0.0583	(0.4279)	0.892
	Production Aid	-0.0355	(0.4011)	0.929

Note: *p<0.1, **p<0.05, ***p<0.01.

C. Additional Callaway–Sant'Anna results

Table 18: Callaway–Sant'Anna unconditional ATT (no doubly-robust covariates), ADM2 PSM-Matched

Level	Outcome	ATT	SE	p-value
Adm2	Total Aid	0.0423	(0.3257)	0.897
	DAC Aid	0.1431	(0.2686)	0.594
	World Bank Aid	-0.0682	(0.3384)	0.840
	Social Aid	-0.0522	(0.3143)	0.868
	Economic Aid	-0.4454	(0.3407)	0.191
	Production Aid	-0.1511	(0.2784)	0.587

Note: *p<0.1, **p<0.05, ***p<0.01.

Figure 9: Callaway–Sant'Anna dynamic event study: ADM2 PSM-Matched, asinh of Total Aid

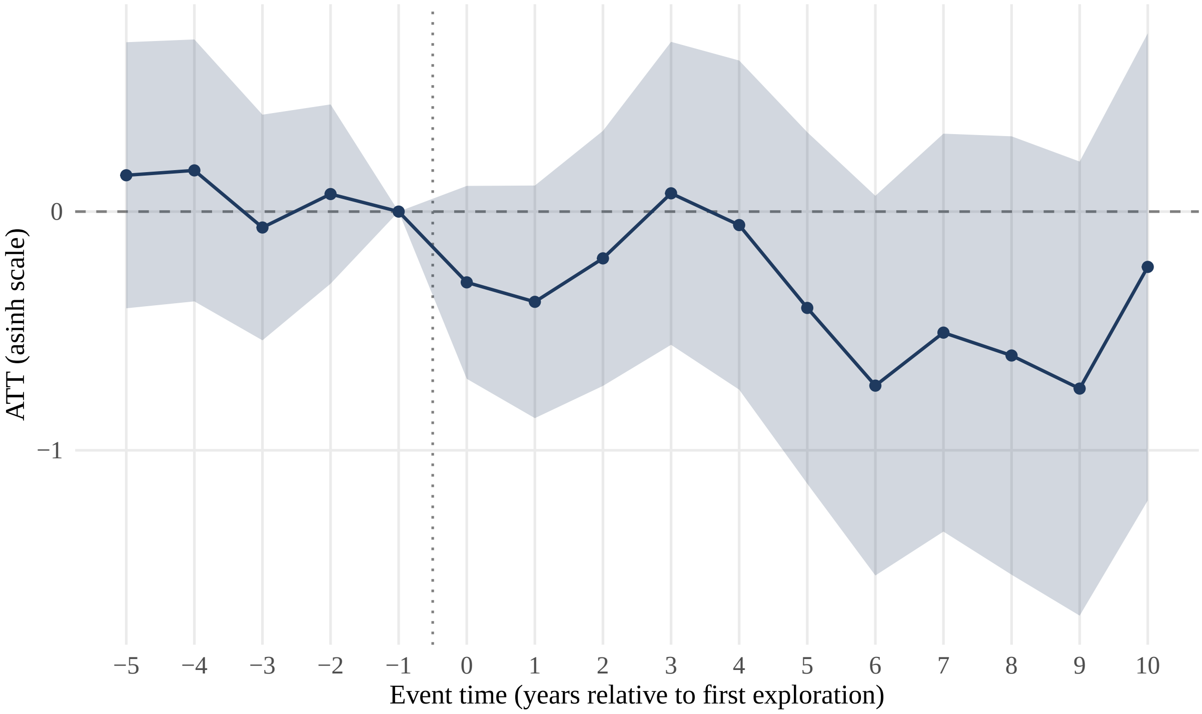

Not-yet-treated controls, universal base period. These are the event-study coefficients that enter the HonestDiD analysis (Figure 5); confidence intervals use the unit-level influence-function covariance, the finest level the influence functions support. In contrast to the never-treated-control ETWFE event study (Figure 4), no coefficient is statistically distinguishable from zero at any horizon. Referenced in Section 5.1.5.

