NET WORK LISER

Discussion Paper Series

IZA DP No. 18698

May 2026

The Economic Impact of the USAID Shutdown

Marcella Nicolini
University of Pavia

Fabio Sabatini
Sapienza University of Rome
and IZA@LISER

ISSN: 2365-9793

IZA DP No. 18698 May 2026

The Economic Impact of the USAID Shutdown

Abstract

On 28 January 2025 the second Trump administration issued a blanket stop-work order on the United States Agency for International Development (USAID), abruptly terminating the largest national bilateral aid programme in the world. We use this natural experiment to estimate the economic impact of the aid cut on two outcomes in Africa: local economic activity, measured through nighttime light radiance within fine spatial buffers around geocoded USAID project sites; and acute food insecurity, measured through the Integrated Food Security Phase Classification (IPC) at the subnational administrative level. We document three findings. First, the cessation of USAID activities produced a sharp and statistically significant decline in nighttime light radiance within 500 m to 10 km of project sites, with the magnitude attenuating monotonically and the effect no longer detectable at 25 km. Second, subnational areas with greater pre-shock exposure to USAID humanitarian assistance experienced relative increases in the share of population classified at IPC Phase 3 (Crisis) or worse and at Phase 4 (Emergency) or worse, with effects building up over the first year after the shock, amplified in regions with higher pre-shock vulnerability, and approximately fourteen times larger in countries with less democratic institutions. Third, both effects are driven by cuts to humanitarian aid (emergency response and food assistance); the nightlight effect is additionally driven by cuts to productive-sector aid (primarily agriculture).

JEL classification

F35, O12, O19, Q18, R11

Keywords

foreign aid, USAID, natural experiment, nighttime lights, food security, IPC, difference-in-differences, Africa

Corresponding author

Fabio Sabatini
fabio.sabatini@uniroma1.it

1 Introduction

On 28 January 2025, shortly after the inauguration of the second Trump administration, the Department of State issued a blanket stop-work order on all activities funded by the United States Agency for International Development (USAID), operationalising the ninety-day foreign-aid review mandated by an Executive Order of 20 January 2025 (The White House, 2025). USAID had operated continuously since 1961 as the principal civilian channel for U.S. bilateral foreign assistance, with total budgetary resources of approximately $45 billion in the fiscal year preceding the shutdown (USAID OIG, 2024). Roughly one-third of that flow went to Sub-Saharan Africa, supporting a portfolio that ranged from antiretroviral drug distribution to emergency food and cash assistance, and agricultural-development activities. The 28 January order terminated activities essentially overnight, with implementing partners notified by email that all work on existing contracts was to cease immediately and that no further disbursements would be made. By February 2025 substantially all USAID staff had been placed on administrative leave; by late February, the Department of Government Efficiency (DOGE) had announced the cancellation of approximately 83 percent of the agency's active programmes (Congressional Research Service, 2025).

This paper exploits the abrupt and exogenous character of the USAID shutdown to estimate its consequences for two outcomes in Africa: local economic activity, measured through satellite radiance of nighttime lights within fine spatial buffers around 1,374 USAID project locations across 50 African countries; and acute food insecurity, measured through the Integrated Food Security Phase Classification (IPC) for 388 subnational analytical units across 27 African countries. The USAID project locations are sourced from the Geocoded Official Development Assistance Dataset (GODAD) (Bomprezzi et al., 2024).

Identification rests on within-country variation in pre-shock USAID exposure (averaged over 2017–2020 disbursements recorded in GODAD, following the strategy of Rohner et al. (2026)) interacted with a post-shock indicator that turns on in February 2025. Our specifications include unit-level fixed effects (project sites for the nightlight outcome; subnational administrative cells for the food-insecurity outcome) that absorb time-invariant determinants of the outcomes, together with country and time fixed effects (interacted in the IPC specification) that absorb any country-wide variation orthogonal to local USAID exposure. The setting is essentially that of a continuous-treatment difference-in-differences design with a single, abrupt, exogenously timed shock; a quasi-experimental configuration whose identifying assumptions we discuss in detail and test through event studies, trend-break decompositions, sample restrictions, and a battery of placebo exercises.

The empirical results yield three central findings. First, the cessation of USAID activities produced a sharp post-shock decline in nighttime light radiance within 500

metres to 10 kilometres of USAID project sites. The magnitude of the effect attenuates monotonically with the buffer radius and falls below statistical detection at 25 kilometres. The effect takes the form of a discrete level shift at the date of the shock, with no detectable change in the pre-shock USAID–nightlight relationship, consistent with the interpretation that USAID had been sustaining a stock-like level of local economic activity (salaries paid to local staff, contractor operations, beneficiary cash and commodity transfers, and the operational footprint of implementing partners) whose immediate cessation produced a contemporaneous reduction in the visible economic footprint of the affected locations. The result is robust to the most conservative country-by-month fixed-effect structure, to sample restrictions including the exclusion of all active-conflict countries and the leave-one-country-out exercise across the fifty African countries in our sample, to a time-shifted placebo, and to the structural placebo provided by the spatial-decay pattern itself.

Second, in the 27 African countries where the Integrated Food Security Phase Classification (IPC) operates, regions historically exposed to USAID experienced relative increases in the share of population classified at Phase 4 (Emergency) or worse after the shutdown. The IPC is a multi-agency initiative that produces standardised classifications of acute food insecurity for the most vulnerable countries worldwide. The event study reveals an escalation pattern in which the post-shock coefficient at the second semester after the shock is larger than at the first. This is consistent with the operational dynamics through which the cessation of humanitarian assistance transmits to household food consumption, through the progressive depletion of in-country commodity stocks held by USAID implementing partners (Bureau for Humanitarian Assistance and Food for Peace operations) following the January 2025 stop-work order.

The effect is concentrated in countries with weaker political institutions — identified through a within-Africa median split of the V-Dem liberal democracy index — where the marginal impact is approximately fourteen times larger than in countries with stronger institutions, and in subnational units with higher pre-shock vulnerability, where the humanitarian-sector effect is essentially absent in low-vulnerability units while substantial in high-vulnerability units. These heterogeneity patterns support the interpretation that the USAID shutdown may push populations already at the margin of food crisis into emergency, while leaving more food-secure populations comparatively unaffected.

Third, both outcomes are driven by the humanitarian sector of the USAID portfolio. Decomposing pre-shock USAID exposure into four standard sectoral categories (social, economic, productive, humanitarian), we find that the nightlight and food-security effects both load on the humanitarian sector; the nightlight effect additionally loads on the productive sector, which operates at a spatial scale that the project-buffer measure captures directly but that the broader IPC averaging dilutes.

The paper contributes to three intersecting strands of the empirical literature on foreign aid and on the local consequences of external shocks. The most direct connection

is to Rohner et al. (2026), who use the same GODAD treatment-data construction to estimate the effect of the shock on violent conflict in Africa using Armed Conflict Location and Event Data (ACLED). They document an increase in conflict incidence in regions historically exposed to USAID, with the effect amplified in countries with weaker political institutions, and propose a theoretical mechanism in which the removal of aid-funded local economic activity reduces the opportunity cost of fighting and therefore increases conflict. Our analysis is complementary in several respects. Whereas Rohner et al. (2026) exploit ADM1-level conflict event data to test the consequences of the shock on violence, we use satellite-derived nightlights at fine spatial resolution and at higher temporal frequency to measure the first step in the proposed causal chain, the change in local economic activity itself, and document a magnitude and timing consistent with the level-shift interpretation that the opportunity-cost channel would predict. We also bring direct evidence on the food-security outcome that Rohner et al. (2026) discuss as motivating context, and we exploit the disaggregation of GODAD by sector to narrow the set of sectoral channels most plausibly responsible for the local effects.

A second strand of the literature is the long-standing body of empirical work on the local effects of foreign aid disbursements. The use of satellite nighttime lights as a proxy for subnational economic activity originates with Henderson et al. (2012) and has since been deployed across a wide range of contexts surveyed by Donaldson and Storeygard (2016). The local political-economy effects of geocoded aid are examined by Dreher et al. (2019a) on Chinese assistance, by Briggs (2017) on aid targeting, and by Isaksson and Kotsadam (2018) on the local corruption consequences of Chinese aid. The within-country political-economy literature is reviewed in Bommer et al. (2022). In recent work, Fazio et al. (2026) use the same GODAD geocoding infrastructure to analyse the relationship between local exposure to development aid projects and citizens' trust in domestic political institutions, documenting that exposure to implemented aid projects is associated with higher institutional trust, with the effect concentrated in socially oriented programmes and in less democratically consolidated countries.

A third strand of the literature is the body of work on the local consequences of foreign aid and humanitarian assistance in specific shock contexts. Nunn and Qian (2014) examine the effect of U.S. food aid on civil conflict using exogenous variation generated by U.S. domestic wheat-production cycles. Crost et al. (2014) document the effects of community-driven development programmes on conflict in the Philippines. Maystadt and Ecker (2014) examine the effects of weather-induced shocks on conflict through livestock-price channels in Somalia. Andrabi and Das (2017) show that U.S. humanitarian assistance after the 2005 Pakistan earthquake durably improved local trust in foreign donors. Dreher et al. (2019b) document the effects of foreign aid on refugee flows. Langlotz (2026) examines the effect of foreign interventions on community cohesion in conflict settings. Our paper differs from this literature in two respects. First, we exploit a sharper exogenous shock, an outright

shutdown rather than year-to-year variation in disbursements or natural-disaster timing. Second, we capture the post-shock effects through two methodologically independent measurement infrastructures automated satellite radiance for local economic activity and IPC classifications produced by multi-agency expert panels for acute food insecurity, thereby moving beyond designs that rely on a single measurement infrastructure.

Our empirical evidence on the early post-shock welfare consequences also complements the projection exercise of Cavalcanti et al. (2025), who extrapolate pre-shock programme-impact estimates to forecast medium-term mortality consequences of the shutdown.

The remainder of the paper is organised as follows. Section 2 describes the institutional context of USAID's African operations, the chronology and implementation of the January 2025 shutdown, and the reasons we treat the shock as plausibly exogenous to recipient-country conditions. Section 3 documents the construction of the datasets. Section 4 presents the empirical strategy. Section 5 reports the main results on local economic activity around USAID project sites. Section 6 reports the results on acute food insecurity. Section 7 discusses the convergent evidence on sectoral channels. Section 8 concludes.

2 Background and setting

In recent years USAID has disbursed approximately eight billion U.S. dollars per year to Sub-Saharan Africa (Congressional Research Service, 2023), the agency's largest regional concentration. By the mid-2020s the agency had built deep operational ties with national governments, local non-governmental organisations, international implementing partners, and beneficiary communities across the continent. A typical USAID project location housed contracted staff (often a mix of expatriate and local personnel), maintained sustained relationships with national ministries of health or agriculture, and delivered services to identifiable beneficiary populations. The local economic footprint of these activities, i.e., the salaries paid to local staff, the procurement of goods and services from local markets, the transfers received by beneficiary households, and the secondary economic flows generated by the presence of an institutional anchor, was substantial and, in some specific regions, dominant.

Aggregating the project-level data recorded in the Geocoded Official Development Assistance Dataset (GODAD) over the 2017–2020 window we use to construct pre-shock exposure, total USAID disbursements to Africa amounted to approximately $4.5 billion in constant 2014 U.S. dollars, distributed across 1,374 geocoded project locations in 50 of the 55 African countries.[1] Each project in GODAD carries the OECD Creditor

[1]This is the location-attributable total, which evenly attributes each project's disbursement across its geographic locations to avoid double-counting in our spatial analysis. The corresponding unsplit project-level total (directly comparable to OECD/CRS reporting) is approximately 10.0 billion, or 2.5 billion annually; we discuss the relationship to broader U.S. assistance flows in Section 3.1.

Reporting System (CRS) purpose code assigned by the donor at the time of disbursement. We classify each project location into one of four OECD-DAC standard macro-sectors on the basis of the first digit of its three-digit CRS purpose code: *social* (CRS codes 1xx, including health, education, governance, water and sanitation, and PEPFAR-funded HIV/AIDS programmes) accounts for approximately 52.9% of disbursements in our sample; *economic* (CRS codes 2xx, including transport, communications, energy, banking and financial services) accounts for approximately 4.4%; *productive* (CRS codes 3xx, including agriculture, forestry, fishing, industry, mining, and trade, encompassing Feed the Future agricultural-development activities) accounts for approximately 6.9%; and *humanitarian* (CRS codes 5xx and 7xx, including emergency response, food assistance, disaster prevention, and the Bureau for Humanitarian Assistance (BHA) and Food for Peace operations) accounts for approximately 34.4%. A small residual share (1.8%) corresponds to multi-sector and unallocated activities (CRS codes 4xx and 9xx).[2]

The second Trump administration took office on 20 January 2025. On 24 January, the U.S. Department of State issued an internal directive freezing all foreign-assistance disbursements for a ninety-day review, with limited exceptions for life-saving humanitarian assistance and military aid to specific partner countries (The White House, 2025). On 28 January the directive was operationalised as a blanket stop-work order applied to all USAID-funded activities worldwide. Implementing partners received notification, typically by email, that all work on existing USAID-funded contracts and cooperative agreements was to cease immediately, that no further disbursements would be made, and that staff should prepare for the discontinuation of activities. The order applied retroactively: activities already in progress and partially completed were treated as terminated as of the order date, with no provision for orderly close-out.

In the weeks that followed, USAID's operational infrastructure was dismantled at a pace without precedent in U.S. foreign-assistance history. Substantially all USAID staff, both in Washington and in overseas missions, were placed on administrative leave; several thousand foreign-service officers were given short notice to return to the United States, with formal reduction-in-force notices issued on 28 March 2025 setting full separation by 1 July 2025 (Congressional Research Service, 2025). The agency's public-facing websites went offline; the implementing-partner contracts database was archived and removed from public access. By mid-March the Department of State announced the cancellation of approximately 83% of USAID-managed programmes, with the remaining 17% transferred

[2]GODAD also distributes pre-aggregated sectoral variables (*disb_soc_split*, *disb_eco_split*, *disb_prod_split*, *disb_other_split*) that follow a project-specific aggregation of CRS purpose codes which does not coincide one-to-one with the OECD-DAC macro-categories described above. For transparency and direct comparability with the international aid literature that uses the OECD-DAC standard, we recompute the four sectoral exposures directly from the three-digit CRS purpose codes recorded for each project record in GODAD, applying the canonical first-digit mapping just described. Appendix A.3 (Table A2) reports the joint sectoral regressions for both outcomes side by side under both classifications.

under Department of State oversight (Congressional Research Service, 2025). Implementing partners, many of which had hired thousands of local staff and built supply chains around multi-year USAID contracts, were given no opportunity to plan for transition, to identify alternative donors, or to wind down activities in an orderly manner. Local salaries went unpaid; in-progress shipments of food and medical commodities were halted at ports of entry; clinics dependent on PEPFAR-funded procurement of antiretroviral drugs faced sudden stockouts; agricultural extension services delivered through Feed the Future were abruptly discontinued mid-growing-season.

The policy decision was motivated by U.S. domestic political dynamics: a perceived misalignment of USAID's programmatic priorities with the administration's foreign-policy objectives, a broader ideological commitment to reducing the federal civilian workforce, and specific objections to the agency's involvement in family-planning and gender-related programmes (The White House, 2025). None of these motivations was plausibly responsive to contemporaneous African economic or food-security conditions.

The shutdown was implemented uniformly across all USAID-funded activities in Africa, irrespective of country, sector, or beneficiary. We are not aware of any substantial geographic or sectoral exemptions to the stop-work order. This uniformity matters for our continuous difference-in-differences identification: the absence of selectivity in the shock means that the differential post-shock evolution of outcomes can be attributed to pre-shock variation in exposure intensity rather than to differential rates of cessation across treated locations.

The timing of the shutdown was unanticipated by both USAID staff and recipients. Implementing partners described the receipt of the stop-work order as a sudden event for which they had received no prior signal, leaving them without time to plan for the transition, identify alternative donors, or wind down activities in an orderly manner (Congressional Research Service, 2025). Forward-looking adjustment of activities in expectation of the shock would have eroded the contrast between pre- and post-shock outcomes; we find no such evidence in the data. Our event studies in Sections 5 and 6 show that the relationship between historical USAID exposure and outcomes was stable through the months immediately preceding January 2025, with the empirical break visible only after the shock date.

Taken together, the U.S.-domestic political origin of the decision, the operational uniformity of its implementation, and its unanticipated character make the shutdown plausibly exogenous to local economic and food-security conditions in the affected African regions.

Two qualifications are worth flagging. First, the formal classification of programmes as "cancelled" versus "transferred to State Department oversight" was carried out, in subsequent weeks, on a programme-by-programme basis. The resulting partial reactivation

of approximately 17 % of activities may have introduced some heterogeneity in actual treatment intensity. Whether this partial reactivation attenuates or amplifies our estimates depends on its correlation with historical USAID exposure: if reactivation is concentrated at high-exposure locations (for instance, PEPFAR-funded activities preserved by emergency waiver), our estimates are attenuated; if concentrated at low-exposure locations, they are amplified. We have no direct evidence on this joint distribution, but the stability of the headline coefficient across the leave-one-country-out and sample-restriction exercises in Section 5.6 argues against selective reactivation being a material driver of our results. Second, secondary cuts to U.S. contributions to multilateral institutions, notably to International Bank for Reconstruction and Development (IBRD) capital, to International Development Association (IDA) replenishments, and to the World Health Organization's regular budget, were announced in the months after the USAID shutdown but operated through different fiscal mechanisms and on different timelines. We discuss the implications of these secondary cuts for the interpretation of our placebo tests in Section 5.7.

3 Data

This section documents the construction of the analytical datasets we use throughout the paper. We describe in turn the treatment data (Section 3.1), the two outcome datasets (Sections 3.2 and 3.3), the institutional moderators (Section 3.4), and the summary statistics (Section 3.5).

3.1 USAID exposure

We measure USAID exposure using the Geocoded Official Development Assistance Dataset (GODAD), version 1.0, released in July 2025 (Bomprezzi et al., 2024). GODAD combines data from the OECD's Creditor Reporting System (CRS), AidData's project-level geocoding, the World Bank's IATI feeds, and donor-specific datasets to produce a comprehensive geocoded panel of aid flows from 23 donors over the 1973–2023 period (the U.S. coverage being effectively complete only through 2020 because of the geocoding lag for recent CRS submissions). For each project, GODAD records the latitude and longitude of the project location, a sector code from the OECD CRS taxonomy, the donor agency, the year of payment, the disbursement amount in nominal and constant 2014 U.S. dollars, and a variety of auxiliary identifiers. We restrict the GODAD project-level data to (i) donor equal to "United States", (ii) project locations within Africa as defined by the Global Administrative Areas database (GADM, version 3.6), the standard reference for subnational administrative boundaries in development economics, and (iii) payment years between 2017 and 2020 inclusive. The restriction to 2017–2020 follows Rohner et al. (2026) and reflects the fact that the geocoded version of the CRS data is effectively complete only through

2020. After applying these filters, we are left with 8,996 USAID project-location-year records corresponding to 1,374 unique geographic points (latitude × longitude rounded to four decimals). Of these 1,374 points, 1,223 received non-zero disbursement during the four-year window; the remaining 151 points correspond to USAID-affiliated locations that were registered in GODAD but did not have disbursements during the window.

Figure 1 shows the geographic distribution of the 1,374 USAID project locations across Africa, with the color intensity of each point reflecting the log of cumulative disbursement during 2017–2020. The geographic concentration is heavily skewed toward Sub-Saharan low-income and conflict-affected regions, with North Africa and parts of Southern Africa visibly under-represented.

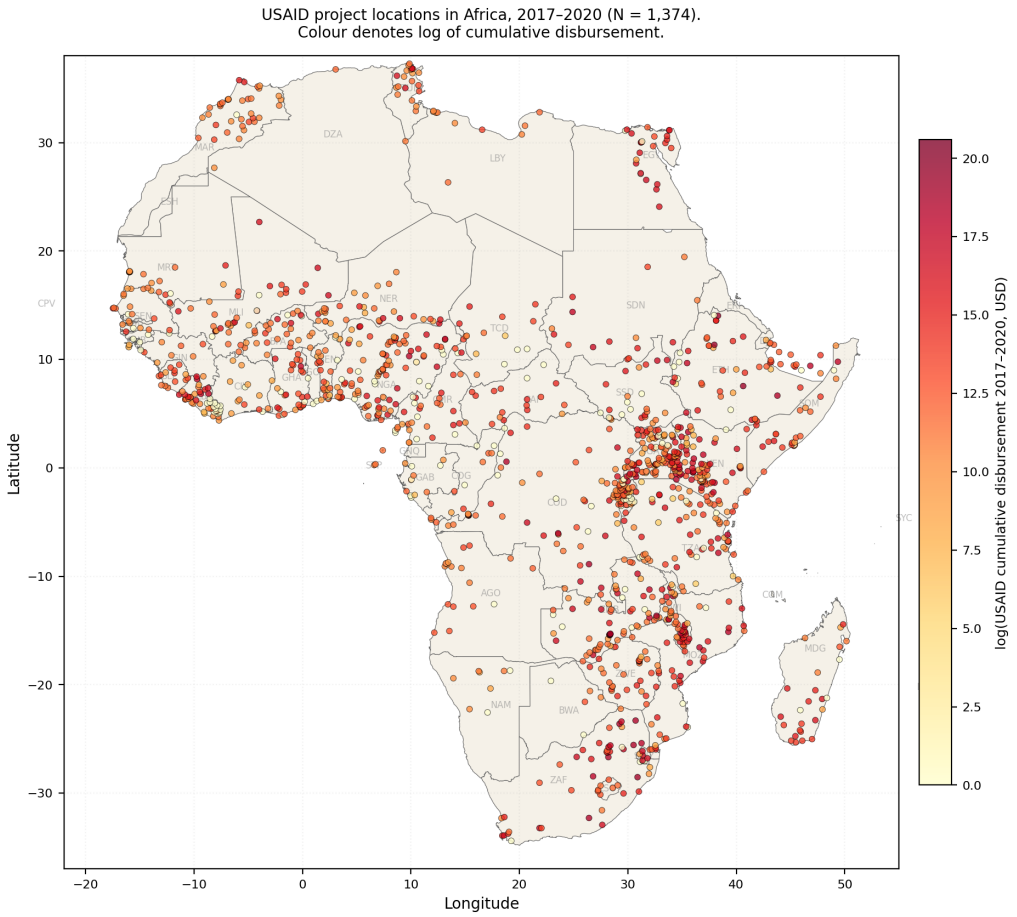

Figure 1. Geographic distribution of USAID project locations in Africa, 2017–2020
Notes. Each point is one of the 1,374 unique USAID project locations recorded in the GODAD project-level dataset with non-zero disbursement during the 2017–2020 window. Point colour denotes the log of cumulative disbursement at the location, in constant 2014 U.S. dollars. Mauritius is not visible at this map projection because of its small geographic footprint relative to the continent. *Source:* authors' calculations from GODAD v1.0.

For each unique project location we compute the total USAID disbursement during 2017–2020, in constant 2014 U.S. dollars. The distribution is heavily right-skewed: the median total disbursement at a project location is approximately $119,000, while the 90th percentile is approximately $15.07 million. Aggregated to the country level, GODAD-

tracked USAID disbursements amount to approximately $2.5 billion annually to Africa during the 2017–2020 window (in constant 2014 U.S. dollars), against an approximate $8 billion in total U.S. assistance to Sub-Saharan Africa per year reported by Congressional Research Service (2023). The gap reflects the fact that GODAD records only transactions that can be geocoded to a specific project location, excluding budget support, multi-country programmes, and contributions to multilateral channels that cannot be attributed to a single coordinate; the spatial granularity that this selection provides is, however, essential for the buffer-level identification strategy we adopt. Figure 2 reports the full distribution of log-transformed disbursements and the country-level totals for the top fifteen recipient countries.

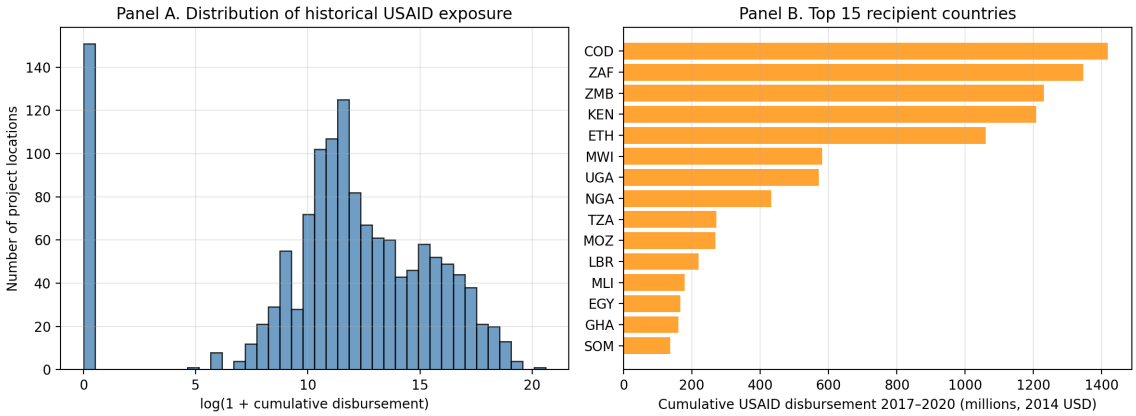

Figure 2. Distribution of USAID exposure, 2017–2020

Notes. Panel A: histogram of $\log(1 + \overline{D}_i^{2017–2020})$ over the 1,374 USAID project locations. Panel B: cumulative USAID disbursement to Africa during 2017–2020, in millions of constant 2014 U.S. dollars, by recipient country. The top 15 countries are shown. *Source:* authors' calculations from GODAD v1.0.

For each location we also construct sector-specific disbursement totals over 2017–2020, using the four OECD-DAC macro-sectors introduced in Section 2 and applying the OECD CRS purpose codes as recorded in GODAD.

In addition to the project-level data, GODAD provides aggregated panels at the ADM1 of GADM 3.6 boundaries, with disbursement and commitment columns for every (unit × year) combination over 1989–2023. We use the ADM1 aggregate to construct treatment variables for the subnational-administrative consistency check (Section 5).

3.2 Nighttime light radiance

We obtain nighttime light radiance from the Visible Infrared Imaging Radiometer Suite (VIIRS) sensor aboard the Suomi-NPP and NOAA-20 satellites. The VIIRS Day/Night Band detects very low light intensities at a native pixel resolution of approximately 463 m and a daily revisit time. We use the monthly stray-light-corrected composites (Elvidge et al., 2017) produced by the Earth Observation Group at the Colorado School of Mines

and distributed through the Google Earth Engine Data Catalog. The stray-light correction removes residual contamination from off-axis lunar light, which is important for the analysis of low-radiance pixels that dominate much of rural Africa.

We extract VIIRS data from January 2017 through April 2026, yielding 112 monthly composites. The temporal coverage spans 97 pre-shock months (January 2017 through January 2025) and 15 post-shock months (February 2025 through April 2026). We retain the `avg_rad` band, which records the average radiance per pixel in nanoWatts per steradian per square centimetre, after dropping pixels with negative radiance values and applying a gas-flare mask: pixels with avg_rad > 300 nW sr^{-1} cm^{-2} are masked out, following the convention in the literature, because such extreme radiance values almost always correspond to oil and gas flares (which dominate the VIIRS signal in certain locations such as the Niger Delta in Nigeria, the gas fields of eastern Algeria, and parts of South Sudan) rather than to human settlement lights and would otherwise contaminate the estimation.

We construct two VIIRS panels, corresponding to the two units of analysis used in the paper.

Project-buffer panel. For each of the 1,374 unique USAID project locations identified in the GODAD project-level extract, we construct circular buffers of five radii: 500 m, 1 km, 5 km, 10 km, and 25 km. For each (project location $\times$ buffer radius $\times$ month) combination, we use Earth Engine's `reduceRegions` operation to compute four summary statistics from the monthly VIIRS composite: the mean radiance per pixel within the buffer (avg_rad), the total radiance summed across all pixels in the buffer (sum_rad), the count of pixels above a luminosity threshold of 0.3 nW sr^{-1} cm^{-2} (lit_count), and the total count of non-flare-masked pixels (total_count). The threshold of 0.3 for the lit-pixel indicator follows Henderson et al. (2012) and is standard in the literature on extensive-margin measures of nighttime light. The five-radius buffer panel contains $1{,}374 \times 5 \times 112 = 769{,}440$ records.

ADM1 panel. As a complementary aggregate-level dataset, we also extract VIIRS statistics at the first-level administrative boundaries of the FAO Global Administrative Unit Layers (GAUL) 2015 simplified-500m feature collection, restricted to the 55 African countries listed in Appendix A.1 (Table A1). The same gas-flare mask and the same four summary statistics are computed at the GAUL ADM1 level. The resulting raw panel contains 859 polygons $\times$ 112 months = 96,208 records. We match GAUL ADM1 units to the GADM 3.6 boundaries used by GODAD via a name-based crosswalk that combines exact, substring, and fuzzy string matching with manual overrides (details in Appendix A.2); the procedure achieves an 83.0% match rate (713 of 859 polygons). The unmatched units are concentrated in small island states with no measurable USAID

exposure during 2017–2020 (Cape Verde, the Comoros, the Seychelles, São Tomé and Príncipe, Mauritius), where GODAD records either zero or trivial disbursements, and in Madagascar and Kenya, for which the structural mismatch between FAO GAUL boundaries and GADM 3.6 (driven by Madagascar's 2009 administrative reorganisation and Kenya's 2013 devolution into 47 counties) prevents matching at the ADM1 level. We emphasise that this restriction applies only to the ADM1 consistency check reported in Section 5.5: the main buffer-level analysis uses the project-location coordinates directly and therefore includes all USAID project sites in Kenya, Madagascar, and the small island states with non-zero USAID activity. The final ADM1 analysis panel comprises 79,856 records covering 672 matched ADM1 units across 53 African countries.

3.3 Food insecurity

We obtain acute food insecurity classifications from the Integrated Food Security Phase Classification (IPC), accessed through the Humanitarian Data Exchange repository maintained by the Joint Research Centre of the European Commission. The IPC releases "Global Acute Food Insecurity Country Data" covering approximately 30 countries with an active IPC Technical Working Group, at three nested levels of aggregation: country, subnational "Level 1", and more granular "Area". Our primary outcome analyses use the Level 1 data, the unit at which IPC classifications can be meaningfully linked to the GADM 3.6 administrative boundaries on which the GODAD treatment data are organised.

An IPC classification is the consensus output of a Technical Working Group, which examines multiple sources of evidence simultaneously, i.e., food consumption scores from WFP mobile vulnerability and analysis monitoring, acute malnutrition data from SMART surveys, mortality estimates where available, market prices, agricultural production data, displacement and conflict information, and assigns each IPC area to a phase using a standardised protocol of cross-referenced thresholds. The integration of multiple data sources through expert judgement, rather than a single mechanically computed statistic, provides robustness to noise in any single source, including noise that itself may be partly attributable to the USAID shutdown (since some of the underlying data are collected by partners funded by USAID). We address this potential endogeneity-of-measurement concern in Section 4.6, where we report direct diagnostics of IPC operational continuity over the post-shock period, and partially mitigate it by focusing on categorical IPC thresholds rather than continuous food consumption indicators.

The IPC Level 1 file contains 53,872 records covering analyses conducted from January 2017 through January 2026, with subsequent projection periods extending to April 2026. For each (country × analysis date × Level 1 unit × validity period) combination, the file reports the share of the unit's population classified in each IPC phase from 1 to 5, the cumulative share in Phase 3 or higher (which we denote "Phase 3+"), and the

total population of the unit. Each analysis records three validity periods: a *current* period covering the months at and around the analysis date (typically 2–3 months), a *first projection* period covering the subsequent 4–6 months (the working group's best forecast at the time of analysis), and a *second projection* period covering 7–12 months ahead. For our baseline analyses we use the *current* period only, as this represents the working group's contemporaneous assessment of food-security conditions rather than a forward-looking forecast.

Pivoting the Level 1 file by phase and restricting to current-period observations in African countries yields an initial panel of 3,056 IPC unit × analysis-date records spanning 29 countries from June 2017 to January 2026. We then drop 42 records corresponding to non-geographic IPC categories that some country working groups publish alongside the geographic units (most prominently in Mozambique, where the working group separately classifies internally displaced persons ("Deslocados") and host communities ("Acolhedores")). The resulting panel contains 3,014 IPC unit × analysis-date observations. We match IPC Level 1 units to the GADM 3.6 boundaries used by GODAD via a name-based crosswalk that combines exact matching, manual overrides for language and spelling differences, and fuzzy string matching (details in Appendix A.2); the procedure achieves a 95.3% match rate (388 of 407 IPC units in non-excluded countries). Madagascar and Kenya are excluded entirely from the analytical sample because the structural mismatch between their post-2009 (Madagascar) and post-2013 (Kenya) subnational reorganisations and the GADM 3.6 boundaries does not permit meaningful crosswalking. Our final analytical IPC sample comprises 2,873 observations covering 388 IPC units in 27 African countries, of which 19 contribute observations both before and after the January 2025 shock.

For each observation, the headline outcomes are the share of population classified in Phase 3+ ("Crisis or worse") and in Phase 4+ ("Emergency or worse"). We use both indicators throughout the paper, but our main results focus on Phase 4+ as the operationally most consequential margin and as the indicator on which our estimates are most robust.

3.4 Institutional environment

The moderation analyses use country-level institutional indicators averaged over the 2017–2020 pre-shock window, drawing on two complementary sources.

The Polity 5 dataset (Marshall and Gurr, 2020), distributed by the Center for Systemic Peace, provides annual measures of political-regime characteristics for 167 countries from 1800 through 2018, with partial coverage through 2020. From it we extract the *polity2* score (ranging from −10 for full autocracy to +10 for full democracy), the *xconst* variable measuring constraints on the chief executive, and the *durable* variable measuring regime

longevity, averaging each over the 2017–2020 window; after excluding small countries not in Polity (Cape Verde, the Comoros, São Tomé and Príncipe, the Seychelles), we have institutional measures for 51 of the 55 African countries in our sample.

The Varieties of Democracy v15 dataset (Coppedge et al., 2025), released in 2025, provides more granular and continuous measures of democratic institutions for 202 countries from 1789 to 2024; we extract seven high-level indices from the Core Country-Year file, i.e., electoral democracy (*v2x_polyarchy*), liberal democracy (*v2x_libdem*), participatory democracy (*v2x_partipdem*), deliberative democracy (*v2x_delibdem*), egalitarian democracy (*v2x_egaldem*), judicial constraints (*v2x_jucon*), and legislative constraints (*v2xlg_legcon*), averaging each over 2017–2020 and obtaining values for all 54 African countries in our sample.

In our headline heterogeneity analyses we use *v2x_libdem* as the primary V-Dem moderator, because it captures the broad notion of constraints on majority rule that the institutional-mitigation hypothesis invokes; we report *v2x_jucon* as a robustness check and the Polity 5 *polity2* score for direct comparability with Rohner et al. (2026).

3.5 Summary statistics

Table 1 reports summary statistics for the key variables in the three analytical panels constructed above: the project-buffer panel (769,440 observations covering 1,374 USAID project locations × five buffer radii × 112 monthly VIIRS composites), the IPC panel (2,873 observations covering 388 IPC units in 27 African countries), and the ADM1 panel used for the consistency check (79,856 observations covering 672 GADM 3.6 ADM1 units).

Three features of the summary statistics warrant attention. First, the distribution of USAID disbursements is heavily right-skewed: a small number of highly-exposed project locations (capital-city offices, major PEPFAR sites, large food-aid distribution centres) receive disbursements orders of magnitude larger than the majority of locations in the sample, as visible in Panel A from the contrast between the mean ($7.3 million) and the 75th percentile ($1.5 million). The logarithmic transformation we apply in equation (1) addresses this skew.

Second, the IPC outcome variables are themselves right-skewed: the 25th and 75th percentiles of the Phase 3+ population share are 6% and 24%, while for Phase 4+ they are 0% and 4% respectively; the maximum values reach 81% (Phase 3+) and 49% (Phase 4+). The Phase 4+ outcome in particular is zero for a substantial part of the sample, with only the most acutely food-insecure units exhibiting non-zero values.

Third, the institutional moderators show wide variation across African countries: the V-Dem *v2x_libdem* index ranges from below 0.05 in the most autocratic countries to above 0.6 in the most democratic, with the 25th and 75th percentiles in the African sample at 0.13 and 0.40 (Panel D), providing ample variation for moderation analyses.

Table 1. Summary statistics

	N	Mean	SD	p25	p75
Panel A: Project-buffer VIIRS panel (1 km radius)					
avg_rad (nW sr^{-1} cm^{-2})	153,888	6.56	16.07	0.27	4.24
log(avg_rad + 0.01)	153,888	0.13	1.65	−1.13	1.45
$\log(1 + \overline{D}_i^{2017\text{-}2020})$	1,374	12.66	4.93	9.02	16.52
Disbursement, USD (2014 constant)	1,374	7,273,049	24,102,415	0	1,506,618
Buffer is "urban" (radiance > median)	1,374	0.50	0.50		
Panel B: IPC analytical panel					
% pop. in Phase 3+ (current period)	2,873	0.177	0.152	0.060	0.240
% pop. in Phase 4+ (current period)	2,873	0.031	0.053	0.000	0.040
$\log(1 + \overline{D}_i^{2017\text{-}2020})$	388	7.35	6.18	0.00	12.78
USAID disb., USD (2014 constant)	388	1,340,070	7,416,350	0	355,235
Panel C: ADM1 panel (GAUL→GADM 3.6 matched units)					
avg_rad (nW sr^{-1} cm^{-2})	79,856	0.45	1.92	0.05	0.27
log(avg_rad + 0.01)	79,856	−1.69	1.32	−2.98	−1.27
$\log(1 + \overline{D}_i^{2017\text{-}2020})$	672	9.84	6.40	4.92	14.94
Panel D: Institutional moderators (country-level, 2017–2020 averages)					
Polity 5 *polity2*	51	1.7	4.9	−2.0	6.0
V-Dem *v2x_libdem*	54	0.27	0.16	0.13	0.40
V-Dem *v2x_polyarchy*	54	0.40	0.19	0.25	0.55
V-Dem *v2x_jucon*	54	0.47	0.22	0.27	0.66

Notes. Panel A reports statistics for the 1 km buffer panel; panels for the other four buffer radii are reported in Appendix Table A3. Panel B is the IPC analytical sample after applying the IPC↔GADM 3.6 crosswalk. Panel C is the ADM1 panel after applying the GAUL→GADM crosswalk. Panel D reports country-level averages of institutional moderators over 2017–2020. "% pop. in Phase k+" is the share of the IPC unit's population classified in IPC Phase k or higher during the current period of each analysis. USAID disbursements are reported in constant 2014 U.S. dollars; mean disbursement values reflect the strong right-skew of the underlying distribution. *Sources:* authors' calculations from GODAD v1.0, VIIRS DNB monthly composites VCMSLCFG, IPC Acute Food Insecurity Country Data via HDX, Polity 5, V-Dem v15, FAO GAUL 2015 simplified.

4 Empirical strategy

This section develops the empirical framework that we use to estimate the effect of the USAID shutdown of January 2025. We organise the discussion in eight subsections. Section 4.1 states the identification framework in general terms and explains why the abrupt nature of the policy intervention permits a difference-in-differences approach. Section 4.2 defines the continuous treatment variable. Section 4.3 describes the two outcome variables and the units of observation. Section 4.4 writes down the baseline regression equation, motivates the asymmetric fixed-effect structures that we use across the two outcomes, and describes the clustering of standard errors. Section 4.5 is dedicated to the parallel-trends assumption. Section 4.7 describes the event study, Section 4.8 introduces the trend-break specification, and Section 4.10 previews the robustness exercises that we report in Section 5.6 and Section 5.7.

4.1 Identification framework

We exploit the abrupt and arguably unanticipated nature of the USAID shutdown to construct a continuous difference-in-differences design. The argument for the exogeneity of the policy shock rests on three observations. First, the blanket stop-work order of 28 January 2025 was the result of a political process internal to the U.S. executive branch, motivated by domestic ideological considerations of the second Trump administration rather than by any contemporaneous change in conditions in the recipient regions; it is therefore plausibly orthogonal to local economic and food-security developments at the subnational level in Africa. Second, the announcement was unanticipated even by USAID staff, contractors, and recipients: programmes that had been operating for decades were terminated within a matter of days, with no period of phasedown that would have allowed local actors to substitute or to anticipate the shock. Third, the shock was applied uniformly to all USAID-funded activities, a near-total cessation of programmes, with no substantial geographic or sectoral exemptions. This comprehensiveness eliminates the possibility of selective implementation that could have generated endogenous variation in treatment intensity: differential post-shock outcomes can therefore be attributed to pre-shock variation in exposure rather than to differential rates of cessation across affected locations.

Given these features, we treat 28 January 2025 as a sharp policy break that approximates an exogenous, simultaneous shock to all regions previously exposed to USAID. Our continuous DiD design uses subnational variation in the intensity of pre-shock USAID exposure to identify the differential post-shock evolution of outcomes.

The implementation of this design differs across our two outcomes because the units of analysis differ: circular buffers around USAID project locations for local economic activity, and IPC Level 1 areas (approximately ADM1) for food insecurity. We describe both in Section 4.3.

4.2 Treatment variable: pre-shock USAID exposure

We measure pre-shock USAID exposure following Rohner et al. (2026), using the Geocoded Official Development Assistance Dataset (GODAD) (Bomprezzi et al., 2024). For each unit of analysis i, either a project location or an ADM1 region, we define the historical exposure

$$\text{USAID}_i \;=\; \log\!\left(1 + \overline{D}_i^{2017\text{–}2020}\right), \tag{1}$$

where $\overline{D}_i^{2017\text{–}2020}$ is the average annual U.S. disbursement to unit i during the four-year window 2017–2020, measured in constant 2014 U.S. dollars. The logarithmic transformation accommodates the strong right-skew of the disbursement distribution and the substantial

mass at zero. We use the four-year average rather than a single year's disbursement to smooth out idiosyncratic timing in project funding cycles.

Two features of the treatment variable warrant emphasis. First, the temporal window is constrained by data availability: the geocoded version of the OECD's Creditor Reporting System, on which GODAD relies for U.S. data, is fully geocoded only through 2020. Following Rohner et al. (2026), we therefore implicitly assume that the spatial distribution of U.S. aid programmes was sufficiently stable between 2020 and the time of the shock in January 2025 that the 2017–2020 average is a good proxy for U.S. exposure on the eve of the shutdown. This assumption is consistent with the institutional reality of USAID activities, which were typically organised around multi-year country compacts (PEPFAR five-year plans, Feed the Future strategy cycles, BHA framework agreements with implementing partners) and which exhibit considerable continuity over time at the level of recipient locations. Second, the treatment intensity varies continuously across units rather than being binary. We discuss how this affects the interpretation of our coefficients in Section 4.4.