Table 19: Callaway–Sant'Anna anticipation sensitivity grid

Sample	Level	Outcome	Anticipation	ATT	SE	p-value
Mining-Restricted	Adm2	Total Aid	0	0.3435	(0.4251)	0.419
			1	0.3266	(0.4501)	0.468
			2	0.1977	(0.5030)	0.694
	Cell	World Bank Aid	0	0.0815	(0.4982)	0.870
			1	0.1757	(0.4569)	0.701
			2	0.3581	(0.4818)	0.457
PSM	Adm2	Total Aid	0	-0.1212	(0.3053)	0.692
			1	-0.3090	(0.3584)	0.389
			2	-0.2029	(0.3558)	0.568
	Cell	World Bank Aid	0	-0.0634	(0.3524)	0.857
			1	-0.0142	(0.3246)	0.965
			2	-0.1831	(0.4083)	0.654

Note: Anticipation window of 0, 1 or 2 years; estimates on the asinh scale. *p<0.1, **p<0.05, ***p<0.01.

D. Cell-level results

Table 20: ETWFE Poisson ATT, Cell-level Mining-Restricted

Level	Outcome	ATT	SE	p-value
Cell	World Bank Aid	-0.3809***	(0.0360)	<0.001
	Social Aid	-0.8396***	(0.0443)	<0.001
	Economic Aid	-0.1869***	(0.0155)	<0.001
	Production Aid	-1.0981***	(0.0976)	<0.001

Note: *p<0.1, **p<0.05, ***p<0.01.

Table 21: ETWFE Poisson ATT, Cell-level PSM-Matched

Level	Outcome	ATT	SE	p-value
Cell	World Bank Aid	-0.1395***	(0.0248)	<0.001
	Social Aid	-0.4241***	(0.0346)	<0.001
	Economic Aid	-0.2858***	(0.0177)	<0.001
	Production Aid	-1.1386***	(0.0608)	<0.001

Note: *p<0.1, **p<0.05, ***p<0.01.

Table 22: Callaway–Sant'Anna conditional, Cell Mining-Restricted

Level	Outcome	ATT	SE	p-value
Cell	World Bank Aid	0.0815	(0.4638)	0.860
	Social Aid	-0.1675	(0.4389)	0.703
	Economic Aid	0.3179	(0.2477)	0.199
	Production Aid	0.0146	(0.1108)	0.895

Note: *p<0.1, **p<0.05, ***p<0.01.

Table 23: Callaway–Sant'Anna conditional, Cell PSM-Matched

Level	Outcome	ATT	SE	p-value
Cell	World Bank Aid	-0.0634	(0.3615)	0.861
	Social Aid	-0.1001	(0.3264)	0.759
	Economic Aid	0.0787	(0.1841)	0.669
	Production Aid	-0.0129	(0.0731)	0.859

Note: *p<0.1, **p<0.05, ***p<0.01.

Table 24: China sub-analysis at the cell level

Estimator	ATT	SE	p-value

ETWFE Poisson (link scale)	-0.4700***	(0.1209)	<0.001
Callaway–Sant'Anna conditional (asinh)	0.1111*	(0.0655)	0.090

Note: Precision-filtered PRIO-GRID cells, 2000+ within-country PSM, outcome Chinese ODA-like commitments. The two estimators diverge in the same way as the other cell-level results (significantly negative under the Poisson estimator, marginally positive under the Callaway–Sant'Anna estimator), so the cell level offers no consistent evidence in either direction; the ADM2 design in Table 8 carries the Chinese verdict (see Section 5.2). *p<0.1, **p<0.05, ***p<0.01.

E. ASCM: pooled ATT and event study, all panels

Table 25: ASCM ATT, Mining-Restricted sample

Level	Outcome	ATT	SE	L2 imb.	Note
Adm2	Total Aid	-0.3511	(0.6796)	0.0000	
	DAC Aid	0.1962	(0.5190)	0.0001	
	World Bank Aid	-0.4614	(0.6781)	0.0000	
	Social Aid	0.0488	(0.7005)	0.0001	
	Economic Aid	-0.0076	(0.6349)	0.0000	
	Production Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration
Cell	World Bank Aid	0.5271	(0.3327)	0.0001	
	Social Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration
	Economic Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration
	Production Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration

Note: The outcome is the asinh-transformed aid measure; outcome names are shown untransformed. The pooled SE is the conservative root-mean-square of the per-period jackknife SEs. An L2 imbalance at or near zero is the desired outcome, not a defect: it means the augmented synthetic control reproduces the treated units' pre-treatment aid trajectory almost exactly, which is the property that lets ASCM stand in for the parallel-trends assumption; the near-zero pre-fit is shown directly in the event study (Figure 10).