4.3 Outcomes and units of analysis

Our first outcome is satellite-measured nighttime light radiance. The visible light emitted by a location at night reflects the underlying density of electrified infrastructure, commercial activity, residential consumption of electricity, and labour-related activity (Henderson et al., 2012). Although nightlights are an imperfect measure of economic activity as they capture some forms of activity (urbanisation, formal sector commerce, electrified housing) better than others (subsistence agriculture, informal markets, household-level consumption), they have the considerable advantages of (i) global coverage at high spatial and temporal resolution, (ii) being available essentially in real time, and (iii) being insensitive to administrative-data reporting capacity, which itself may be affected by the USAID shutdown.

The unit of analysis for the nightlights outcome is a circular buffer of fixed radius around each USAID project location recorded in GODAD. We extract from GODAD the latitude and longitude of every USAID-disbursed project location active in Africa during 2017–2020, identifying 1,374 unique sites. For each site i and each month t, we compute four statistics within the buffer: the mean radiance per pixel (*avg_rad*), the total radiance summed across pixels (*sum_rad*), the count of pixels above a luminosity threshold of 0.3 nW/sr/cm^2 (*lit_count*), and the total count of non-masked pixels (*total_count*). The same gas-flare mask described in Section 3.2 is applied throughout. Our baseline outcome variable is $\log(avg\_rad_{i,t} + 0.01)$, which is interpretable as the log of average radiance per pixel and is comparable across buffers of different size.

The buffer radius is itself a parameter of interest. We compute the same statistics

for five radii: 500 metres, 1 kilometre, 5 kilometres, 10 kilometres, and 25 kilometres. The motivation for examining multiple radii is twofold. First, the radii of plausible economic spillover from a USAID project (an office, a clinic, an agricultural extension station, a food-distribution depot) are uncertain ex ante, and the literature does not provide a definitive answer for the spatial scale at which the local economic ripple of an aid intervention attenuates. Second, and more importantly, the pattern of coefficients across radii provides an internal placebo: an effect that is genuinely local to USAID activities should be strongest at the smallest radii and should attenuate as the radius grows to include unrelated economic activity. A pattern in which the coefficient decays monotonically with radius is hard to manufacture with confounding factors that operate at a broader spatial scale. We use the project-buffer level as the *primary* unit of analysis for the nightlights outcome because it offers the most precise spatial identification, and report results at the first administrative subnational level (ADM1) as a complementary aggregate-level check.

Our second outcome is acute food insecurity. To assess the impact of the shutdown on food insecurity, we focus on the IPC Phase 3+ ("Crisis or worse") and Phase 4+ ("Emergency or worse") thresholds because these are the operational thresholds for humanitarian action: Phase 3+ identifies populations in need of immediate intervention to prevent further deterioration, and Phase 4+ identifies populations already experiencing emergency-level food consumption gaps and elevated mortality. The unit of analysis is the IPC Level 1 area, which approximates ADM1; our analytical sample comprises 388 IPC areas in 27 African countries.

4.4 Baseline difference-in-differences specification

For both outcomes, our baseline regression takes the form

$$Y_{i,c,t} \;=\; \beta\ \mathrm{USAID}_i \times \mathrm{Post}_t \;+\; \alpha_i \;+\; \gamma_{c,t}^{(\cdot)} \;+\; \varepsilon_{i,c,t}, \tag{2}$$

where i indexes units (project locations for the nightlights outcome, IPC units for the food-security outcome), c indexes the country containing unit i, and t indexes time (a calendar month for nightlights, an IPC analysis date for food security). The coefficient of interest, β, is the coefficient on the interaction between pre-shock USAID exposure and an indicator Post_t equal to one for observations from February 2025 onward.

The term α_i captures unit-specific fixed effects, which absorb any time-invariant determinant of the outcome at the unit level: the local geography, the long-run economic structure, the baseline level of food insecurity, the underlying electrification and built-up footprint, and so on. The term $\gamma_{c,t}^{(\cdot)}$ is a placeholder for the country-by-time effects that absorb country-level shocks; the precise structure of this term differs across our two

outcomes, in a manner we now explain.

For the nightlights outcome at the project-buffer level, we take $\gamma_{c,t}^{(\cdot)} = \gamma_{c,t}$, that is, a full set of country-by-month fixed effects. This is the most conservative specification possible in our setting: it absorbs the entirety of country-level variation at the monthly frequency, leaving identification to come purely from within-country, within-month variation across project locations of different historical USAID intensity. We are able to afford this aggressive fixed-effect structure at the buffer level because the high spatial density of USAID project locations within each country, on average about twenty-five locations per country, and over one hundred in countries like Uganda, Nigeria and Kenya, provides ample within-country-month variation to identify β.

For the food-security outcome at the IPC-unit level, the same country-by-time fixed-effects strategy would be too restrictive. The IPC publishes only one to four analyses per country per year, and the number of IPC units per country in our sample averages about fourteen but in some smaller countries is as low as five or six. Imposing country-by-analysis-date fixed effects in this setting would erode the identifying variation to the point that the standard errors become uninformative. We therefore weaken the country-by-time structure for the food-security outcome to $\gamma_{c,t}^{(\cdot)} = \gamma_c + \gamma_t + \delta_c \cdot t$, that is, country fixed effects, time fixed effects, and a country-specific linear trend. This specification absorbs (i) all time-invariant country-level differences, (ii) common time effects shared by all countries, and (iii) country-specific linear deviations from the common time trend. It preserves more identifying variation than a country-by-analysis-date specification while still controlling for the most common forms of country-level confounding.

The asymmetry between the two specifications is deliberate and reflects the different temporal structure of the two outcomes: the nightlight outcome is observed at monthly frequency across 112 months, while the IPC outcome is observed only on sparse, country-specific analysis dates. The available identifying variation, and the fixed-effect structure that best preserves it, differ correspondingly. To verify that our results are not sensitive to this asymmetry, we report in Appendix A.8 the trend-break decomposition under the country-trends specification (which serves as a less-conservative robustness for the buffer outcome) and in Figure 10 the IPC event-study coefficients under both the country-trends and unit-specific-trends specifications. The estimated coefficients are larger in magnitude when the less conservative specification is used, as is expected, since more identifying variation is preserved, but the sign and statistical significance of the central conclusions are unchanged.

Standard errors are clustered throughout at the level at which the treatment varies: at the project-location level for the nightlights buffer specifications, at the IPC Level 1 unit for the food-security specifications, and at the GADM 3.6 ADM1 unit for the aggregate-level checks. The number of clusters in our preferred specifications (1,374 for the buffer outcome, 388 for the IPC outcome, 672 for the ADM1 robustness) is comfortably above the

threshold beyond which conventional clustered standard errors are well-behaved (Abadie et al., 2023). As an additional robustness exercise we report standard errors clustered at the country level.

4.5 Parallel trends

In our continuous-treatment setting with a single simultaneous shock, the parallel-trends assumption translates into a stability-of-slope requirement: the slope of the relationship between pre-shock USAID exposure and the outcome would have remained constant absent the shutdown. The pre-shock relationship is itself non-zero by design as the central rationale of foreign aid is to alter economic and humanitarian outcomes in exposed locations. Thus, high-exposure locations may have been growing more rapidly in economic activity (because USAID-funded clinics, contractors and commercial activity were directly supporting local economic flows) or experiencing lower food insecurity (because USAID-funded humanitarian assistance was buffering vulnerable populations) while the agency was operating. The identifying question is therefore not whether the pre-shock relationship exists in levels but whether its slope would have remained constant in the absence of the shock.

We test this in two complementary ways: an event study (Section 4.7) that traces the period-by-period coefficient on USAID exposure, and a trend-break decomposition (Section 4.8) that decomposes the post-shock movement into a continuation of the pre-shock slope versus a discrete level break at the shock date. In addition, our preferred specifications include country-specific linear trends (for the IPC outcome) or country-by-month fixed effects (for the nightlights outcome), which absorb any country-wide linear or unrestricted secular trend that might otherwise contaminate the estimate of β.

We anticipate, as a substantive matter, that the relationship between USAID_i and the outcomes in the pre-shock period reflects the active support that USAID was providing to high-exposure locations. The interpretation of a post-shock break, whether a level shift, a slope reversal, or both, is then that the removal of this support broke a relationship that USAID had been actively sustaining. We return to this interpretive point in the results sections.

4.6 Measurement integrity of the IPC outcome

A concern specific to the IPC outcome, distinct from the parallel-trends assumption of Section 4.5, is whether the IPC measurement infrastructure has itself been degraded by the post-shock loss of USAID funding. The IPC classifications draw on multiple sources of evidence: WFP mobile vulnerability surveys, FEWS NET seasonal monitoring, humanitarian-partner needs assessments, and country-specific administrative inputs

assembled by national Technical Working Groups. Several of these inputs are partly USAID-funded. If the post-shock degradation of these inputs leads to less frequent, less granular, or systematically biased IPC classifications, particularly in the most USAID-exposed countries, then our estimates of the effect on Phase 4+ food insecurity may be biased relative to the true post-shock change in food security.

We address this concern empirically. In Appendix A.5 (Table A4) we report three diagnostics on IPC operational continuity, broken down by country USAID-exposure quartile: the cadence of IPC analyses (analyses per year, pre vs post), the coverage per analysis (number of IPC units classified, pre vs post), and the country-level correlation between USAID exposure and the cadence ratio. The diagnostics yield three findings. First, the mean cadence of IPC analyses has been broadly preserved across the post-shock window, with a median post/pre ratio of approximately 0.95 across the countries with both pre- and post-shock analyses. Second, where post-shock cadence has been reduced, the reduction is concentrated in the lowest, not the highest, USAID-exposure quartile; the country-level correlation between USAID exposure and the post/pre cadence ratio is positive, the opposite of the sign expected under the degradation hypothesis. Third, the coverage per analysis has been preserved or expanded in all nineteen countries with both pre- and post-shock observations, with no country falling below 80% of its pre-shock coverage.

We interpret these patterns as empirical evidence that the most direct manifestations of measurement degradation, namely reduced cadence, reduced coverage, and concentration of the reduction in the most USAID-exposed countries, are not present in the data. Residual concerns about the unobserved quality of the underlying survey inputs to IPC classifications, which we cannot rule out directly, would bias our estimates toward zero (under-detection of acute cases) and therefore leave our findings as conservative lower bounds.

4.7 Event study

The event-study version of our baseline specification estimates the relationship between USAID exposure and the outcome separately for each time bin relative to the shock:

$$Y_{i,c,t} \;=\; \sum_{k \neq -1} \beta_k \; \text{USAID}_i \;\times\; \mathbf{1}\{T_t = k\} \;+\; \alpha_i \;+\; \gamma_{c,t}^{(\cdot)} \;+\; \varepsilon_{i,c,t}, \tag{3}$$

where T_t is a coarsened time index relative to the shock, measured in quarters for the nightlights outcome (so that $T_t = 0$ corresponds to February–April 2025) and in semesters for the food-security outcome (where IPC analyses are typically published twice a year). The bin $k = -1$ is omitted as the reference period; all other coefficients β_k are interpreted as the differential level of the outcome in bin k relative to the immediate pre-shock bin,

for a one-log-unit increase in USAID exposure.

We report event studies for both outcomes and discuss them in Sections 5 and 6. The reader should bear two interpretive points in mind. First, as discussed above, pre-shock event-study coefficients in our setting do not test the standard "no pre-trends" null in the way that they would in a binary-treatment DiD without an active prior policy. They test instead whether the relationship between USAID exposure and the outcome is stable through the pre-shock period, a different and more demanding test. Second, the event study provides direct information about the timing of any post-shock effect: an immediate level break (as we find for the buffer outcome) is consistent with a sudden change in local economic flows; a gradual escalation (as we find for the food-security outcome) is consistent with a slower-acting channel such as the progressive depletion of in-country food-aid commodity stocks following the cessation of new shipments.

4.8 Trend-break decomposition

To formalise the interpretive issue raised in Section 4.5, we estimate an explicit trend-break specification:

$$(4)\quad Y_{i,c,t} \;=\; \beta_1\,\mathrm{USAID}_i{\cdot}\tau_t \;+\; \beta_2\,\mathrm{USAID}_i{\cdot}\tau_t{\cdot}\mathrm{Post}_t \;+\; \beta_3\,\mathrm{USAID}_i{\cdot}\mathrm{Post}_t \;+\; \alpha_i \;+\; \gamma_{c,t}^{(\cdot)} \;+\; \varepsilon_{i,c,t},$$

where τ_t is a linear time trend centered at the shock date (so $\tau_t = 0$ in February 2025), and Post_t is the post-shock indicator as before.

The three coefficients have clear and complementary interpretations. The coefficient β_1 captures the pre-shock slope of the relationship between USAID exposure and the outcome: a non-zero β_1 indicates that the USAID exposure was associated with a differential time-trend in the outcome *while USAID was operating*. The coefficient β_2 captures the change in this slope at the shock date: a negative β_2 indicates that the pre-shock trend was reversed (or attenuated) when the programme was discontinued. The coefficient β_3 captures any discrete level discontinuity at the shock date, on top of any slope change.

The trend-break specification therefore provides a direct test of two distinct hypotheses about how the USAID shutdown affected outcomes. If USAID's effect on the outcome operated primarily through sustaining a level of activity, then we should observe $\beta_3 \neq 0$ and $\beta_2 \approx 0$: a sudden break at the shock date, no change in trend. If USAID's effect operated primarily through sustaining a growth trajectory, then we should observe $\beta_2 \neq 0$ and $\beta_3 \approx 0$: the slope reverses at the shock, but the level adjusts only gradually. The two hypotheses are not mutually exclusive: a programme might support both a level and a growth trajectory, in which case both coefficients would be non-zero.

Section 5.2 reports the estimated coefficients and discusses which of the two hypotheses is consistent with the data.

4.9 Heterogeneity

We probe heterogeneity along four dimensions, each motivated by an economic hypothesis about the channels through which the USAID shutdown might affect local outcomes.

First, we decompose the treatment variable by USAID sector, re-estimating equation (2) with each of the four OECD-DAC macro-sector exposures (social, economic, productive, humanitarian; defined in Section 2) entered individually and then jointly. This decomposition is informative both for the interpretation of the effect — which programmes mattered — and as a constraint on alternative explanations: an effect operating uniformly across all sectors would suggest confounding by general USAID presence, whereas an effect concentrated in sectors with a plausible mechanism narrows the space of plausible confounders.

Second, we examine moderation by pre-shock institutional quality, following the heterogeneity analysis in Rohner et al. (2026) on the conflict outcome. We use two complementary measures: the Polity 5 polity2 score (Marshall and Gurr, 2020) and the V-Dem v15 liberal democracy index v2x_libdem and judicial constraints index v2x_jucon (Coppedge et al., 2025), all averaged over the 2017–2020 pre-shock window. The hypothesis under test is that countries with stronger institutions are better able to absorb an external shock and to mobilise compensatory domestic policy responses, so that the marginal effect of the USAID shutdown on the outcome is dampened.

Third, for the food-security outcome we examine moderation by pre-shock vulnerability. We split IPC units at the within-sample median of their pre-shock average share of population in Phase 3+, and re-estimate the baseline DiD separately on the high-vulnerability and low-vulnerability sub-samples. The hypothesis is that the USAID shutdown pushes already-vulnerable populations further into emergency, while leaving more food-secure populations relatively unaffected.

Fourth, for the nightlights outcome at the project-buffer level we examine moderation by pre-shock urbanisation. We classify each USAID project location as urban or rural based on whether the pre-shock mean radiance in its 1 km buffer exceeds the within-sample median, and we estimate the DiD separately for each sub-sample. This split tests whether the local economic effect of the shutdown is concentrated in urban project sites (where USAID offices, contractors, and beneficiary distribution centres are most concentrated and where the economic activity supported by USAID is most directly visible in nightlights) or in rural sites (where USAID's agricultural and humanitarian activities operate but where the visible-light proxy may be a less sensitive measure of the underlying activity).

4.10 Robustness

We report a battery of robustness checks for both outcomes. We preview them here and refer the reader to Sections 5 and 6 for the full results.

First, we estimate alternative fixed-effect specifications: more conservative (country-by-time fixed effects for both outcomes, where feasible), less conservative (unit fixed effects only), and unit-specific linear trends in the style of Goodman-Bacon (2021). The pattern of results across these specifications provides information about the sensitivity of our estimates to the absorption of country-level secular variation.

Second, we estimate the baseline specification on systematically restricted samples: dropping active-conflict countries, dropping chronic crisis countries, dropping the country with the highest USAID exposure, dropping the top one percent of treated units by exposure, and leaving each country out in turn. The pattern of coefficients across these restrictions documents the extent to which the main result is driven by particular countries or particular outliers.

Third, we report a battery of placebo specifications, with an asymmetric design that reflects the structural differences between the two outcomes. For the buffer outcome, the natural placebo is geographic: we test whether random African locations away from USAID activity exhibit the same post-shock pattern. Specifically, we exploit three placebo exercises. (i) A *spatial-decay* test that is implicit in our headline empirical analysis: because we estimate the DiD at five buffer radii (500 m to 25 km) as part of the main result, an effect that is genuinely local to USAID activities should attenuate monotonically with the radius, vanishing by 25 km. A pattern that does not exhibit such spatial decay would suggest contamination by influences that operate at broader spatial scales. This is therefore not a separate exercise but a structural placebo embedded in the main specification. (ii) A *time-shifted* test in which we re-estimate equation (2) using a fake shock date two years earlier (February 2023) and a sample restricted to the pre-shock period only. A significant coefficient at the fake shock would indicate that our estimated effect is an artefact of timing-unrelated fluctuations. (iii) A *random-point* test in which we replicate the buffer extraction at random points drawn from across the African continent, conditioning on the points being at least fifty kilometres from any USAID project location. A pooled difference-in-differences that compares the USAID buffers with the random buffers identifies the USAID-specific component of any post-shock change.

For the IPC outcome a random-point placebo is not structurally available, because IPC units are administrative entities rather than arbitrary points on a map; there is no analogue to "random IPC areas". The natural placebo at the IPC level is therefore donor-substitution: we re-estimate the headline Phase 4+ DiD with the USAID exposure replaced, in turn, by exposure to France, the United Kingdom, Germany, the World Bank, and China. If the post-shock food-security effect is USAID-specific rather than a generic

correlate of donor presence, these placebo coefficients should be null. The interpretation of donor placebos is complicated by donor interdependence: some U.S. funding is channelled through the World Bank, and the 2025 fiscal climate may have affected several donors contemporaneously. We discuss this explicitly when interpreting the results in Section 6.7.

Taken together, the buffer-side placebo battery (spatial decay, time-shifted, random points) and the IPC-side donor-placebo battery probe two complementary identification threats: the buffer placebos rule out generic post-shock patterns at random African locations, while the IPC donor placebos rule out generic donor-presence confounders at the IPC analytical units.

5 Results — Part I: Local economic activity around USAID project sites

This section reports our main results on the effect of the USAID shutdown on local economic activity. We organise the discussion in seven subsections. Section 5.1 presents the main difference-in-differences result and documents the spatial decay of the effect across buffer radii. Section 5.2 characterises the dynamic structure of the effect using an event study and a trend-break decomposition. Section 5.3 decomposes the effect by USAID sector. Section 5.4 examines heterogeneity by pre-shock urbanisation. Section 5.5 reports an aggregate-level consistency check at the ADM1 administrative scale. Sections 5.6 and 5.7 present a battery of robustness exercises based on sample restrictions and placebo tests.

Before turning to the systematic results, Figure 3 offers a visual illustration of the underlying phenomenon at two of the highest-USAID-exposure locations on the continent: Kinshasa (Democratic Republic of the Congo) and Lusaka (Zambia). Both are major hubs of PEPFAR clinical activity (U.S. Department of State, 2024), with the Kinshasa project location alone accounting for cumulative USAID disbursement of approximately $877 million during 2017–2020 (authors' calculation from GODAD). For each location, the figure shows the pixel-by-pixel difference in nighttime light radiance between the fifteen months following the shock (February 2025 through April 2026) and the year preceding it (calendar year 2024), across a 14 km square centred on the project site. Both panels are predominantly red, indicating broad-based declines in nighttime activity around the project sites: across the 14 km square at Kinshasa, the mean radiance falls by 25% and 74% of pixels exhibit a decline of more than 0.5 nW sr^{-1} cm^{-2}; at Lusaka, the mean radiance falls by 32% and 98% of pixels show a decline of comparable magnitude. The illustrative cases are consistent with the systematic statistical evidence that follows.

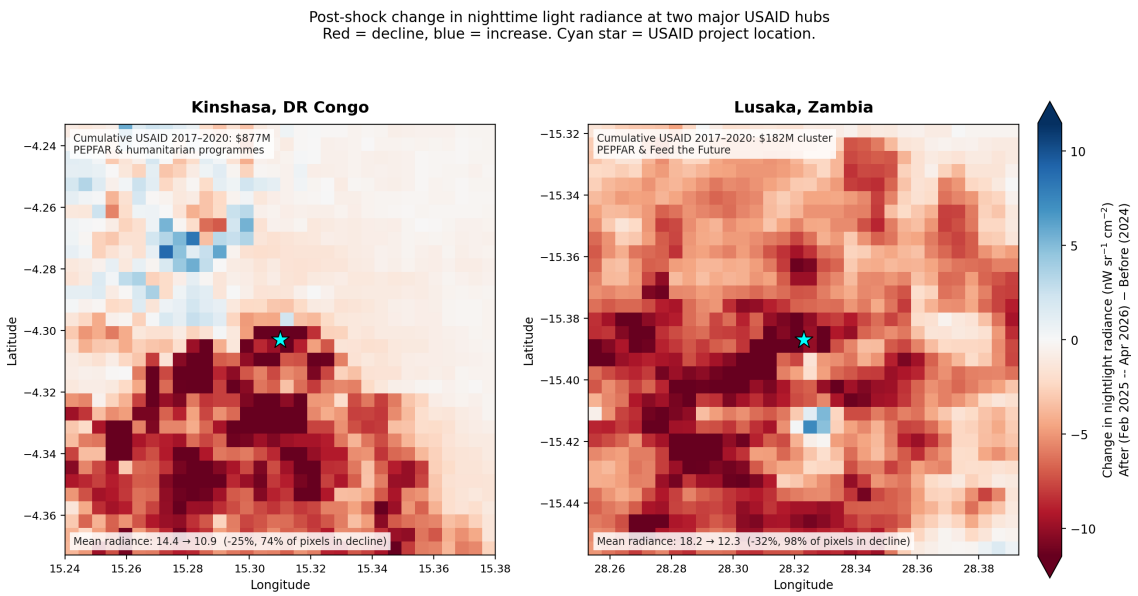

Figure 3. Illustrative cases: post-shock change in nightlight radiance at two major USAID hubs

Notes. Pixel-by-pixel change in VIIRS monthly mean nighttime light radiance between the post-shock period (mean of monthly composites from February 2025 through April 2026) and the pre-shock period (mean of monthly composites for calendar year 2024), for 14 km squares centred on two USAID project locations. The left panel shows Kinshasa, Democratic Republic of the Congo. The right panel shows a USAID project location in the Lusaka, Zambia, metropolitan area. The cyan star in each panel marks the USAID project location. The diverging colour scale is shared across the two panels and clipped at the 95th percentile of |difference|: red denotes pixels with a post-shock decline in radiance, blue denotes pixels with a post-shock increase. *Source:* authors' calculations from VIIRS DNB monthly composites VCMSLCFG accessed via Google Earth Engine.

5.1 Spatial decay of the post-shock effect

Table 2 reports estimates of the baseline difference-in-differences specification of equation (2) at each of the five buffer radii. Panel A reports the baseline specification with project fixed effects and country × month fixed effects, the most conservative specification that absorbs the entirety of country-level temporal variation and identifies the coefficient purely from within-country, within-month variation across project locations of different pre-shock USAID exposure. Panel B reports the primary specification, which replaces the country × month effects with country fixed effects, month fixed effects, and country-specific linear trends; this specification preserves slightly more identifying variation and is the specification we use throughout the rest of the paper as our main reference.

The headline pattern is the monotonic attenuation of the coefficient with the buffer radius. In Panel B, $\widehat{\beta}$ declines from approximately -0.005 at the smallest buffers (500 m and 1 km, both statistically significant at the 5% level) to approximately -0.001 at 25 km, no longer statistically distinguishable from zero. The intermediate radii (5 km and 10 km) fall between these bounds and remain statistically significant. The decline is roughly a factor of four across the radius range, from the buffer where the treated area is most tightly identified with the project site itself to the buffer where it is large enough to include

substantial economic activity unrelated to the project. Panel A produces an essentially identical pattern in shape, with marginally smaller magnitudes and wider standard errors, as expected from the more aggressive absorption of country-level temporal variation through country × month fixed effects.

Table 2. Spatial decay: baseline DiD for the project-buffer VIIRS outcome

	500 m	1 km	5 km	10 km	25 km
Panel A. Baseline (project + country×month FE)					
$\log(\text{USAID})_i \times \text{Post}_t$	−0.0047*	−0.0043**	−0.0033**	−0.0025*	−0.0009
	(0.0024)	(0.0022)	(0.0015)	(0.0013)	(0.0009)
Panel B. Primary (project + country FE + month FE + country trends)					
$\log(\text{USAID})_i \times \text{Post}_t$	−0.0050**	−0.0047**	−0.0038**	−0.0028**	−0.0012
	(0.0022)	(0.0020)	(0.0014)	(0.0012)	(0.0009)
Observations	153,888	153,888	153,888	153,888	153,888
Project clusters	1,374	1,374	1,374	1,374	1,374

Notes. Dependent variable is $\log(\text{avg\_rad}_{i,t}+0.01)$, the log of mean VIIRS nightlight radiance within the indicated buffer radius. Treatment variable is $\log(1+\overline{D}_i^{2017\text{–}2020})$, the log of average annual USAID disbursement to project location i during 2017–2020. *Post* is an indicator for observations from February 2025 onward. Sample: 1,374 USAID project locations, 112 monthly observations from January 2017 to April 2026. Standard errors clustered at the project-location level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We discuss the placebo properties of the decay in detail in Section 5.7; here we note that the decay is theoretically expected from any economic activity that is generated by USAID at the project location and that diffuses to surrounding areas with declining intensity. A pattern that did not exhibit such decay, one in which the coefficient remained large at the 25 km radius or actually increased with the radius, would suggest that the estimated effect is contaminated by influences that operate at a broader spatial scale than the USAID project itself.

Figure 4 visualises the spatial decay graphically. The decay is observable to the naked eye and survives both fixed-effect specifications. Under the primary specification (Panel B), the 95% confidence intervals at the 500 m and 1 km radii contain only the negative half of the real line; the interval at the 25 km radius spans both sides of zero.

Magnitude. At the 1 km buffer, $\widehat{\beta} = -0.0047$ means that, comparing two USAID project locations within the same country at the same month, a one-log-unit (a factor of $e \approx 2.72$) increase in cumulative pre-shock USAID disbursement is associated with a decrease of approximately 0.47% in mean post-shock nightlight radiance. To put this in perspective: the interquartile range of $\log(1 + \overline{D}_i)$ in our sample is approximately 4.4 log

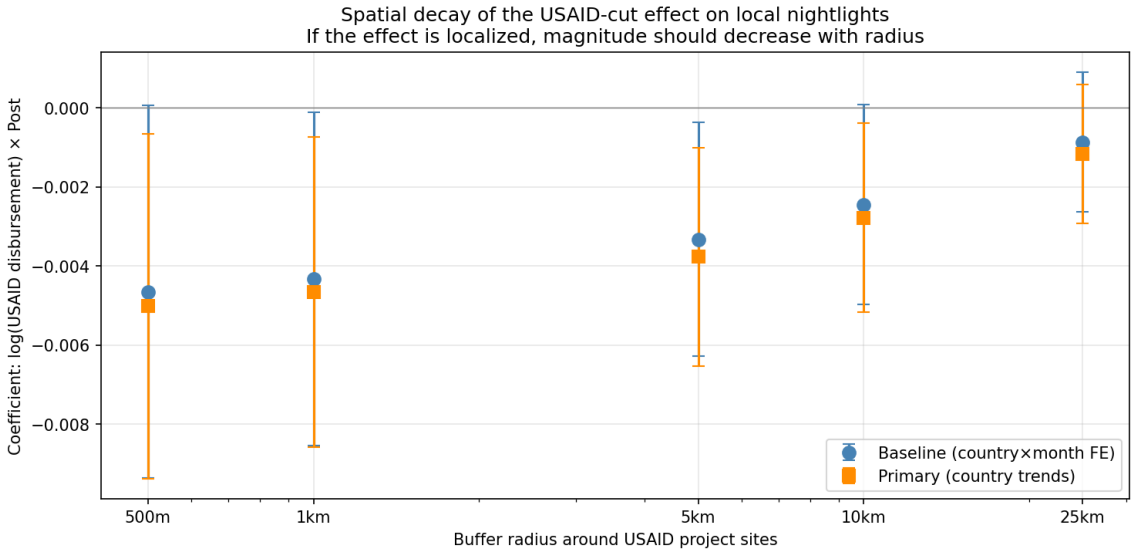

Figure 4. Spatial decay of the USAID-shutdown effect on local nighttime light radiance
Notes. The figure plots the difference-in-differences coefficient on $\log(\text{USAID})_i \times \text{Post}_t$ from equation (2), with $\log(\text{avg\_rad}_{i,t} + 0.01)$ as the dependent variable, separately for each buffer radius around the 1,374 USAID project locations in Africa. Two specifications are reported: *Baseline* includes project fixed effects and country × month fixed effects; *Primary* replaces the country×month effects with country fixed effects, month fixed effects, and country-specific linear trends. Standard errors are clustered at the project-location level; bars represent 95% confidence intervals.

units (from 10.1 to 14.4), giving a differential effect of approximately 2.0% between a project at the 75th and a project at the 25th percentile of pre-shock exposure. For a project at the 90th percentile of exposure (cumulative disbursement of approximately $15.07 million) relative to a project at the median (approximately $119,000), the differential effect on post-shock nightlights is approximately 2.3%.

Visual evidence. Figure 5 provides a visualisation of the within-country variation in nightlight changes that the difference-in-differences exploits, together with the spatial and quantitative distribution of USAID exposure. The top panel maps each USAID project location as a scatter point whose colour denotes the post-shock change in log nightlight radiance relative to the average post-shock change of all USAID locations in the same country (red: below the country mean change; blue: above), and whose size denotes the quartile of pre-shock USAID exposure $\log(1 + \overline{D}_i)$. The figure shows changes relative to the country mean rather than absolute changes because the secular upward trend in African nighttime light intensity driven by ongoing urbanisation and rural electrification (Falchetta et al., 2019) means that the majority of locations exhibit positive raw changes; the DiD identifies the within-country, within-month *differential* effect. The bottom panel summarises the relationship visually: it reports the mean country-demeaned post-shock change in log nightlight radiance for each quartile of pre-shock USAID exposure, with ±1.96 standard-error whiskers. The means are positive in the two lower-exposure quartiles (Q1: +0.0115; Q2: +0.0171) and negative in the two higher-exposure quartiles

(Q3: −0.0128; Q4: −0.0158), with a Q1−Q4 differential of approximately +0.027 log units. The pattern is the visual signature of the negative DiD coefficient that Table 2 identifies on all 153,888 unit-month observations.

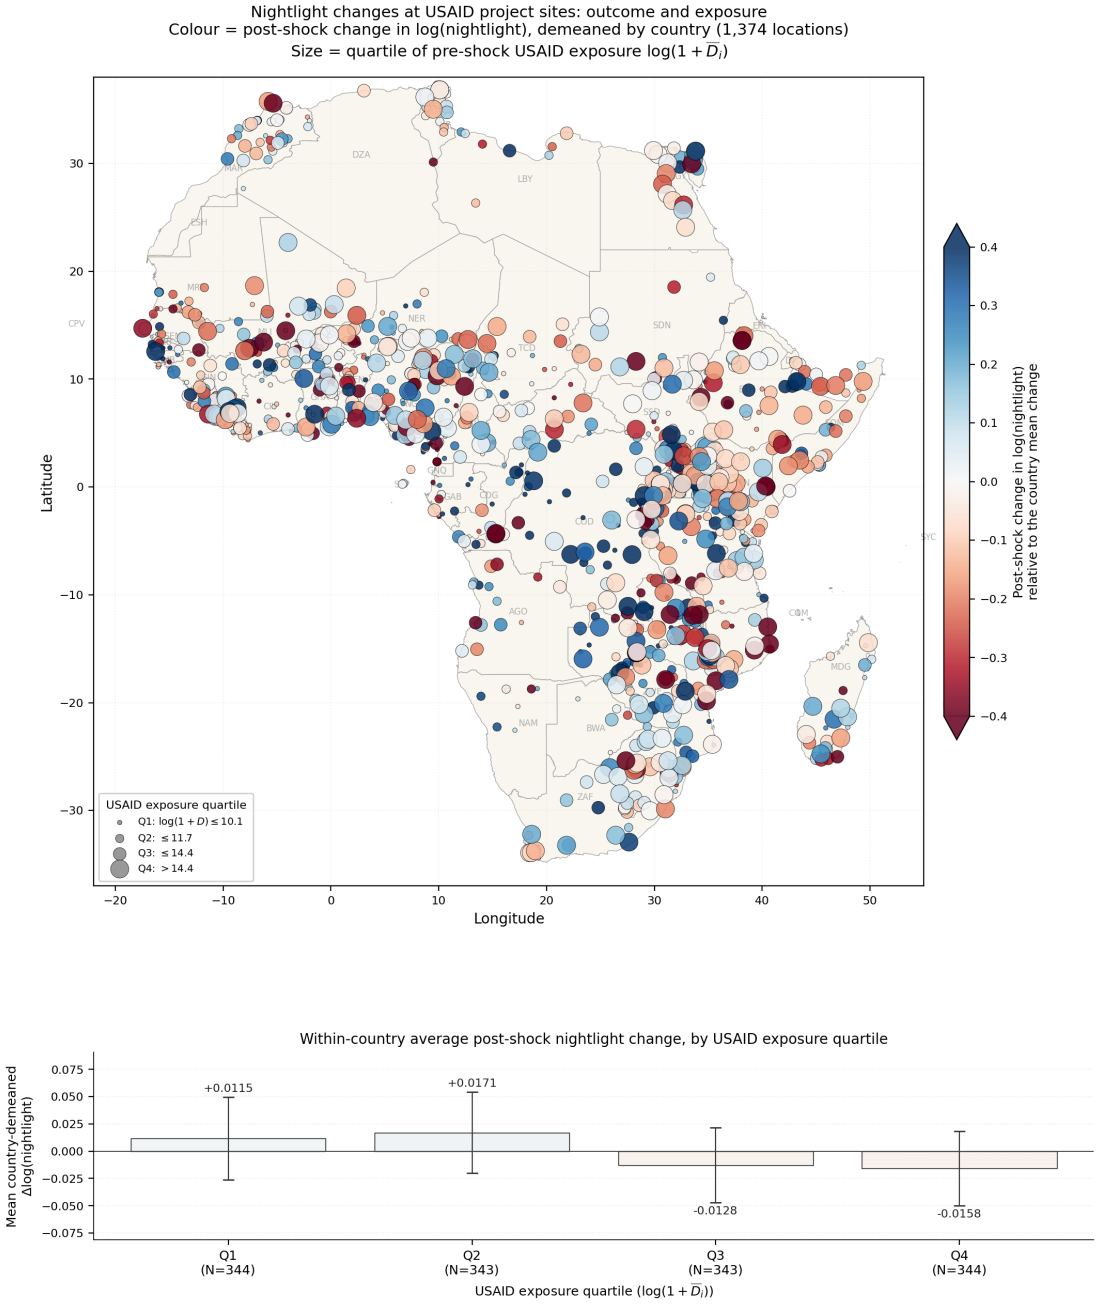

Figure 5. Post-shock change in nightlights at USAID project sites, relative to the country mean

Notes. Each point is one of the 1,374 USAID project locations. Colour denotes the post-shock change in mean log nightlight radiance within the 1 km buffer, demeaned by the mean post-shock change of all USAID project locations in the same country. Red points are project sites whose post-shock change in nightlights was below their country's average; blue points are sites whose change was above their country's average. Country averages and the secular continental upward trend in African nightlights are absorbed into the country-demeaning, so that the figure visualizes the within-country, within-month variation on which the DiD coefficient relies. The color scale is clipped at ±0.4 log units. *Source:* authors' calculations from VIIRS DNB monthly composites VCMSLCFG and GODAD v1.0.

Time-series evidence. A second visualisation of the underlying pattern is provided in Figure 6, which plots the monthly evolution of mean log nightlight radiance separately for the top 20% (deciles 9-10), middle 60% (deciles 3-8), and bottom 20% (deciles 1-2) of USAID project locations by pre-shock disbursement. Each series is plotted as a deviation from its own pre-shock average and smoothed with a six-month rolling mean to reduce noise. The three series track each other closely through the entire 2017–2024 period, with no detectable differential trend; immediately after the January 2025 shock, the top-20% series diverges downward relative to the bottom-20% series, opening a persistent gap that remains visible through 2025 and 2026. This visual pattern — a discrete divergence at the shock date followed by a roughly stable gap — is the time-series counterpart of the level-shift evidence formalised by the trend-break decomposition (Section 4.8).

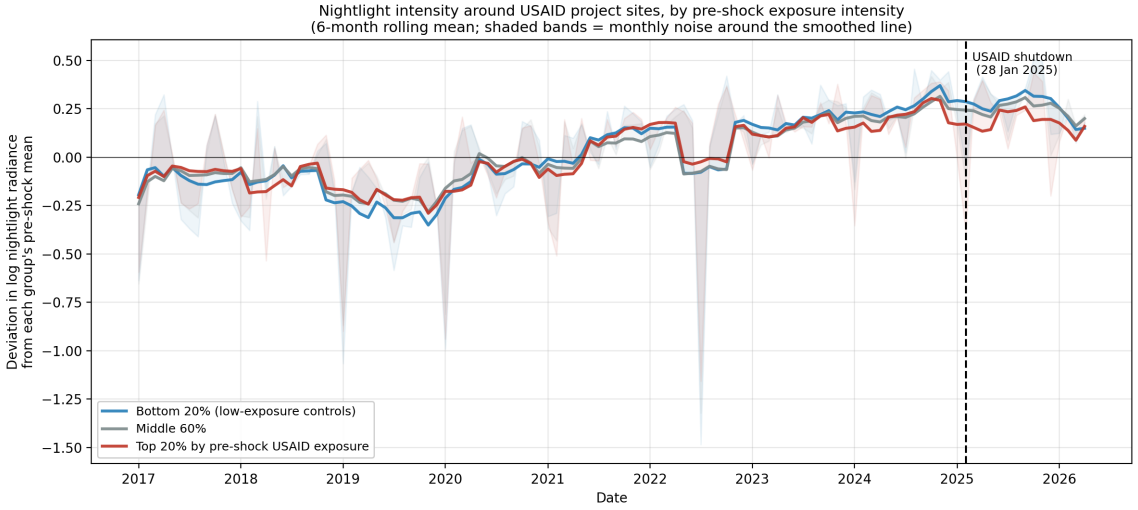

Figure 6. Nightlight evolution at USAID project sites, by pre-shock exposure intensity
Notes. The figure plots monthly mean log(avg_rad + 0.01) within the 1 km buffer around USAID project locations, separately for three groups defined by the within-sample decile of cumulative USAID disbursement during 2017–2020: top 20% (deciles 9–10), middle 60% (deciles 3–8), and bottom 20% (deciles 1–2). Each series is normalised to its own pre-shock mean (so that the value at any month is the deviation from the pre-shock baseline of that group) and smoothed with a six-month rolling mean; the shaded bands around each line are the unsmoothed monthly values. The vertical dashed line marks the USAID shutdown of 28 January 2025. The three series track each other closely through the entire pre-shock period; post-shock, the top-20% series begins to diverge downward.

5.2 Dynamic effects: event study and trend-break decomposition

The headline DiD coefficient pools all post-shock months into a single estimate. We resolve its temporal structure with two complementary analyses introduced in Sections 4.7 and 4.8: an event study tracing the semester-by-semester evolution of the USAID–nightlight relationship, and a trend-break decomposition distinguishing a discrete level shift from a slope reversal.

Event study. Figure 7 plots the event-study coefficients β_k from equation (3) at the 1 km buffer, with bins defined in calendar semesters relative to the shock. Coefficients are re-centered so that the average of the pre-shock bins is normalised to zero, making the visual baseline coincide with the average pre-shock USAID–nightlight relationship rather than with the arbitrarily chosen reference bin.

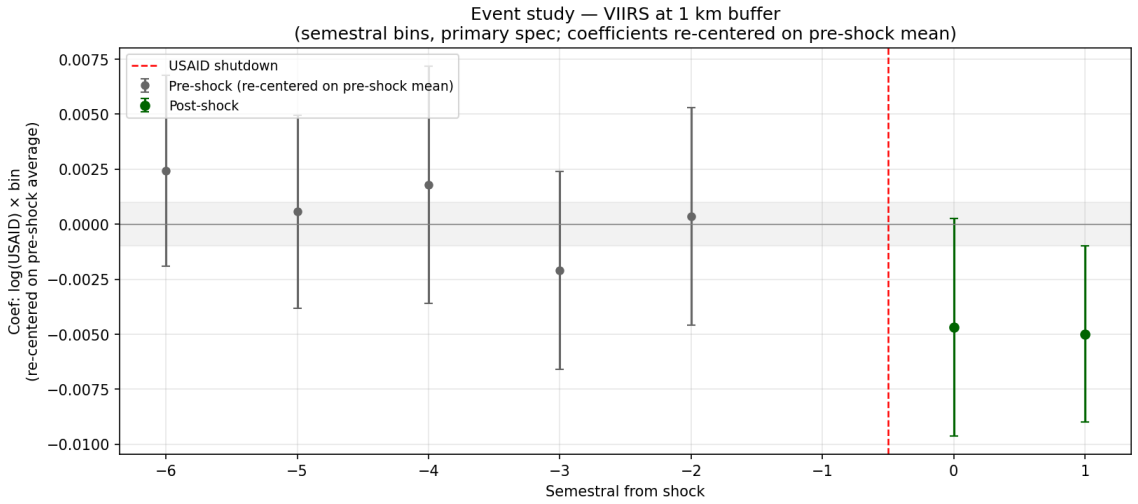

Figure 7. Event study at the 1 km buffer (semestral bins, re-centered on pre-shock mean)
Notes. Coefficients β_k from equation (3) at the 1 km buffer, with bins defined in calendar semesters relative to the USAID shock of January 2025. Each coefficient is re-centered by subtracting the average of the six pre-shock coefficients ($k = -6$ through $k = -1$, with $k = -1$ entered as zero by construction as the omitted reference of the regression); the resulting estimates represent deviations of each bin from the pre-shock average rather than from $k = -1$ alone. Whiskers are ± 1.96 standard errors, conservative since they reflect the original SEs of the regression rather than the slightly smaller SEs of the linearly transformed coefficients. Post-shock bin $k = +2$ (February–April 2026) is excluded from the plot because it covers only three of the six calendar months of a complete semestral bin; its data continue to enter the regression through the unit and time fixed effects, and the re-centered coefficient is +0.0028 (SE 0.0039), statistically indistinguishable from zero or from the preceding negative bins. Primary specification (project FE, country FE, month FE, country-specific linear trends). Dashed vertical line: USAID shutdown date.