Table 26: ASCM ATT, PSM-Matched sample

Level	Outcome	ATT	SE	L2 imb.	Note
Adm2	Total Aid	-0.4726	(0.5139)	0.0000	
	DAC Aid	-0.3703	(0.4197)	0.0001	
	World Bank Aid	-0.0766	(0.5140)	0.0000	
	Social Aid	-0.2789	(0.4983)	0.0000	
	Economic Aid	-0.4289	(0.4659)	0.0000	
	Production Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration
Cell	World Bank Aid	0.1219	(0.2429)	0.0000	
	Social Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration
	Economic Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration
	Production Aid	–	–	–	did not converge: near-degenerate (heavily zero-valued) outcome series; the optimisation is infeasible in every cascade configuration

Note: Outcome scale, pooled-SE construction and the reading of the L2 imbalance are as in Table 25.

Figure 10: ASCM event study for total aid: ADM2 Mining-Restricted (top) and ADM2 PSM-Matched (bottom)

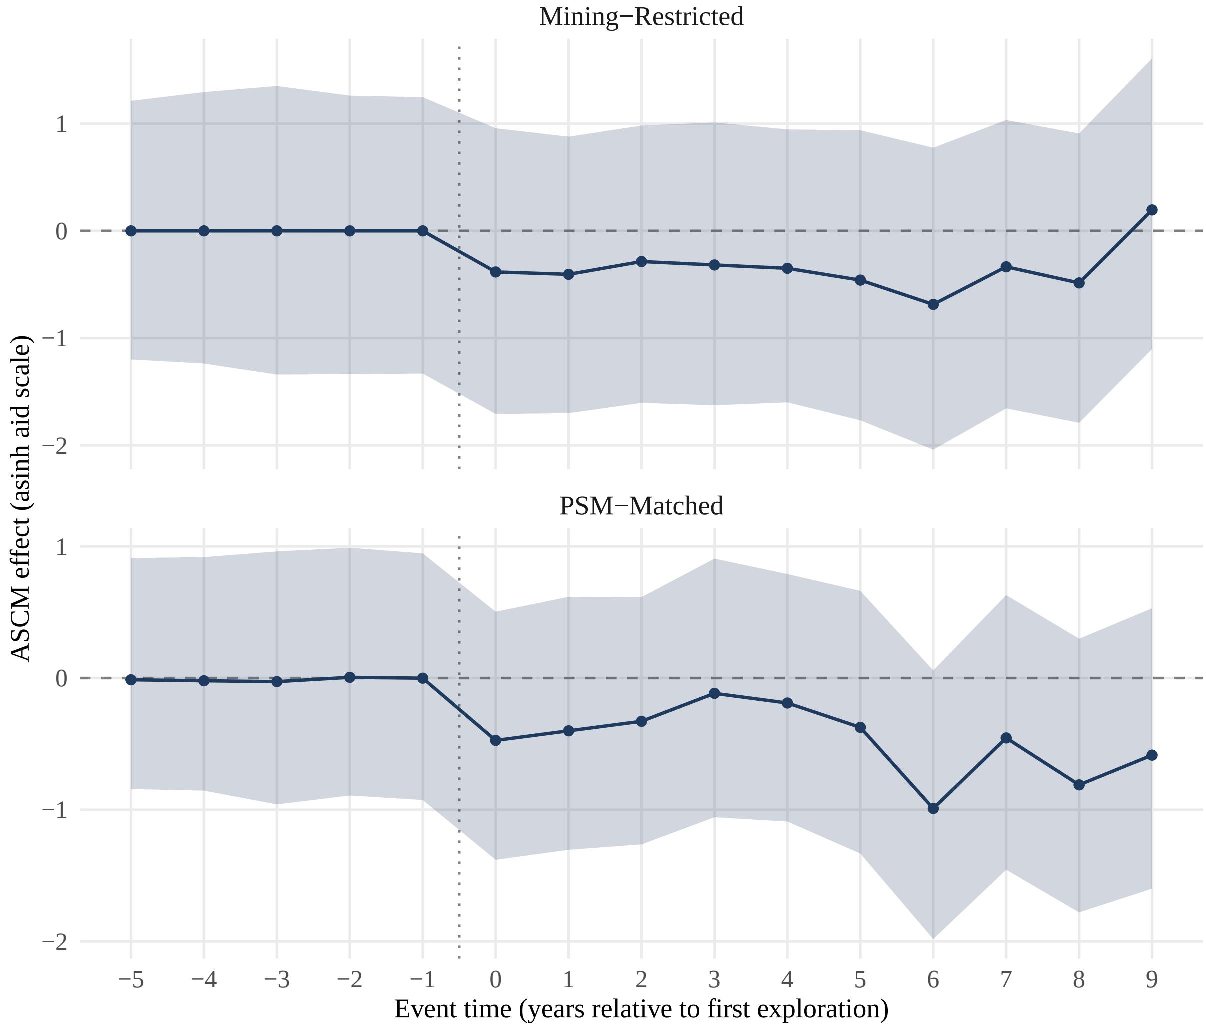

Average effect across treated cohorts. Event time 0 is the first exploration year. Negative event times are the five pre-treatment periods the synthetic control is balanced on; positive event times are the post-treatment effect. Shaded bands are 95% jackknife confidence intervals. The near-zero balancing window is the visual counterpart of the near-zero L2 imbalance in Tables 25 and 26; the post-treatment path is statistically indistinguishable from zero on both samples, and its average equals the pooled ASCM ATT in those tables. Estimated on the asinh aid scale, using the same configuration as the pooled ATT.

F. Estimation diagnostics

Figure 11: Distant-lead coverage diagnostic: 2000+ cohorts, ADM2 PSM-Matched, Total Aid