Two features of the pattern are visible. First, the pre-shock coefficients scatter around zero with confidence intervals all crossing zero, consistent with the stability-of-slope identifying assumption of Section 4.5: high-USAID-exposure locations were not on a differential trend relative to low-USAID-exposure locations in the years preceding the shock. Second, both complete post-shock semestral bins ($k = 0$: February–July 2025; $k = 1$: August 2025–January 2026) settle at approximately -0.005, separating from the pre-shock cloud and matching in magnitude the headline DiD coefficient reported in Table 2.

The event study is well-suited to a visual pre-trend test but, with only fifteen post-shock months distributed across complete semestral bins, has limited per-bin statistical power: individual post-shock bin coefficients are not significant at conventional levels. The statistically powerful complement is the trend-break decomposition of Section 4.8, which exploits all fifteen post-shock months simultaneously in a single parametric model and

confirms the level shift visible in the figure with high significance. Appendix A.6 reports a leave-one-country-out exercise on the event-study estimates that verifies the post-shock pattern is not driven by any single country in the sample. Additionally, Appendix A.7 replicates the event study at the 5 km buffer, where tighter standard errors render the pre-shock pattern even cleaner while preserving the same post-shock level shift.

Trend-break decomposition. Table 3 reports the trend-break decomposition of equation (4), which separates the post-shock effect into a pre-shock slope (β_1, capturing any pre-existing time-trend of the USAID–nightlight relationship), a slope break at the shock (β_2, capturing a change in that trend at the shock date), and a level discontinuity at the shock (β_3, capturing a discrete change in the level of the relationship at the shock date). Under both the baseline and the conservative specifications, the pre-shock slope β_1 is statistically indistinguishable from zero, indicating no detectable pre-shock time-trend in the USAID–nightlight relationship once unit fixed effects are absorbed. The slope-break coefficient β_2 is positive and either marginally significant or null across the two specifications, providing no evidence of a downward break in the slope at the shock date; if anything, the post-shock slope tilts very slightly upward. The level-shift coefficient β_3 is negative and statistically significant in both specifications, identifying a discrete downward shift in the USAID–nightlight relationship at the shock date.

A natural reading of this pattern is that USAID may have been sustaining a steady-state level of local economic activity at its project locations, and that the cessation of activities produced a discrete downward break in that level when the shutdown took effect. We find no evidence for an alternative interpretation under which USAID had been driving a positive growth trajectory in local activity that the shock interrupted: the pre-shock slope is effectively flat. The level-shift reading is consistent with the operational reality of the cessation: salaries paid monthly stopped, contractor offices closed within weeks, food and medical commodity shipments halted at ports of entry. These are stock-like changes in the underlying activity rather than changes in its rate of growth.

5.3 Sectoral channels

We next examine which USAID sectors drive the post-shock effect on nightlights. We re-estimate equation (2) replacing the total-disbursement treatment with each of the four OECD-DAC sector-specific disbursements (social, economic, productive, humanitarian; see Section 2 for definitions). Table 4 reports four regressions in which each sector is entered separately (columns 1–4) and a joint regression in which all four sectors are entered together (column 5).

The most informative column is the joint regression. Two coefficients are negative and statistically significant: the productive sector and the humanitarian sector. The

Table 3. Trend-break decomposition at the 1 km buffer

	Baseline [(a)]	Conservative [(b)]
β_1: $\log(\text{USAID})_i \cdot \tau_t$ (pre-shock slope)	−0.00004	−0.00003
	(0.00003)	(0.00003)
β_2: $\log(\text{USAID})_i \cdot \tau_t \cdot \text{Post}_t$ (slope break)	0.00063*	0.00031
	(0.00035)	(0.00029)
β_3: $\log(\text{USAID})_i \cdot \text{Post}_t$ (level shift at shock)	−0.00743***	−0.00460*
	(0.00283)	(0.00264)
Observations	153,888	153,776
Project clusters	1,374	1,374

Notes. Trend-break specification, equation (4). The dependent variable is $\log(\text{avg\_rad}_{i,t} + 0.01)$ in the 1 km buffer; τ_t is a linear trend centered at the shock date (zero in February 2025). (a) "Baseline" includes project and month fixed effects only. (b) "Conservative" adds country × month fixed effects (the most restrictive specification). Trend-break results under the primary specification (project + country fixed effects + month fixed effects + country-specific linear trends) are very similar to (a) and are reported in Appendix Table A5. Standard errors clustered at the project-location level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

social-sector coefficient is essentially zero, and the economic-sector coefficient is small and statistically insignificant.

Figure 8 extends the sectoral decomposition across all five buffer radii. The pattern in the joint regression at each radius is consistent: the productive and humanitarian sectors show negative significant coefficients at all five radii, with magnitudes that attenuate mildly with the radius but remain significant through 25 km; the social and economic sectors show coefficients statistically indistinguishable from zero throughout, with the single exception of a marginally significant positive social-sector coefficient at the 25 km buffer.

Of the four sectoral channels through which the USAID portfolio could plausibly affect local nightlight activity, two contribute measurably to the post-shock decline within the first fifteen months: the productive sector (agriculture, irrigation, smallholder market access, Feed the Future) and the humanitarian sector (food assistance, cash transfers, emergency response, BHA, Food for Peace). Both effects survive across all buffer radii in the joint specification, with monotonic spatial decay.

For both sectors, a plausible reading of the mechanism through which the cessation registers in nighttime light radiance is the disruption of an operational footprint that is spatially concentrated at or near the project location.

Table 4. Sectoral heterogeneity at the 1 km buffer: joint regression

	(1)	(2)	(3)	(4)	(5)
$\log(\text{USAID}^{\text{social}})_i \times \text{Post}_t$	0.0006				0.0006
	(0.0017)				(0.0017)
$\log(\text{USAID}^{\text{economic}})_i \times \text{Post}_t$		0.0009			0.0009
		(0.0030)			(0.0030)
$\log(\text{USAID}^{\text{productive}})_i \times \text{Post}_t$			−0.0043**		−0.0043**
			(0.0021)		(0.0021)
$\log(\text{USAID}^{\text{humanitarian}})_i \times \text{Post}_t$				−0.0052**	−0.0051**
				(0.0023)	(0.0023)
Observations	153,888	153,888	153,888	153,888	153,888
Project clusters	1,374	1,374	1,374	1,374	1,374

Notes. Dependent variable is $\log(\text{avg\_rad}_{i,t}+0.01)$ in the 1 km buffer. Sectoral exposure variables are constructed by classifying each project-location-year record by the first digit of its OECD-CRS purpose code: *social* (1xx), *economic* (2xx), *productive* (3xx, agriculture/forestry/industry/trade), *humanitarian* (5xx and 7xx, food assistance and emergency response). Each column reports a separate regression: columns (1)–(4) include only the indicated sector's exposure; column (5) enters all four sectors jointly. All specifications include project fixed effects, country fixed effects, month fixed effects, and country-specific linear trends. Standard errors clustered at the project-location level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

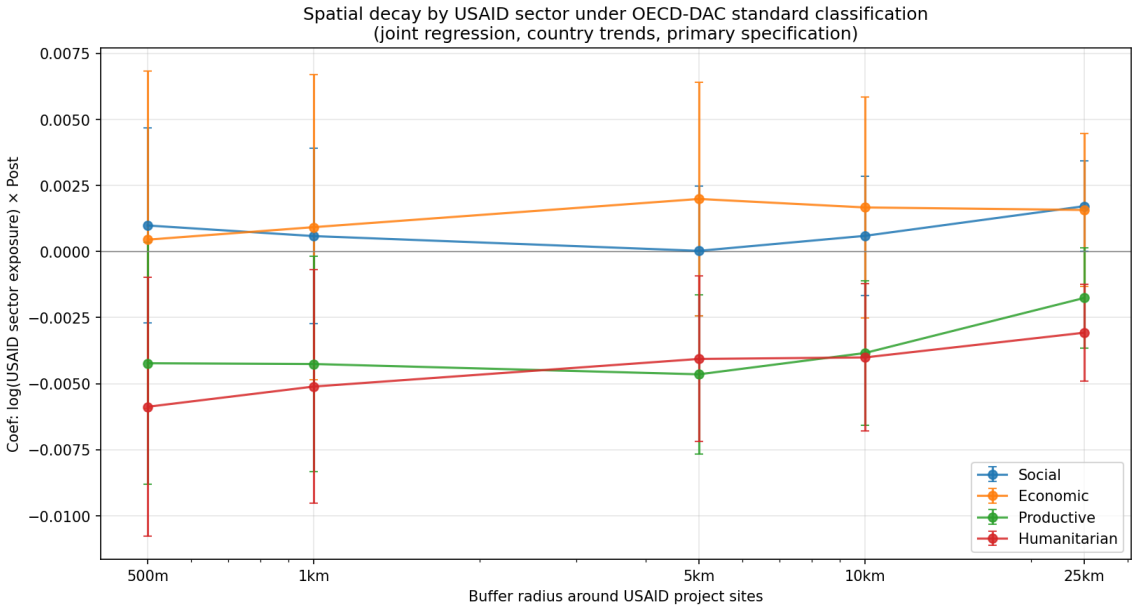

Figure 8. Spatial decay by USAID sector — joint regression at the project-buffer level

Notes. The figure plots the coefficients on $\log(\text{USAID}^{(s)})_i \times \text{Post}_t$ for each USAID sector $s \in$ {social, economic, productive, humanitarian}, entered jointly in a single regression at each buffer radius. Sector definitions follow the OECD-DAC standard classification of CRS purpose codes (Section 2). The primary specification (country trends) is used. Bars represent 95% confidence intervals.

Productive-sector activities are not confined to agricultural production in the fields:

they encompass extension-service offices, seed multiplication centres, irrigation infrastructure maintenance, storage and processing facilities, market hubs, and the salaries paid to local agricultural-development staff and contractors. To the extent that these activities involve electrified facilities or sustain local consumption by staff and contractors, they may plausibly contribute to localised nightlight emission. Humanitarian-sector activities operate through a comparable infrastructure: BHA emergency operations, Food for Peace warehousing and distribution, and the cash and commodity transfers themselves run through implementing-partner offices, depots, and beneficiary-receiving centres. While we cannot directly observe the channels through which the shutdown is transmitted to nighttime radiance, the pattern of monotonic spatial decay documented in Section 5.1 is consistent with this interpretation: an effect propagating from a localised operational footprint should attenuate sharply with distance.

5.4 Heterogeneity by pre-shock nightlight intensity

Table 5 reports the DiD estimates separately on two sub-samples defined by the pre-shock mean radiance in the 1 km buffer: project locations with pre-shock mean radiance above the within-sample median ($N = 76{,}944$) and project locations at or below the median ($N = 76{,}944$). The split is on the baseline level of the nightlights outcome itself: we make no claim that the two groups correspond to externally validated urban/rural categories, although the higher-baseline group is naturally interpretable as comprising more electrified and more economically dense locations and the lower-baseline group as comprising less electrified ones.

Table 5. Heterogeneity by pre-shock nightlight intensity at the 1 km buffer

	Total USAID	Productive sector	Humanitarian sector
Above-median baseline radiance ($N = 76{,}944$)			
coef.	-0.0073^{**}	-0.0055^{*}	-0.0068^{*}
s.e.	(0.0036)	(0.0032)	(0.0041)
Below-median baseline radiance ($N = 76{,}944$)			
coef.	-0.0018	0.0002	-0.0017
s.e.	(0.0016)	(0.0015)	(0.0014)

Notes. Each cell is from a separate regression on the indicated sub-sample, with $\log(\text{avg\_rad}_{i,t} + 0.01)$ as dependent variable and the indicated USAID treatment interacted with Post_t. Sectoral exposures follow the OECD-DAC standard classification of CRS purpose codes: productive (CRS 3xx, agriculture/forestry/fishing/industry/trade, including Feed the Future), humanitarian (CRS 5xx and 7xx, food assistance and emergency response, including BHA and Food for Peace). Sub-samples are defined by whether the project location's pre-shock mean radiance in the 1 km buffer is above the within-sample median ("above-median baseline") or at or below it ("below-median baseline"). All regressions include project FE, country FE, month FE, and country-specific linear trends. Standard errors clustered at the project-location level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Three patterns are visible. First, the total-USAID coefficient is markedly larger in above-median baseline locations: the effect is statistically significant at the 5% level in the above-median sub-sample, the below-median effect is not, and the above-median effect is roughly four times the below-median effect in absolute magnitude. Second, the productive-sector coefficient mirrors this pattern: negative and marginally significant in above-median locations, essentially zero in below-median locations. Third, the humanitarian-sector coefficient is also concentrated in above-median baseline locations, where it is negative and marginally significant; in below-median locations the humanitarian coefficient is negative but statistically indistinguishable from zero. The concentration of the effect in above-median baseline locations is therefore consistent across the total-USAID, productive-sector, and humanitarian-sector treatments.

The pattern of stronger effects in above-median baseline locations is consistent with two features of the data, both of which suggest a natural interpretation in terms of the more electrified component of the USAID-funded portfolio. First, project sites with higher baseline radiance tend to host the institutional infrastructure of the agency — offices, contractor headquarters, central distribution points for commodities, and beneficiary registration centres — which are concentrated in more urbanised settings; the cessation of activities at such sites produces a discrete drop in the visible economic footprint that VIIRS measures directly. Second, the nightlights proxy itself is more sensitive to changes in economic activity in well-lit settings, where there is more electrified activity to lose, than in dimly-lit settings, where much of the underlying economic activity (subsistence agriculture, informal markets, household-level consumption) is invisible to the satellite to begin with. The pattern is therefore consistent with the IPC food-security results reported in the next section (Section 6), which detect humanitarian-sector effects on food insecurity that the VIIRS proxy is less sensitive to in low-baseline locations.

5.5 Aggregate-level consistency check at ADM1

As a complementary check, we re-estimate the baseline DiD at the level of the first administrative subnational unit (ADM1) of GADM 3.6, using the GAUL → GADM crosswalk described in Section 3.2. The unit of analysis becomes the ADM1 polygon (722 polygons in 54 African countries) and the monthly observation, for a panel of 80,864 unit-month observations. The treatment variable is the log of average annual USAID disbursement to the ADM1 region during 2017–2020, computed by aggregating GODAD project-level disbursements.

This aggregate-level analysis serves three purposes. First, it provides a consistency check that the effect detected at the buffer level is not an artefact of the buffer-construction methodology: if the effect that we detect within 1 km of USAID project sites is the local manifestation of a real change in economic activity, then the same change should be visible

(in attenuated form, because of aggregation dilution) when we aggregate to the larger administrative scale. Second, it provides a unit of analysis that is comparable both with the IPC food-security analysis of Section 6 (which is also at the subnational administrative scale) and with Rohner et al. (2026), who use ADM1 ACLED data on conflict outcomes. Third, it allows us to verify the within-country variation argument by examining whether the relationship between USAID exposure and nightlights at the ADM1 scale displays the spatial-economic pattern we would expect.

The choice of specification at the ADM1 scale deserves particular attention. At the project-buffer level, country × month fixed effects are the appropriate baseline: the within-country density of treated units (approximately 25 project locations per country) is high enough that the interaction with continuous USAID exposure preserves substantial identifying variation, and the saturated FE structure absorbs any country-specific seasonal or macroeconomic shock that might otherwise confound the estimates. At the ADM1 scale, however, the within-country density of treated units is considerably lower (12–13 ADM1 units per country on average), so that country × month FE absorb a large share of the identifying variation in the post-shock interaction. For this reason, the country-trends specification, which controls for country-specific linear trends while preserving the cross-sectional and time variation in the post-shock interaction, is our preferred specification at this scale. We nonetheless report both specifications for transparency and as a robustness check.

Table 6 reports the two estimates. Under the country-trends specification, our preferred baseline at this scale, the estimated coefficient is negative and statistically significant at the 1% level. Under the country × month FE specification, the coefficient is statistically null. The attenuation between the two specifications is consistent with the expected dilution of the local-buffer effect when the treatment is averaged over an entire administrative region and a substantial fraction of the identifying variation is absorbed by the saturated FE structure.

Table 6. Aggregate-level consistency check at ADM1

	(1) Country trends	(2) Country×month FE
$\log(\text{USAID})_i \times \text{Post}_t$	-0.0034^{***}	-0.0006
	(0.0009)	(0.0011)
Observations	80,864	80,752
ADM1 clusters	722	722

Notes. Dependent variable is $\log(\text{avg\_rad}_{i,t} + 0.01)$ averaged over the ADM1 polygon. Column (1) uses ADM1 fixed effects, country fixed effects, month fixed effects, and country-specific linear trends. Column (2) uses ADM1 fixed effects and country×month fixed effects. Standard errors clustered at the ADM1 level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The order of magnitude of the country-trends ADM1 coefficient is comparable to the country-trends coefficient at the 10 km buffer (-0.0028 in Table 2), consistent with the ADM1 spatial scale being roughly comparable to a 10–15 km radius for typical African subnational units. Under the more saturated country $\times$ month fixed-effect specification, the point estimate remains negative (-0.0006) but is no longer statistically distinguishable from zero. We interpret the loss of significance as the mechanical consequence of two structural features of the ADM1 aggregation rather than as a substantive reversal of the buffer-level evidence. First, the localised effect identified at the 1 km buffer is averaged over an ADM1 polygon two to three orders of magnitude larger; once country-wide variation is absorbed by the saturated fixed effects, the residual identifying variation captures only the within-country differential changes across ADM1 units, which are necessarily a small fraction of the aggregate signal. Second, the ADM1 panel contains 672 clusters compared to the buffer panel's 1,374, halving the cross-sectional dimension on which the within-country-month coefficient is identified. Read jointly with the buffer-level results, the ADM1 estimates indicate that the post-shock decline in local economic activity survives aggregation to the administrative-region scale under the country-trends specification, while under the most conservative specification the sign of the effect is preserved but its statistical precision is attenuated as expected.

5.6 Sample restrictions and leave-one-country-out

Table 7 reports the baseline DiD coefficient at the 1 km buffer under a series of sample restrictions. The first four restrictions test whether the result is driven by specific categories of countries or projects that might have particular dynamics in the post-shock period. Dropping six countries with active armed conflict during the estimation window (Sudan, South Sudan, the Central African Republic, Mali, Niger, Burkina Faso) leaves the coefficient essentially unchanged and statistically significant at the 5% level. Dropping five countries with chronic food crises (Somalia, Sudan, South Sudan, the Central African Republic, and the Democratic Republic of the Congo) attenuates the coefficient slightly while preserving statistical significance at the 5% level. Dropping the three largest recipient countries by the number of USAID project sites (Uganda, Nigeria, Kenya) actually *strengthens* the effect, with statistical significance rising to the 1% level. Dropping the top one percent of treated units by exposure (the fourteen most exposed project locations) attenuates the coefficient, which remains marginally significant at the 10% level.

The leave-one-country-out exercise, summarised in the bottom panel of Table 7 and reported in full in Appendix Table A6, drops each of the fifty countries in the sample in turn and re-estimates the baseline DiD. The estimated coefficient remains negative across all fifty subsets, with the smallest absolute magnitude obtained when the Democratic

Republic of the Congo is dropped (the largest single-country contribution to the estimated effect) and the largest absolute magnitude when Liberia is dropped. Forty-seven of the fifty subsets produce a coefficient that is statistically significant at the 5% level, and all fifty subsets produce a coefficient that is significant at the 10% level. No single country drives the result.

Table 7. Robustness: sample restrictions at the 1 km buffer

Restriction	β	SE	p	# proj.	# countries
R1: Full sample	-0.0047^{**}	0.0020	0.020	1,374	50
R2: Drop active-conflict[(a)]	-0.0052^{**}	0.0021	0.014	1,218	44
R3: Drop chronic crisis[(b)]	-0.0039^{**}	0.0020	0.049	1,218	45
R4: Drop top-3 by # projects[(c)]	-0.0060^{***}	0.0023	0.008	1,073	47
R5: Drop top-1% exposed projects	-0.0032^{*}	0.0019	0.097	1,362	50
Leave-one-country-out range:					
Min β (drop COD)	-0.0034^{*}	0.0019	0.070		
Max β (drop LBR)	-0.0058^{***}	0.0021	0.007		
Subsets with $p < 0.05$	47 / 50				
Subsets with $p < 0.10$	50 / 50				

Notes. (a) Countries dropped: SDN, SSD, CAF, MLI, NER, BFA. (b) Dropped: SOM, SDN, SSD, CAF, COD. (c) Dropped: UGA, NGA, KEN. All regressions use the primary specification with project FE, country FE, month FE, and country-specific linear trends; dependent variable is $\log(\text{avg\_rad}_{i,t}+0.01)$ in the 1 km buffer; treatment is $\log(1+\overline{D}_i^{2017\text{–}2020})\times \text{Post}_t$. Standard errors clustered at the project-location level. The leave-one-country-out exercise re-estimates the baseline regression dropping each of the 50 countries in turn; the full table is in Appendix Table A6. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The general pattern of these sample-restriction exercises is that the buffer-level effect on nightlights is broadly robust to the exclusion of any particular subset of countries, including the most-conflict-affected and the most-exposed subsets. This is in marked contrast to the IPC food-security results that we will report in Section 6, where the headline effect is more concentrated in conflict-prone and chronically vulnerable countries. The implication is that the local economic activity that the buffer-level analysis detects is a relatively uniform feature of USAID's African operation, while the food-security effect is concentrated in the populations and contexts where USAID's withdrawal pushes already-fragile food systems past their margin.

5.7 Placebo tests

We close this section by reporting three main placebo exercises that, taken together, identify the post-shock effect we estimate as the USAID-specific local economic effect and not as a generic post-shock pattern in aid-receiving regions. We complement these with an auxiliary donor-substitution check that we describe in full in Appendix A.10.

Placebo I: spatial decay. The first placebo is the spatial decay of the effect across the buffer radii, already described in Section 5.1. The monotonic attenuation of the coefficient from its largest absolute magnitude at the 500 m buffer to a near-zero, statistically insignificant value at the 25 km buffer is itself a placebo test against confounding by influences that operate at broader spatial scales. A confounder that operates at the regional or national scale, such as a country-wide macroeconomic shock, a national electrification program, or a regional drought, would not produce a coefficient that systematically decays with the buffer radius. The observed decay is consistent only with an effect that is genuinely local to the USAID project location and that diffuses to surrounding areas with declining intensity.

Placebo II: time-shifted shock. The second placebo replaces the actual shock date (28 January 2025) with a fictitious shock date two years earlier (1 February 2023) and re-estimates the baseline DiD on a sample restricted to the pre-shock period (January 2021 through January 2025), so that no portion of the post-shock period contaminates the test. Under the country-trends specification, the time-shifted placebo produces a coefficient that is statistically indistinguishable from zero. Table 8 reports the placebo coefficient side by side with the actual-shock coefficient estimated on the full sample under the same specification; the contrast supports the interpretation that the post-shock effect is specific to the timing of the actual shock and is not the artefact of a general timing-unrelated drift in USAID-exposed regions.

Table 8. Time-shifted placebo at the 1 km buffer

Shock date	$\widehat{\beta}$	SE	p	N
Placebo (1 February 2023)	−0.0024	0.0016	0.137	67,326
Actual (1 February 2025)	−0.0047**	0.0020	0.020	153,888

Notes. Both rows estimate the baseline DiD at the 1 km buffer with $\log(\text{avg\_rad}_{i,t} + 0.01)$ as dependent variable and $\log(1 + \overline{D}_i^{2017\text{-}2020}) \times \text{Post}_t$ as treatment, and both include project, country, and month fixed effects plus country-specific linear trends. The placebo row restricts the sample to the pre-actual-shock window (January 2021 – January 2025) and uses a fictitious shock at 1 February 2023; the actual-shock row uses the full sample (January 2017 – April 2026) and reproduces the baseline of Table 2. Standard errors clustered at the project-location level. The placebo coefficient is statistically indistinguishable from zero; the actual-shock coefficient is negative and statistically significant at the 5% level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Placebo III: random points across Africa. The third placebo compares the post-shock evolution of USAID buffers with that of random African buffers that lack any USAID exposure. We generate 123 random points distributed across Africa, conditional on each being at least 50 km from any USAID project location and on landing on land (point-in-polygon test against Natural Earth country boundaries); the points span 31 African

countries. We extract VIIRS statistics at 1 km and 5 km buffers around each random point using the same gas-flare mask, monthly composites, and threshold conventions as for the USAID buffers, and pool the two samples to estimate a difference-in-differences in which a treated_i indicator (1 for USAID buffers, 0 for random buffers) is interacted with the post-shock dummy. The coefficient on $\text{treated}_i \times \text{Post}_t$ identifies the USAID-specific component of the post-shock decline, after netting out any continent-wide pattern that random and USAID buffers share in common. Identification of a USAID-specific effect requires this coefficient to be significantly negative; a null coefficient would indicate that USAID buffers behave indistinguishably from comparable African locations, contradicting our headline identification.

The pooled DiD is estimated on the 29 African countries in which both USAID project sites and random points are present. At the 1 km buffer, the coefficient on $\text{treated}_i \times \text{Post}_t$ is negative and statistically significant at the 5% level, confirming the USAID-specific post-shock decline. At the 5 km buffer the coefficient is negative but not statistically distinguishable from zero, consistent with the spatial-decay pattern documented in Section 5.1. Full results, including the regression equation, the sample composition, and the coefficient table, are reported in Appendix A.10 (Table A8).

Auxiliary check: alternative donors. We also conducted a donor-substitution placebo at the ADM1 scale, replicating the baseline DiD specification with disbursements from major non-US bilateral and multilateral donors (France, the United Kingdom, Germany, the World Bank, and Chinese commitments) that GODAD makes available for all 23 donors in its panel. The exercise is informative but only partly conclusive, because some non-US donors are themselves directly or indirectly affected by US policy: the World Bank relies substantially on US contributions to IBRD and IDA, and Germany announced a contemporaneous reduction in bilateral ODA in its 2025 federal budget. A significant negative coefficient on these donors therefore cannot be cleanly interpreted as a placebo failure. France and Chinese commitments yield null estimates; the United Kingdom is only marginally significant; Germany and the World Bank produce significant negative coefficients consistent with the donor-interdependence channels just described. We report the full results, with the interpretive caveats, in Appendix A.10. The random-points placebo described above is methodologically cleaner precisely because it sidesteps these donor-interdependence complications by comparing USAID-exposed buffers directly to non-treated African locations rather than to other-donor disbursements.

6 Results — Part II: Acute food insecurity

This section reports our second main set of results, on the effect of the USAID shutdown on acute food insecurity as measured by IPC classifications. We organise the discussion in

seven subsections that broadly parallel those of Section 5, with adjustments that reflect the differences between the two outcomes and units of analysis. Section 6.1 presents the baseline DiD estimates. Section 6.2 characterises the dynamic structure of the effect using an event study. Section 6.3 decomposes the effect by USAID sector. Section 6.4 examines heterogeneity by pre-shock institutional quality, following the strategy of Rohner et al. (2026). Section 6.5 examines heterogeneity by pre-shock baseline vulnerability. Section 6.6 reports sample-restriction and leave-one-country-out exercises, and Section 6.7 reports placebo tests.

6.1 Baseline difference-in-differences

Table 9 reports the baseline DiD estimates from equation (2) at the IPC Level 1 unit, with two outcome variables (the share of population in IPC Phase 3+ and Phase 4+) and two treatment variables (the log of total USAID disbursement and the log of USAID disbursement in the humanitarian sector specifically). The decision to anticipate the humanitarian-sector treatment in the headline table reflects the result, developed in detail in Section 6.3, that the food-security response is concentrated in this single sectoral channel; the full sectoral decomposition into social, economic, productive, and humanitarian exposures is reported there. The four columns are arranged so that the comparisons of interest are visible at a glance.

The pattern in Table 9 is informative on two dimensions. First, on the outcome dimension: the Phase 4+ outcome (population in *Emergency* or worse, columns 3–4) yields a statistically significant positive coefficient in both treatment specifications, while the Phase 3+ outcome (population in *Crisis* or worse, columns 1–2) shows a significant positive coefficient only when the treatment is restricted to the humanitarian sector (column 2). One reading of this pattern is that the post-shock deterioration is more readily detectable in the share of the population entering the more extreme food-security categories, particularly Phase 4+ Emergency, than in the broader Phase 3+ category, which by construction contains both populations close to crisis and populations already deep in crisis and is therefore mechanically less sensitive to marginal post-shock variation.

Second, on the treatment dimension: the humanitarian-sector treatment produces a substantially larger and more precisely estimated coefficient than the total-disbursement treatment for the Phase 3+ outcome, and a coefficient of comparable magnitude but tighter precision for the Phase 4+ outcome. This pattern is consistent with the humanitarian-sector channel, i.e., the cessation of food aid, emergency response, and BHA/Food for Peace operations, being the dominant driver of the food-security response to the shock.

At the headline specification (Phase 4+ outcome, humanitarian-sector treatment, column 4), a one-log-unit increase in pre-shock humanitarian-sector USAID disbursement is associated with a small but statistically significant increase in the share of post-shock

Table 9. Baseline DiD on IPC food insecurity

Outcome:	(1) % Phase 3+	(2) % Phase 3+	(3) % Phase 4+	(4) % Phase 4+
$\log(\text{USAID})_i \times \text{Post}_t$	0.0008		0.0010***	
	(0.0007)		(0.0003)	
$\log(\text{USAID}^{\text{hum}})_i \times \text{Post}_t$		0.0017**		0.0010***
		(0.0007)		(0.0004)
Observations	2,800	2,800	2,800	2,800
IPC-unit clusters	388	388	388	388
Unit FE	✓	✓	✓	✓
Country FE	✓	✓	✓	✓
Month FE	✓	✓	✓	✓
Country-specific trends	✓	✓	✓	✓

Notes. Dependent variable is the share of the IPC unit population classified in Phase 3+ (*Crisis or worse*, columns 1–2) or in Phase 4+ (*Emergency or worse*, columns 3–4) during the analysis's current period. Treatment is the log of average annual USAID disbursement to IPC unit i during 2017–2020, either total (columns 1, 3) or restricted to the humanitarian sector (columns 2, 4) under the OECD-DAC standard classification described in Section 2 (CRS codes 5xx and 7xx: food assistance, emergency response, disaster prevention). The interaction is taken with Post_t. The primary specification includes IPC-unit fixed effects, country fixed effects, month fixed effects, and country-specific linear trends. Sample restricted to the 27 African countries where IPC operates; Madagascar and Kenya are excluded because of irreconcilable boundary mismatches with GADM 3.6. Standard errors clustered at the IPC-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

population in Phase 4+ or worse. Multiplied across the within-sample interquartile range of humanitarian-sector log disbursement, the implied effect is on the order of a one-percentage-point increase in Phase 4+ population share at units in the 75th percentile of humanitarian-sector exposure relative to units in the 25th percentile. In a typical African country in our sample, this corresponds to tens of thousands of additional individuals classified in Emergency or worse for each IPC unit at the upper end of the exposure distribution.

6.2 Dynamic effects: event study with escalation pattern

Figure 9 reports event-study estimates from equation (3) for four combinations of outcome (Phase 3+ and Phase 4+) and treatment (total USAID disbursement and humanitarian-sector USAID disbursement). The bins are defined in semesters relative to the shock, reflecting the typical IPC analysis cadence of two analyses per country per year.

Several features of the event study deserve attention. First, in three of the four panels the post-shock coefficient at $k = +1$ (covering August 2025 to January 2026) is *larger* than the coefficient at $k = 0$ (covering February to July 2025). The escalation pattern, in which the effect builds in the second post-shock semester rather than peaking in the first, is consistent with the operational dynamics through which the cessation

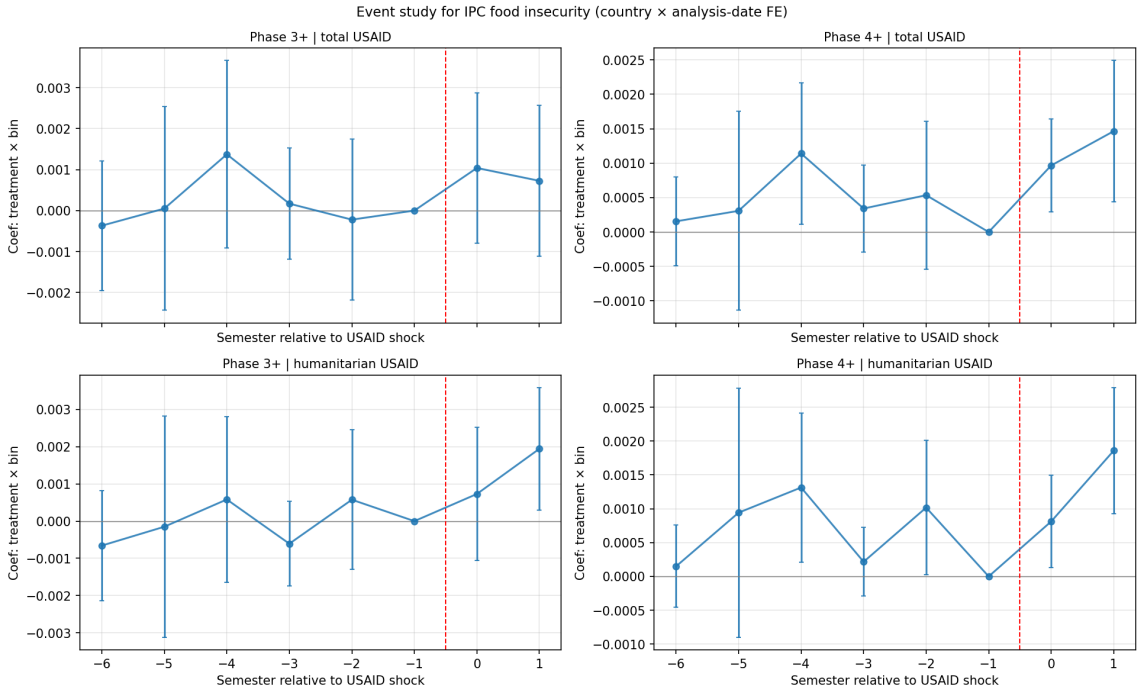

Figure 9. Event study for IPC food insecurity, four specifications

Notes. The figure plots the event-study coefficients from equation (3) on IPC food-insecurity outcomes, in four panels combining two outcomes (Phase 3+ and Phase 4+) with two treatment variables (total USAID disbursement, humanitarian-sector USAID disbursement under the OECD-DAC standard classification: CRS codes 5xx and 7xx, food assistance and emergency response). The bins are in semesters relative to the USAID shock (semester 0 = February–July 2025); $k = -1$ is the reference period. Bars represent 95% confidence intervals; standard errors are clustered at the IPC-unit level.

of humanitarian assistance transmits to household food consumption. The immediate cessation of new shipments and disbursements in February 2025 did not produce an instantaneous collapse of beneficiary food access, because pre-existing commodity stocks held by USAID implementing partners (in-country warehouses, regional pre-positioning facilities, and ongoing distribution rounds funded by prior obligations) continued to be drawn down through the first semester after the shock. Reports from BHA and Food for Peace implementing partners describe a progressive exhaustion of these stocks over the months following the stop-work order, with the resulting interruption of beneficiary distributions producing a measurable increase in acute food insecurity that builds up over the second and third post-shock semesters. The escalation between $k = 0$ and $k = +1$ is consistent with the time-window over which this commodity-pipeline depletion would be expected to materialise in IPC classifications, which themselves operate on a multi-month cycle of evidence gathering and consensus assessment.

Second, the pre-shock coefficients in the bottom-right panel (Phase 4+ outcome, humanitarian-sector treatment) exhibit one statistically significant deviation from zero at $k = -4$ (February–July 2023). This single-bin deviation coincides temporally with the onset of the Sudanese civil war in April 2023, an event that disproportionately affected

several countries in our sample (notably Sudan itself, but also through spillover effects on neighbouring Chad and South Sudan, which jointly host a substantial share of high-humanitarian-sector-exposure IPC units). The pattern is consistent with an interpretation under which the conflict produced a localised increase in Phase 4+ classifications in high-exposure IPC units relative to low-exposure ones for that single semester, although we cannot identify the causal contribution of the conflict directly. The event-study coefficients in the immediately adjacent semesters ($k = -3$, $k = -5$, $k = -6$) are not significantly different from zero, suggesting that the $k = -4$ deviation is an isolated episode rather than a continuous pre-shock trend. We address this episode in two ways. First, the country-trends robustness specification reported below in Figure 10 absorbs country-specific linear trends; under that specification the $k = -4$ coefficient becomes statistically indistinguishable from zero while the post-shock coefficients remain qualitatively similar. Second, the sample-restriction exercise of Section 6.6 drops Sudan, South Sudan, the Central African Republic, and three other active-conflict countries; the post-shock effect survives the restriction, though attenuated, which is consistent with the bulk of the post-shock evidence not being driven by Sudan-related dynamics.

Robustness of the event study to alternative trend specifications. Figure 10 reports the event-study coefficients under three alternative fixed-effect structures: the baseline (unit + country × date FE), the country-trends specification (unit + country + date + country trends), and a unit-specific-trends specification in the style of Goodman-Bacon (2021). The three specifications produce qualitatively similar event-study profiles. The $k = -4$ deviation that is visible under the baseline specification is substantially attenuated under both alternative trend specifications. The post-shock coefficients at $k = 0$ and $k = +1$ remain positive and broadly significant under all three specifications, with the escalation pattern (a larger coefficient at $k = +1$ than at $k = 0$) preserved across specifications.

6.3 Sectoral channels

The baseline result in Section 6.1 previewed the relative importance of the humanitarian sector in driving the food-security effect. Table 10 reports the formal sectoral decomposition for the Phase 3+ outcome, with each OECD-DAC sector entered separately (columns 1–4) and all four sectors entered jointly (column 5).

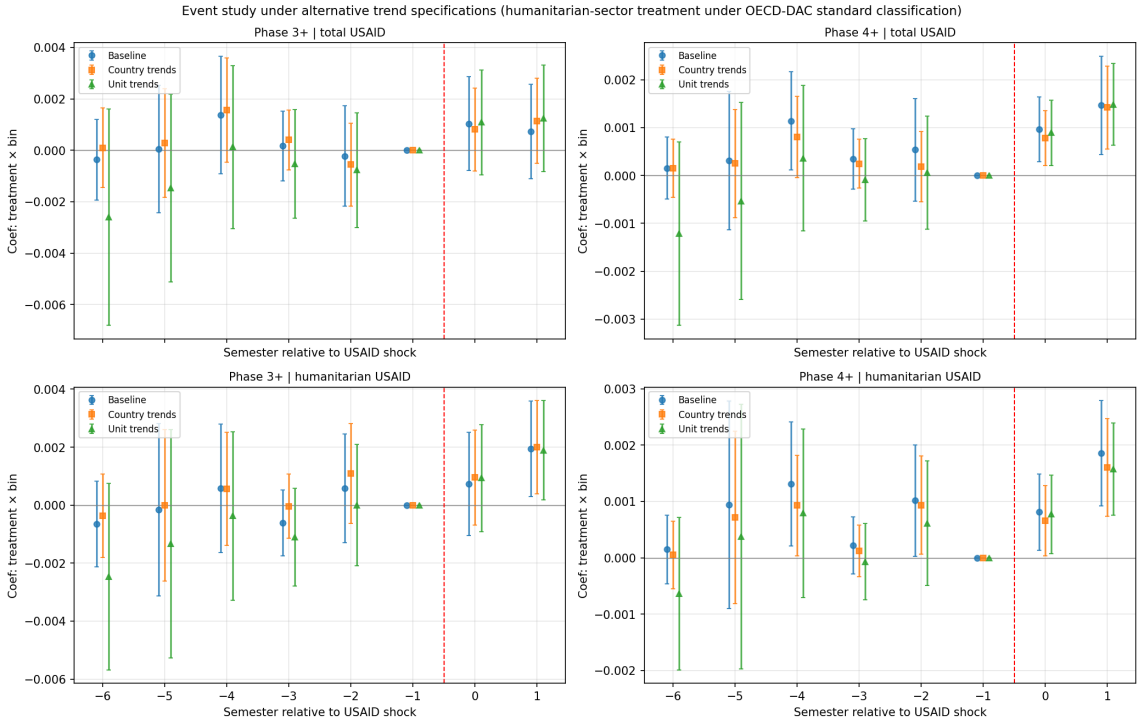

Figure 10. IPC event study under alternative trend specifications

Notes. The four panels report event-study coefficients for the IPC food-security outcome under three alternative fixed-effect structures, for the two outcome variables (Phase 3+ and Phase 4+) combined with two treatment variables (total USAID disbursement and humanitarian-sector USAID disbursement under the OECD-DAC standard classification). The baseline specification (blue) includes IPC-unit fixed effects and country × analysis-date fixed effects. The country-trends specification (orange) replaces the country × date effects with country and date effects plus country-specific linear trends. The unit-trends specification (green) retains country × date effects and adds IPC-unit-specific linear trends (Goodman-Bacon, 2021). Bars represent 95% confidence intervals; standard errors are clustered at the IPC-unit level. The pattern of post-shock escalation between $k = 0$ and $k = +1$ is preserved under all three specifications.

Three patterns are immediately apparent. First, the humanitarian sector dominates the joint regression: its coefficient is the largest by absolute magnitude among the four sectors and is the only one significant at the 5% level. Second, the productive sector, which captures agricultural and industrial activities, including Feed the Future, enters with a coefficient statistically indistinguishable from zero in both the standalone and the joint specifications, with point estimates close to zero. Third, the social and economic sectors also enter with statistically null coefficients.