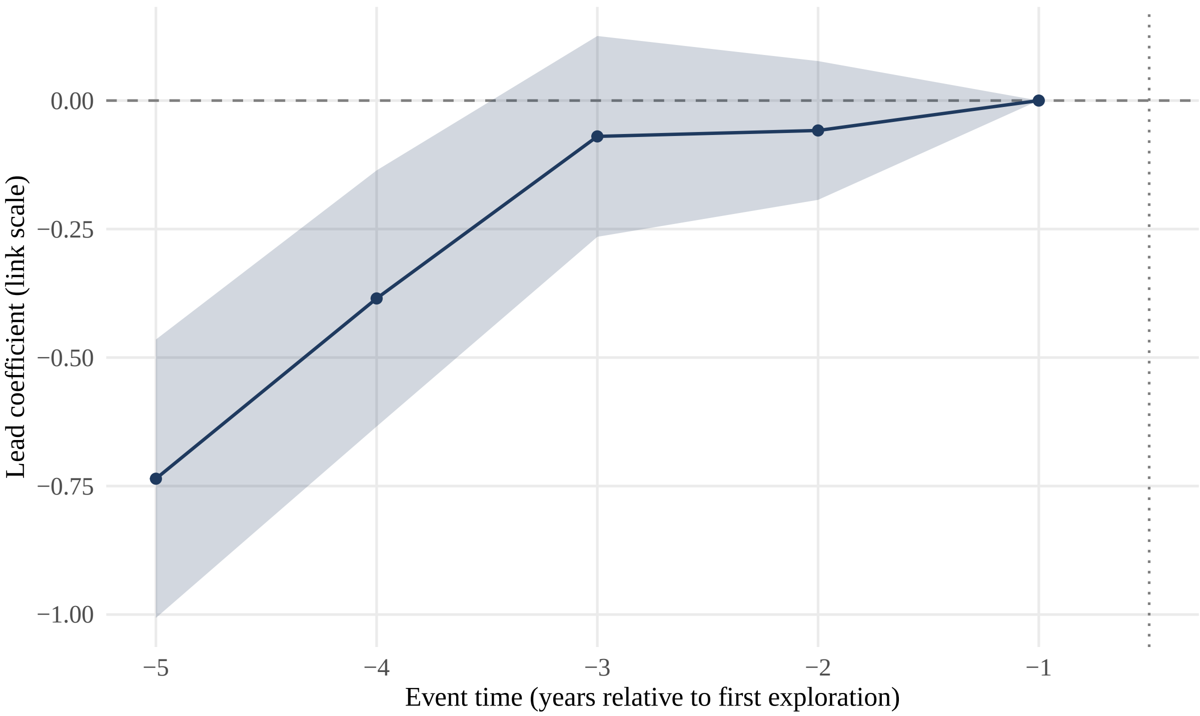

ETWFE Poisson event-study lead coefficients for the 2000+ cohorts (never-treated controls), whose pre-treatment years all lie in full-coverage calendar years. Link-scale estimates with 95% confidence intervals from province-clustered SEs; the reference period is $t = -1$. Compared with the all-cohorts leads in Figure 4, the distant leads roughly halve and the leads at -3 and -2 are statistically indistinguishable from zero.

Table 27: Pearson dispersion of headline Poisson fits

Level	Sample	Estimator	Outcome	dispersion
Adm2	Mining-Restricted	ETWFE-Poisson	Total Aid	14.1
Adm2	Mining-Restricted	ETWFE-Poisson	DAC Aid	7.9
Adm2	Mining-Restricted	ETWFE-Poisson	World Bank Aid	12.8
Adm2	Mining-Restricted	ETWFE-Poisson	Social Aid	10.0
Adm2	Mining-Restricted	ETWFE-Poisson	Economic Aid	13.3
Adm2	Mining-Restricted	ETWFE-Poisson	Production Aid	9.7
Cell	Mining-Restricted	ETWFE-Poisson	World Bank Aid	34.4
Cell	Mining-Restricted	ETWFE-Poisson	Social Aid	25.7
Cell	Mining-Restricted	ETWFE-Poisson	Economic Aid	19.9
Cell	Mining-Restricted	ETWFE-Poisson	Production Aid	9.3
Adm2	PSM	ETWFE-Poisson	Total Aid	19.7
Adm2	PSM	ETWFE-Poisson	DAC Aid	13.3
Adm2	PSM	ETWFE-Poisson	World Bank Aid	15.7