The sectoral estimates indicate that the effect of the USAID shutdown on acute food insecurity is concentrated in the humanitarian-aid component of the agency's portfolio: the Bureau for Humanitarian Assistance, Food for Peace, and the OFDA emergency-response operations. This is the channel most plausibly linking the shock to household food consumption: the immediate cessation of food-aid shipments and emergency cash transfers translates, with the lags imposed by the depletion of in-country commodity stocks and the cancellation of subsequent rounds of assistance, into measurable post-shock increases in the share of population in IPC Phase 4+ Emergency or worse.

Table 10. Sectoral heterogeneity for IPC Phase 3+ outcome

	(1)	(2)	(3)	(4)	(5)
$\log(\text{USAID}^{\text{social}})_i \times \text{Post}_t$	0.0000				−0.0003
	(0.0007)				(0.0007)
$\log(\text{USAID}^{\text{economic}})_i \times \text{Post}_t$		−0.0012			−0.0010
		(0.0008)			(0.0008)
$\log(\text{USAID}^{\text{productive}})_i \times \text{Post}_t$			0.0003		0.0003
			(0.0007)		(0.0008)
$\log(\text{USAID}^{\text{humanitarian}})_i \times \text{Post}_t$				0.0019**	0.0018**
				(0.0007)	(0.0008)
Observations	2,795	2,795	2,795	2,795	2,795
IPC-unit clusters	388	388	388	388	388
IPC-unit FE	✓	✓	✓	✓	✓
Country × analysis-date FE	✓	✓	✓	✓	✓

Notes. Dependent variable is the share of IPC unit population classified in Phase 3+ (Crisis or worse) during the analysis's current period. Sectoral exposure variables follow the OECD-DAC standard CRS classification described in Section 2. Each column reports a separate regression: columns (1)–(4) include only the indicated sector's exposure; column (5) enters all four sectors jointly. The fixed-effect structure (IPC-unit and country × analysis-date) is more saturated than the country-specific linear trends used in the baseline IPC specification of Table 9, reflecting the greater identification demands of the joint four-sector decomposition. Standard errors clustered at the IPC-unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The contrast between the IPC sectoral pattern and the buffer-level sectoral pattern is informative. In the buffer-level analysis (Section 5.3, Table 4), *both* the productive and the humanitarian sectors drive the post-shock effect, with coefficients of comparable magnitude. In the IPC analysis, only the humanitarian sector is detectable; the productive sector is empirically null. Two properties of the two measures account for this difference. First, the spatial scale differs sharply: productive-sector USAID activities (agricultural extension stations, seed-system support, irrigation infrastructure) operate at a fine geographic resolution that the 1 km buffer captures directly, but that the ADM1-scale IPC averaging substantially dilutes. Second, humanitarian disbursements affect household food consumption through a direct channel: when BHA and Food for Peace shipments stop, the beneficiary population's food intake is directly reduced. Productive-sector disbursements, by contrast, affect food consumption indirectly through the agricultural cycle and with substantial lags that may exceed our fifteen-month observation window. The two outcomes therefore capture different but complementary dimensions of the post-shock effect: the buffer-level nightlight measure may pick up both the operational footprint of productive-sector activities and the more immediate humanitarian-cessation channel, while the IPC measure primarily captures the direct humanitarian channel transmitting through household food consumption.

6.4 Heterogeneity by institutional environment

Following Rohner et al. (2026), who document that the effect of the USAID shutdown on conflict is concentrated in countries with weaker political institutions, we examine whether the same institutional moderation applies to the food-security effect. We define a binary dummy $HighInst_c$ equal to one if the country's average institutional-quality score over 2017–2020 is above the within-Africa median, using three alternative measures of institutional quality: the Polity 5 *polity2* index, the V-Dem v15 liberal-democracy index *v2x_libdem*, and the V-Dem judicial-constraints index *v2x_jucon*. We then estimate the moderation specification

$$Y_{i,c,t} = \beta_1 \text{USAID}_i^{\text{hum}} \cdot \text{Post}_t + \beta_2 \text{USAID}_i^{\text{hum}} \cdot \text{Post}_t \cdot HighInst_c + \alpha_i + \gamma_{c,t} + \varepsilon_{i,c,t}. \quad (5)$$

The coefficient β_1 measures the effect of humanitarian-sector USAID exposure in low-institution countries; the sum $\beta_1 + \beta_2$ measures the effect in high-institution countries; and β_2 alone measures the differential moderation by institutional quality.

Table 11 reports the estimates separately for each of the three measures of institutional quality. The pattern is striking and uniform across the three measures. The main effect $\widehat{\beta_1}$ is positive and significant at the 1% level in all three columns, confirming that the humanitarian-sector effect on Phase 4+ food insecurity is concentrated in low-institution countries. The interaction $\widehat{\beta_2}$ is negative and statistically significant in all three columns, indicating that the effect is substantially weaker in countries with stronger institutions. Under the V-Dem liberal-democracy and judicial-constraints indices, the implied effect in high-institutional-quality countries is approximately 7% of the corresponding low-institutional-quality effect. Under the Polity 5 measure, it is approximately one-quarter of the low-institutional-quality effect.

Countries with stronger executive constraints, judicial independence, and democratic institutions appear better positioned to absorb the USAID shock and to mobilise compensatory domestic policy responses, with the marginal effect of the shutdown on acute food insecurity attenuated in those settings. This pattern echoes the analogous finding in Rohner et al. (2026) on the conflict outcome, and is consistent with a broader literature on the role of institutions in moderating the local consequences of external shocks. The magnitude of the moderation in our setting is substantial: the humanitarian-sector food-security effect estimated in Table 9 is, in effect, the average of a sizeable effect concentrated in low-institution countries and an effect close to zero in high-institution ones. The populations most exposed to the consequences of the USAID withdrawal, then, tend to be those in countries with the weakest domestic capacity to compensate.

Table 11. Institutional heterogeneity: IPC Phase 4+ effect by political-institution moderators

	Polity 5	V-Dem libdem	V-Dem jucon
$T_i \times \text{Post}_t$ (low-institution reference)	0.0019***	0.0016***	0.0016***
	(0.0007)	(0.0005)	(0.0005)
$T_i \times \text{Post}_t \times \textit{HighInst}_c$	−0.0015*	−0.0015***	−0.0015***
	(0.0008)	(0.0006)	(0.0006)
Implied effect in high-institution countries	0.0004	0.0001	0.0001
Observations	2,795	2,795	2,795

Notes. Dependent variable is the share of IPC unit population in Phase 4+. $T_i \equiv \log(\text{USAID}^{\text{hum}})_i$ denotes the log of humanitarian-sector USAID disbursement at IPC unit i (CRS codes 5xx and 7xx under OECD-DAC standard classification). Treatment is $T_i \times \text{Post}_t$, interacted with a within-Africa median split of the indicated institutional moderator (Polity 5 *polity2* score, V-Dem v15 *v2x_libdem* index, V-Dem *v2x_jucon* index of judicial constraints, all averaged over 2017–2020). All specifications include unit fixed effects and country × analysis-date fixed effects. Standard errors clustered at the IPC-unit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.5 Heterogeneity by baseline vulnerability

A complementary analysis splits the sample by pre-shock baseline vulnerability rather than by country-level institutional quality. For each IPC unit we compute the average share of population in Phase 3+ over the pre-shock period (2017–2024). We then split the sample at the within-sample median (approximately 14.7% Phase 3+ population pre-shock) into "high-vulnerability" units (above median) and "low-vulnerability" units (at or below median), and re-estimate the baseline DiD separately on each sub-sample.

The results in Table 12 are stark. In high-vulnerability IPC units, the humanitarian-sector effect on Phase 4+ is positive and significant at the 1% level, almost identical in magnitude to the all-sample headline coefficient. In low-vulnerability units, the same coefficient is essentially zero. The pooled regression with an interaction term confirms the formal moderation: the interaction coefficient is positive and significant at the 1% level.

The estimates align with the prediction anticipated in the introduction and in Section 4.9: the shutdown does not push food-secure populations into emergency; rather, it shifts populations already at the margin of food crisis across that threshold. The food-security effect of the shutdown is therefore concentrated in the populations and the regions whose food-security baseline was already fragile pre-shock. In high-vulnerability IPC units, the humanitarian-sector channel of the shock, namely the cessation of food assistance, cash transfers, and emergency response delivered through BHA and Food for Peace, transmits to the share of population in Phase 4+ Emergency or worse. In low-vulnerability units, the same channel operates but the populations supported by USAID activities are not close enough to the Emergency threshold for the shock to move them across it.

Table 12. Baseline-vulnerability heterogeneity: IPC Phase 4+ effect

	Coefficient	N obs.
High-vulnerability units (pre-shock % Phase 3+ above sample median)	0.0018***	1,287
	(0.0007)	
Low-vulnerability units (pre-shock % Phase 3+ at or below median)	0.0000	1,471
	(0.0001)	
Interaction term ($HighVuln \times \log(\text{USAID}^{\text{hum}})_i \times \text{Post}_t$, pooled regression)	0.0013***	2,795
	(0.0004)	

Notes. Dependent variable is the share of IPC unit population in Phase 4+. The two top rows report separate regressions on subsamples split at the within-sample median of pre-shock (2017–2020 average) % Phase 3+. The bottom row reports the pooled regression with the interaction term. Treatment is $\log(\text{USAID}^{\text{hum}})_i$, the log of humanitarian-sector USAID disbursement (CRS codes 5xx and 7xx). All specifications include unit FE and country × analysis-date FE; standard errors clustered at the IPC-unit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The interaction with the institutional quality of Section 6.4 is informative. The two dimensions of heterogeneity are partially overlapping (countries with weak institutions also tend to host populations with higher baseline vulnerability), but capture different aspects of the underlying mechanism. The heterogeneity by institutional quality is consistent with the absence of a domestic policy response that could compensate for the shock; the heterogeneity by baseline vulnerability captures the immediate proximity of the affected population to the Emergency threshold. Both dimensions point in the same direction: the effect of the shutdown is concentrated where compensating mechanisms are weakest and where the marginal effect on the food-insecure population is largest.

6.6 Sample restrictions and leave-one-country-out

Table 13 reports the headline coefficient on Phase 4+ food insecurity under several sample restrictions, with the upper panel using the total-USAID treatment and the lower panel using the humanitarian-sector treatment. The pattern is mixed, in a manner that is informative.

Dropping the six countries with active armed conflict (Sudan, South Sudan, the Central African Republic, Mali, Niger, Burkina Faso) attenuates the total-USAID coefficient substantially, while preserving statistical significance at the 5% level. Dropping five chronic-crisis countries (Somalia, Sudan, South Sudan, the Central African Republic, the Democratic Republic of the Congo) attenuates the coefficient further, also remaining significant at the 5% level. In both cases the magnitude is approximately one-third of the full-sample headline.

In contrast, the humanitarian-sector coefficient (lower panel) is much more robust to less restrictive sample drops. Dropping the top-exposed country (Mozambique) leaves the coefficient essentially unchanged and significant at the 1% level; dropping the top one

percent of USAID-exposed units yields a similar coefficient, also significant at the 1% level; restricting to countries with at least two post-shock observations leaves the coefficient essentially unchanged at the 1% level. The humanitarian-sector coefficient is, however, sensitive to dropping conflict and chronic-crisis countries, where the coefficient attenuates and loses statistical significance. This result is consistent with humanitarian-aid operations being most relevant where pre-existing crises concentrate.

Table 13. Robustness of the IPC food-insecurity effect to sample restrictions

Restriction	β	SE	p	N obs.
Phase 4+ outcome, log USAID total × Post				
R1 Full sample	0.0010***	0.0003	0.0061	2,800
R2 Drop active-conflict (SDN, SSD, CAF, MLI, NER, BFA)	0.0004**	0.0002	0.0121	2,101
R3 Drop chronic crisis (SOM, SDN, SSD, CAF, COD)	0.0003**	0.0001	0.0180	1,746
R4 Drop top-exposed country (MOZ)	0.0009**	0.0003	0.0111	2,697
R5 Drop low-power countries (1 post-shock obs)	0.0010***	0.0003	0.0062	2,478
R7 Drop top-1% USAID-exposed units	0.0010***	0.0004	0.0063	2,782
Phase 4+ outcome, log USAID humanitarian × Post				
R1 Full sample	0.0010***	0.0004	0.0044	2,800
R2 Drop active-conflict	0.0002	0.0002	0.4174	2,101
R3 Drop chronic crisis	0.0002	0.0001	0.2016	1,746
R4 Drop top-exposed country (MOZ)	0.0011***	0.0004	0.0033	2,697
R5 Drop low-power countries	0.0010***	0.0004	0.0045	2,478
R7 Drop top-1% USAID-exposed units	0.0010***	0.0004	0.0046	2,782

Notes. Each row reports the coefficient on the indicated USAID exposure × Post interaction, estimated under the primary specification (IPC unit FE, country FE, month FE, country-specific linear trends), with the indicated sample restriction applied. The humanitarian-sector treatment is based on the OECD-DAC standard classification of CRS codes (5xx food assistance + 7xx emergency response). Standard errors clustered at the IPC-unit level. The leave-one-country-out exercise, reported in Appendix Table A7, finds that the result holds in every subsample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The total-USAID coefficient appears to be partially driven by the most-conflict-affected and chronic-crisis countries, consistent with the literature on the differential vulnerability of food-insecure populations in those contexts. The humanitarian-sector coefficient is robust to less restrictive sample drops (top-exposed country, low-power countries, top-1% USAID-exposed units), but is sensitive to the drop of conflict and chronic-crisis countries. This is consistent with humanitarian-aid operations being concentrated, and most operationally consequential, precisely in the countries where pre-existing crises generate the highest demand for emergency assistance.

The leave-one-country-out exercise (Appendix Table A7) is consistent with the sample-restriction patterns. The full-sample Phase 4+ × total-USAID coefficient remains positive in all nineteen subsets, with the smallest absolute magnitude obtained when Sudan, the

country whose exclusion most affects the estimate, is dropped and the largest when smaller countries are dropped. All nineteen subsets produce a coefficient that is statistically significant at the 5% level. Statistical significance is therefore robust to the exclusion of any individual country, although the point estimate magnitude varies non-trivially with the inclusion of Sudan.

6.7 Placebo tests

We close the IPC analysis by reporting two placebo tests designed to identify the post-shock effect as USAID-shutdown-specific rather than as a generic post-shock pattern in aid-receiving regions.

Placebo I: alternative donors. Replicating the IPC baseline specification with the historical exposure to disbursements from other major bilateral donors yields the placebo pattern reported in Table 14. France, the United Kingdom, Germany, and Chinese commitments all yield coefficients statistically indistinguishable from zero. The World Bank coefficient is positive and marginally significant at the 10% level, a result plausibly linked to the substantial U.S. presence in multilateral funding (approximately 17% of IBRD financial commitments and 19.6% of cumulative IDA contributions; Congressional Research Service, 2024), which makes World Bank operations difficult to treat as cleanly independent of U.S. policy decisions. Overall, four of the five donor placebos produce null estimates and the fifth (World Bank) is subject to the documented dependence on U.S. contributions just described; the pattern is therefore consistent with the post-shock food-security effect being USAID-specific rather than a generic consequence of donor presence in vulnerable regions. The IPC sample (27 African countries) is smaller than the buffer sample (50 African countries), which reduces the statistical power of the placebo tests; the marginal World Bank coefficient should be interpreted with this caveat in mind.

Table 14. IPC donor-placebo: Phase 4+ DiD with non-US donor exposures

Donor	$\widehat{\beta}$ (T × Post)	SE	p	N
FRA	0.00024	(0.00028)	0.3885	2,478
GBR	0.00023	(0.00036)	0.5285	2,478
GER	0.00004	(0.00041)	0.9288	2,478
WB	0.00139*	(0.00082)	0.0919	2,478
CHN	−0.00001	(0.00021)	0.9499	2,478

Placebo II: time-shifted shock. The second placebo replaces the actual shock date (28 January 2025) with a fictitious shock date two years earlier (1 February 2023) and re-estimates the baseline DiD on a sample restricted to the pre-shock period (January 2021 through January 2025), so that no portion of the actual post-shock period contaminates the

test. Under the country-trends specification, the time-shifted placebo produces coefficients that are statistically indistinguishable from zero across all four baseline specifications, Phase 3+ and Phase 4+ outcomes, total-USAID and humanitarian-sector treatments, and that are several times smaller than the corresponding actual-shock coefficients. The full results, reported in Appendix Table A9, show the placebo coefficients side by side with the actual-shock coefficients estimated on the full sample under the same country-trends specification. The contrast between the two columns is consistent with the post-shock effect on acute food insecurity being specific to the timing of the actual shock rather than the artefact of a pre-existing differential trend in USAID-exposed IPC units.

Overall, the two placebo tests, in combination with the alternative trend specifications discussed in Section 6.2 (Figure 10) and the analogous patterns documented in the buffer-level analyses, are consistent with the post-shock effect on acute food insecurity at the IPC-unit level being a USAID-shutdown-specific phenomenon, transmitting through the humanitarian-assistance channel of the disrupted portfolio rather than a generic consequence of donor presence in vulnerable regions or a generic post-shock dynamic in fragile contexts.

7 Discussion

The two outcomes analysed in this paper, local economic activity and acute food insecurity, capture substantively distinct phenomena and are constructed from entirely independent data sources, sharing no overlap in their measurement infrastructure. Any common pattern across them is therefore a substantive finding rather than a mechanical artefact of the data.

A striking feature of our results is that, when each of the two outcomes is decomposed into the four OECD-DAC macro-sectors of the USAID portfolio, both outcomes are driven by the humanitarian sector (food assistance, emergency response, and the BHA/Food for Peace operational complex). The convergence on this single sectoral channel is the central piece of substantive evidence in the paper. Table 15 reports the sectoral coefficients of the joint regression specifications from Tables 4 (local economic activity) and 10 (food insecurity), side by side.

The convergence on the humanitarian sector is reinforced by a complementary finding at the buffer-level outcome: the productive sector (agriculture, forestry, fishing, industry, including Feed the Future) drives a significant component of the nightlight decline at the project-buffer level, while showing no measurable effect on the IPC food-security outcome. This asymmetry is itself interpretable. Productive-sector activities, i.e., agricultural extension stations, seed multiplication centres, irrigation infrastructure, market hubs, operate at a fine geographic resolution that the 1 km buffer captures directly, where

they sustain a localised operational footprint visible to satellite nightlight measurement. At the ADM1 spatial scale at which IPC classifications operate, this productive-sector footprint is geographically too dispersed and population-wise too concentrated to move aggregate food-insecurity classifications in the fifteen-month window we observe. The IPC measure, by contrast, is sensitive to direct food-consumption channels that operate at the population scale, and the cessation of BHA and Food for Peace commodity flows is precisely such a channel.

Table 15. Convergent sectoral channels: joint regression coefficients across the two outcomes

	VIIRS nightlights	IPC food insecurity	
USAID sector (× Post)	1 km buffer outcome	Phase 3+ outcome	Phase 4+ outcome
Social (PEPFAR, health, education, governance)	+0.0006	−0.0003	+0.0006
	(0.0017)	(0.0007)	(0.0004)
Economic (transport, energy, banking)	+0.0009	−0.0010	−0.0001
	(0.0030)	(0.0008)	(0.0002)
Productive (agriculture, Feed the Future)	−0.0043**	+0.0003	−0.0001
	(0.0021)	(0.0008)	(0.0003)
Humanitarian (BHA, Food for Peace, OFDA)	−0.0051**	+0.0018**	+0.0011***
	(0.0023)	(0.0007)	(0.0004)

Notes. The table reports the coefficients on sector-specific log-USAID-disbursement interacted with the post-shock indicator, from joint regressions in which all four OECD-DAC standard macro-sectors are entered simultaneously. The dependent variables are $\log(\text{avg\_rad}_{i,t} + 0.01)$ in the 1 km buffer for the nightlights outcome (column 1; from Table 4, column 5), and the share of population in IPC Phase 3+ for the food-security outcome (column 2; from Table 10, column 5); column 3 reports the corresponding Phase 4+ joint regression. Sectoral exposure variables are constructed by classifying each project-location record by the first digit of its OECD-CRS purpose code (Section 2). Note that the signs differ between columns 1 and columns 2–3 because the two outcomes go in opposite directions when affected by the shock: nightlights *decline* (negative coefficient) and food insecurity *rises* (positive coefficient). Standard errors clustered at the unit level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

By narrowing the space of plausible confounders to those that would have to operate at the level of specific OECD-DAC sectors jointly across two independent measurement infrastructures, the sectoral convergence is difficult to reconcile with a generic "USAID presence" alternative explanation. A confounder correlated with total USAID disbursement would produce coefficients across the four sectors roughly proportional to dollar share; instead, the largest sector (social, 52.9% of total disbursements) shows no effect on either outcome, while the humanitarian sector (34.4% of total) drives both outcomes, and the productive sector (6.9% of total) drives the nightlight outcome only. The asymmetric pattern is consistent with two distinct operational mechanisms transmitting the shock — one fine-grained and visible only at the buffer scale (productive-sector operations), one broader and visible on both the buffer and population-aggregate scales (humanitarian food and cash assistance) — rather than with a single generic correlate of USAID exposure.