Adm2	PSM	ETWFE-Poisson	Social Aid	12.3
Adm2	PSM	ETWFE-Poisson	Economic Aid	14.3
Adm2	PSM	ETWFE-Poisson	Production Aid	6.6
Cell	PSM	ETWFE-Poisson	World Bank Aid	51.1

Cell	PSM	ETWFE-Poisson	Social Aid	41.0
Cell	PSM	ETWFE-Poisson	Economic Aid	24.2
Cell	PSM	ETWFE-Poisson	Production Aid	9.4
Adm2	Mining-Restricted	ETWFE-ES	Total Aid	12.9
Adm2	Mining-Restricted	ETWFE-ES	DAC Aid	6.4
Adm2	Mining-Restricted	ETWFE-ES	World Bank Aid	13.3
Adm2	Mining-Restricted	ETWFE-ES	Social Aid	11.5
Adm2	Mining-Restricted	ETWFE-ES	Economic Aid	12.4
Adm2	Mining-Restricted	ETWFE-ES	Production Aid	11.2
Adm2	PSM	ETWFE-ES	Total Aid	20.9
Adm2	PSM	ETWFE-ES	DAC Aid	14.6
Adm2	PSM	ETWFE-ES	World Bank Aid	16.7
Adm2	PSM	ETWFE-ES	Social Aid	14.5
Adm2	PSM	ETWFE-ES	Economic Aid	14.4
Adm2	PSM	ETWFE-ES	Production Aid	7.1

Table 28: SE diagnostics across specifications

Sample	Level	Outcome	% Zero	ATT (cluster)	SE (cluster)
Mining-Restricted	Adm2	Total Aid	59.7%	0.046	0.042
		Production Aid	82.4%	0.531	0.094
	Cell	World Bank Aid	92.6%	-0.381	0.036
		Production Aid	99.3%	-1.098	0.098
PSM	Adm2	Total Aid	60%	0.019	0.027
	Cell	World Bank Aid	92.2%	-0.140	0.025

Note: Outcome zero shares with the headline point estimate and its cluster-robust (province-level) standard error.