The magnitudes we estimate imply a reduction of measured nighttime activity on the order of several percentage points in the immediate vicinities of high-USAID-exposure

project locations, and an aggregate additional burden of acute food insecurity (Phase 4+) on the order of one to several million individuals across the 19 countries with both pre- and post-shock IPC observations. The detailed IQR-based calculations underlying these orders of magnitude are reported in Appendix A.12.

To put these effects on an intuitive scale, we apply standard conversion factors from the lights-to-GDP literature to the buffer coefficient and use IPC population denominators to convert food-security shares into headcounts. On the nightlights side, the conservative 1-km estimate of Section 5.1 implies a post-shock change in log radiance of -0.06 at a project with the sample-mean USAID exposure (annualised disbursements of approximately USD 314,000 in the 2017–2020 window). Applying the GDP-lights elasticity $\psi \approx 0.30$, the round midpoint of Henderson et al. (2012)'s benchmark estimates of 0.277 (world fixed-effects, their Table 2, column 1) and 0.327 (low- and middle-income long-difference, their Table 4, column 3), this maps into a cumulative GDP shortfall of about 1.8 per cent over the 15-month post-shock window, or approximately 1.4 percentage points of annual GDP growth in the affected buffers. For reference, Henderson et al. (2012) report cross-quartile growth gaps in malaria-exposed sub-Saharan Africa of one-third to two-thirds of a percentage point per year. Our implied shortfall is roughly three times larger, which is consistent with the abrupt and concentrated nature of the cessation relative to a gradual cross-area gradient.

On the food-security side, the Phase 4+ coefficient of Section 6.1 implies, for an IPC unit at the sample-mean treated USAID humanitarian exposure (annualised disbursements of approximately USD 293,000 and population of 2.2 million), an additional 1.3 percentage points of population in IPC Phase 4 or higher and an additional 2.1 percentage points in Phase 3 or higher in each post-shock analytical wave. Aggregating across the 110 IPC units that report in the September 2025 analytical wave (covering approximately 400 million people in sub-Saharan Africa), these per-unit effects sum to roughly 1.6 million additional people in Phase 4 (humanitarian emergency) and 2.7 million additional people in Phase 3 or higher (acute food crisis or worse), in that single wave, attributable to the USAID cessation. Across the nine post-shock analytical waves covered by our sample, the cumulative person-wave count exceeds 6 million for Phase 3+ and approaches 4 million for Phase 4+. This figure sums wave-specific additional headcounts rather than tracking distinct individuals across waves: a population classified in the same phase across consecutive analytical waves contributes to each wave's count separately. The cumulative person-wave count should therefore be read as an upper bound on the count of distinct people whose food-security classification is affected by the shock; tighter de-duplication across waves is not supported by the IPC data structure, which is constructed at the unit-wave level rather than at the individual level.

We acknowledge several limitations of the analysis. The treatment variable is constructed from USAID disbursements during 2017–2020, the most recent four years for

which the geocoded CRS is essentially complete in GODAD. Following Rohner et al. (2026), we assume that the spatial distribution of U.S. aid programmes was sufficiently stable between 2020 and the shock to make the 2017–2020 average a good proxy for pre-shock exposure. To the extent that activities shifted spatially in the intervening years, our treatment variable is a noisy proxy and the coefficient is attenuated; the direction of any bias is toward zero, and our estimates are conservative.

Our post-shock window of fifteen months is short relative to the time scale over which several of the effects unfold, particularly for the humanitarian-sector channel that operates with a six- to twelve-month lag through the progressive depletion of in-country commodity stocks following the cessation of new shipments. The IPC event study, with the effect at $k = +1$ materially larger than at $k = 0$, suggests we are observing a partial and still-unfolding effect rather than steady-state consequences; the magnitudes we report should be regarded as lower bounds.

A separate concern arises from the absence of Ethiopia from our IPC analytical sample, because the national Technical Working Group has not published a post-shock IPC analysis for our window. Ethiopia was among the top USAID recipients during 2017–2020, and the humanitarian-sector channel could plausibly be active there; inclusion of Ethiopia would likely have increased rather than decreased the food-security effect we estimate.

Two further IPC-specific caveats — the integrity of IPC measurement in the post-shock period, and the limited scope for clean placebo testing at the IPC unit level — are documented in detail in Sections 4.6 and 6.7 and in the corresponding appendices. The diagnostics there indicate that IPC operational continuity has been preserved across the post-shock window, with any reduction in cadence or coverage concentrated in the least USAID-exposed countries rather than the most; any residual measurement bias would therefore push our IPC estimates toward zero. The donor-substitution placebos produce null estimates for four of five donors, the marginally significant World Bank result being plausibly attributable to U.S.-dependent multilateral funding rather than to a generic donor-substitution effect. The USAID-specific identification of the headline IPC effect remains substantially intact.

8 Conclusion

The sudden shutdown of the United States Agency for International Development on 28 January 2025 represents an unprecedented natural experiment in the history of bilateral development assistance. A development programme that had operated continuously for more than six decades and disbursed several billion dollars per year to Africa alone was terminated within weeks, by political decision of an incoming U.S. administration, without warning or phase-down. This paper has exploited the timing and magnitude

of the shock to estimate its consequences for two outcomes: local economic activity, measured through satellite nighttime light radiance at fine spatial buffers around 1,374 USAID project locations, and acute food insecurity, measured through the Integrated Food Security Phase Classification at 388 subnational analytical units in 27 African countries. Both outcomes exploit within-country variation in pre-shock exposure to USAID under continuous difference-in-differences specifications that are robust to a wide range of alternative fixed-effect structures, sample restrictions, and placebo tests.

The empirical results converge on three central findings. First, the cessation of USAID activities produced a sharp post-shock decline in nighttime light radiance within 500 metres to 5 kilometres of USAID project sites, with the magnitude of the effect attenuating monotonically with the buffer radius and falling below statistical detection at 25 kilometres. The trend-break decomposition characterises the post-shock effect as a discrete level break at the date of the shock, rather than a slope reversal of a pre-existing growth trajectory, consistent with USAID's having sustained a stock-like level of local economic activity whose immediate cessation produced an instantaneous reduction in the visible economic footprint of the affected locations. The result is robust to alternative fixed-effect specifications, to sample restrictions including the exclusion of all active-conflict countries and the leave-one-country-out exercise across the fifty African countries in our sample, and to three independent placebo tests.

Second, in the 27 African countries where IPC analyses are produced, regions historically exposed to USAID experienced relative increases in the share of population classified at Phase 4 (Emergency) or worse after the shutdown. The effect is small in magnitude, but it is statistically robust across alternative fixed-effect specifications, sample restrictions, and donor placebo tests. The event-study evidence reveals an escalation pattern in which the post-shock effect grows over the first year after the shock, consistent with the gradual transmission of disrupted humanitarian-assistance flows to household food consumption through the progressive depletion of in-country commodity stocks held by BHA and Food for Peace implementing partners.

Third, both the nightlight effect and the food-security effect are driven by the humanitarian sector of the USAID portfolio (food assistance, emergency response, and BHA/Food for Peace operations), with the nightlight effect additionally reinforced by a contribution of the productive sector (agriculture, including Feed the Future) that is visible at the 1 km buffer scale but not at the ADM1 scale of the IPC analysis. The much larger social sector (PEPFAR and related health programmes) shows no measurable effect on either outcome in the first fifteen months after the shock. The convergence of two independently measured outcomes on the same humanitarian sectoral channel constitutes a central piece of evidence in the paper. By narrowing the space of plausible confounders to those that would have to operate at the level of a specific OECD-DAC sector jointly across two independent measurement infrastructures, the convergence is

difficult to reconcile with a generic "USAID presence" alternative explanation. It also identifies the specific operational activities, such as emergency food and cash assistance, commodity distribution, and the agricultural-sector activities that maintain a localised operational footprint, through which the local welfare consequences of the shutdown are immediately transmitted.

The findings carry implications both for the academic literature on foreign aid and for the policy debate that the shutdown has generated. For the academic literature, the paper provides the first empirical (as opposed to projected) estimates of the local economic and food-security consequences of an exogenous foreign-aid shock of unprecedented scale and abruptness. Our results provide a natural complement to those of Rohner et al. (2026), who document the conflict consequences of the same shock. Whereas their analysis focuses on the violent-conflict response, our evidence speaks to an intermediate step in the proposed opportunity-cost mechanism: the transmission of the aid cessation to local economic activity (through both humanitarian and productive operational channels at fine spatial resolution) and to household food consumption (through the humanitarian channel specifically). Our findings also complement the projections of Cavalcanti et al. (2025), who use programme-impact estimates to forecast the medium-term mortality consequences of the shutdown, by providing direct evidence of the early-stage welfare consequences that motivate those projections.

For the policy debate, the institutional-quality pattern documented in Section 6.4 suggests that the welfare costs of an abrupt aid cessation are borne disproportionately by populations in countries with weaker political institutions. The marginal effect of the USAID shock on food insecurity is approximately fourteen times larger in countries below the within-Africa median of the V-Dem liberal democracy index than in countries above. This pattern is consistent with the institutional-mitigation hypothesis advanced by Rohner et al. (2026) on the conflict outcome and indicates that the marginal cost of the cut falls most heavily on populations least equipped to mobilise compensatory domestic policy responses.

We close with three caveats and forward-looking observations. First, the post-shock observation window of fifteen months is short relative to the time scale over which several of the relevant economic and humanitarian effects unfold. The point estimates we report should be regarded as lower bounds on the medium-term effects, particularly for the humanitarian-sector channel whose transmission to household food security operates with lags of six to twelve months as in-country commodity stocks held by implementing partners are progressively depleted. Second, several of the most affected populations in Africa are not directly observed in our IPC sample, most notably in Ethiopia, where post-shock IPC analyses have not yet been published. Third, the same political process that produced the USAID shutdown has also produced subsequent secondary cuts to U.S. contributions to multilateral channels (the International Bank for Reconstruction and Development, the

International Development Association, the World Health Organization), and to bilateral aid from other major donors (notably Germany), whose effects unfold on a different time scale and whose welfare consequences will compound those of the USAID shock as the post-shock period lengthens.

Subsequent research will be in a position to estimate, with longer observation windows and with additional outcome measures, the medium-term consequences of what is on track to be the largest-scale negative shock to bilateral development assistance in the history of the modern international aid system. The findings we report here suggest that those consequences will be substantial, concentrated in identifiable populations and operational channels, and unevenly distributed across the African subnational geography that the U.S. bilateral aid programme has supported over the past several decades.

References

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *Quarterly Journal of Economics*, 138(1):1–35.

Andrabi, T. and Das, J. (2017). In aid we trust: Hearts and minds and the Pakistan earthquake of 2005. *Review of Economics and Statistics*, 99(3):371–386.

Bommer, C., Dreher, A., and Perez-Alvarez, M. (2022). Home bias in humanitarian aid: The role of regional favoritism in the allocation of international disaster relief. *Journal of Public Economics*, 208:104604.

Bomprezzi, P., Dreher, A., Fuchs, A., Hailer, T., Kammerlander, A., Kaplan, L. C., Marchesi, S., Masi, T., Perlik, K., and Robert, C. (2024). Wedded to prosperity? spousal favoritism in foreign aid and regional development. CEPR Discussion Paper 18878.

Briggs, R. C. (2017). Does foreign aid target the poorest? *International Organization*, 71(1):187–206.

Cavalcanti, D. M., de Oliveira Ferreira de Sales, L., da Silva, A. F., Landin Basterra, E., Pena, D., Monti, C., Barreix, G., Silva, N. J., Vaz, P., Saute, F., Fanjul, G., Bassat, Q., Naniche, D., Macinko, J., and Rasella, D. (2025). Evaluating the impact of two decades of USAID interventions and projecting the effects of defunding on mortality up to 2030: a retrospective impact evaluation and forecasting analysis. *The Lancet*, 406(10500):283–294.

Congressional Research Service (2023). U.S. assistance for sub-Saharan Africa: An overview. Technical Report R46368, Congressional Research Service, Library of Congress, Washington, DC. Retrieved from https://www.congress.gov/crs-product/R46368.

Congressional Research Service (2024). The World Bank. Technical Report IF11361, Congressional Research Service, Library of Congress, Washington, DC. Retrieved from https://www.congress.gov/crs-product/IF11361.

Congressional Research Service (2025). U.S. Agency for International Development: An overview. Technical Report IF10261, Version 20, Congressional Research Service, Library of Congress, Washington, DC. Retrieved from https://www.congress.gov/crs-product/IF10261.

Coppedge, M., Gerring, J., Knutsen, C. H., Lindberg, S. I., Teorell, J., Altman, D., Angiolillo, F., Bernhard, M., Cornell, A., Fish, M. S., Fox, L., Gastaldi, L., Gjerløw, H., Glynn, A., Good God, A., Grahn, S., Hicken, A., Kinzelbach, K., Krusell, J., Marquardt, K. L., McMann, K., Mechkova, V., Medzihorsky, J., Natsika, N., Neundorf, A., Paxton, P., Pemstein, D., von Römer, J., Seim, B., Sigman, R., Skaaning, S.-E., Staton, J., Sundström, A., Tannenberg, M., Tzelgov, E., Wang, Y.-T., Wiebrecht, F., Wig, T., Wilson, S. L., and Ziblatt, D. (2025). V-Dem [Country-Year/Country-Date] Dataset v15.

Crost, B., Felter, J., and Johnston, P. (2014). Aid under fire: Development projects and civil conflict. *American Economic Review*, 104(6):1833–1856.

Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–198.

Donor Tracker (SEEK Development) (2024). Germany's draft 2025 budget: Downward ODA trends confirmed. Retreived from https://donortracker.org/publications/germanys-draft-2025-budget-downward-oda-trends-confirmed.

Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., and Tierney, M. J. (2019a). African leaders and the geography of China's foreign assistance. *Journal of Development Economics*, 140:44–71.

Dreher, A., Fuchs, A., and Langlotz, S. (2019b). The effects of foreign aid on refugee flows. *European Economic Review*, 112:127–147.

Elvidge, C. D., Baugh, K., Zhizhin, M., Hsu, F.-C., and Ghosh, T. (2017). VIIRS night-time lights. *International Journal of Remote Sensing*, 38(21):5860–5879.

Falchetta, G., Pachauri, S., Parkinson, S., and Byers, E. (2019). A high-resolution gridded dataset to assess electrification in sub-Saharan Africa. *Scientific Data*, 6(1):110.

Fazio, A., Nicolini, M., and Sabatini, F. (2026). Foreign aid and trust in institutions. SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6667678.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028.

Isaksson, A.-S. and Kotsadam, A. (2018). Chinese aid and local corruption. *Journal of Public Economics*, 159:146–159.

Langlotz, S. (2026). Foreign interventions and community cohesion in times of conflict. *Journal of Development Economics*, 182:103751.

Marshall, M. G. and Gurr, T. R. (2020). *Polity5: Political Regime Characteristics and Transitions, 1800–2018*. Center for Systemic Peace, Vienna, VA.

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

Nunn, N. and Qian, N. (2014). US food aid and civil conflict. *American Economic Review*, 104(6):1630–1666.

Rohner, D., Sunde, U., Vanden Eynde, O., Wright, A. L., and Zeng, J.-R. (2026). Aiding peace or conflict? the impact of USAID cuts on violence. *Science*, 392(6799):eaed6802.

The White House (2025). Executive order 14169: Reevaluating and realigning united states foreign aid. Federal Register, Vol. 90, No. 21, pp. 8619–8621. Signed January 20, 2025. Retrieved from https://www.whitehouse.gov/presidential-actions/2025/01/reevaluating-and-realigning-united-states-foreign-aid/.

U.S. Department of State (2024). PEPFAR latest global results factsheet. Technical report, Office of the U.S. Global AIDS Coordinator and Health Diplomacy, U.S. Department of State, Washington, DC. Retrieved from https://www.state.gov/pepfar-latest-global-results-factsheet-dec-2024.

USAID OIG (2024). Audit of USAID's financial statements for fiscal years 2024 and 2023. Technical Report Audit Report 0-000-25-001-C, Office of Inspector General, U.S. Agency for International Development, Washington, DC. Retrieved from https://oig.usaid.gov/sites/default/files/2024-12/0-000-25-001-C.pdf.

A Additional figures and tables

This Appendix collects supplementary material referenced in the main text.

A.1 List of African countries

Our universe consists of the 55 African countries listed in Table A1, identified by their ISO 3166-1 alpha-3 codes. The number of countries that enter each analytical panel varies with the available data: 50 countries appear in the project-buffer panel (those with at least one geocoded USAID project location in GODAD over 2017–2020), 27 in the IPC panel (countries with an active IPC Technical Working Group whose Level 1 units match GADM 3.6 boundaries), and 53 in the ADM1 panel (those matched by the GAUL→GADM 3.6 crosswalk).

Table A1. The 55 African countries

ISO3	Country	ISO3	Country	ISO3	Country
DZA	Algeria	GIN	Guinea	RWA	Rwanda
AGO	Angola	GNB	Guinea-Bissau	STP	São Tomé and Príncipe
BEN	Benin	KEN	Kenya	SEN	Senegal
BWA	Botswana	LSO	Lesotho	SYC	Seychelles
BFA	Burkina Faso	LBR	Liberia	SLE	Sierra Leone
BDI	Burundi	LBY	Libya	SOM	Somalia
CMR	Cameroon	MDG	Madagascar	ZAF	South Africa
CPV	Cape Verde	MWI	Malawi	SSD	South Sudan
CAF	Central African Republic	MLI	Mali	SDN	Sudan
TCD	Chad	MRT	Mauritania	TZA	Tanzania
COM	Comoros	MUS	Mauritius	TGO	Togo
COG	Republic of the Congo	MAR	Morocco	TUN	Tunisia
COD	DR Congo	MOZ	Mozambique	UGA	Uganda
CIV	Côte d'Ivoire	NAM	Namibia	ZMB	Zambia
DJI	Djibouti	NER	Niger	ZWE	Zimbabwe
EGY	Egypt	NGA	Nigeria	ESH	Western Sahara
GNQ	Equatorial Guinea				
ERI	Eritrea				
SWZ	Eswatini				
ETH	Ethiopia				
GAB	Gabon				
GMB	Gambia				
GHA	Ghana				

A.2 Boundary crosswalks

The GODAD treatment data and the FAO GAUL VIIRS extractions both reference GADM, but at slightly different vintages: GODAD v1.0 uses GADM 3.6 (the version available at the time of the original geocoding effort in 2021–2024), while Earth Engine's FAO GAUL 2015 simplified collection corresponds to a closely related but not identical set of administrative boundaries. The IPC Technical Working Groups define their analytical units ("Level 1" in the IPC nomenclature) using a country-specific mix of GADM-like administrative divisions and traditional analytical sub-regions that do not always coincide with the GADM 3.6 boundaries used by GODAD. We address these mismatches with two separate crosswalks.

IPC Level 1 ↔ GADM 3.6. The crosswalk between IPC Level 1 names and GADM 3.6 ADM1 units is constructed in three passes. First, we drop a small set of non-geographic IPC categories that some working groups publish (e.g., "Refugee Settlements" in Uganda, "Urban Analysis" and "Others" in Kenya, "Centro" and "Sul" in Mozambique, "Mainland" in Tanzania), since these do not correspond to administrative units that can be assigned a unique GODAD treatment value. Second, we apply manual translation overrides for cases of language difference or alternative spelling: for example, the Somali names of administrative units in IPC publications ("Lower Juba", "Middle Shabelle") are translated to their corresponding GADM 3.6 Somali-language names ("Jubbada Hoose", "Shabeellaha Dhexe"), and the Swahili names used by IPC for the Tanzanian Zanzibar regions ("Kaskazini Unguja", "Mjini Magharibi") are translated to their GADM English-language equivalents. We also apply manual mappings for a small number of countries where the IPC uses traditional sub-regions that aggregate multiple GADM districts (the cultural sub-regions of Uganda such as Acholi, Karamoja, Bunyoro, Busoga, Ankole; the new Ghanaian regions of Bono and Bono East mapped to the GADM 3.6 parent region Brong Ahafo, which was split in 2019). In these many-to-one cases, we aggregate GODAD treatment over the constituent GADM units to obtain the IPC zone's total exposure. Third, we apply an automatic fuzzy-matching procedure to the remaining IPC Level 1 names, using normalised string distance with a similarity threshold of 0.6.

After this procedure, 388 of the 407 IPC Level 1 units in non-excluded countries are matched to GADM 3.6 ADM1 units, a coverage rate of 95.3%. Twenty-one of the 27 African countries with IPC analyses achieve a 100% match rate; the remaining six countries (Mozambique, Tanzania, South Sudan, Togo, Central African Republic, Uganda) have partial coverage of 79–94%, principally because of IPC's use of country-specific sub-units that do not correspond cleanly to any GADM 3.6 entity. Madagascar and Kenya are excluded entirely from the IPC analytical sample because the structural mismatch between IPC's units (the 22 new Malagasy regions created in 2009 and the 47 Kenyan counties

created in 2013) and the GADM 3.6 boundaries (the 6 historical Malagasy provinces and the 8 historical Kenyan provinces) does not permit a meaningful crosswalk at the ADM1 level. Eritrea is also excluded as it has no IPC data.

FAO GAUL ↔ GADM 3.6. The crosswalk between the FAO GAUL 2015 administrative boundaries used by our Earth Engine VIIRS extractions and the GADM 3.6 boundaries used by GODAD is constructed for ADM1 only. For each GAUL 2015 ADM1 polygon in an African country, we attempt to match its administrative name to a GADM 3.6 name within the same ISO 3 country code, using a four-step procedure: (i) exact match on normalised strings, (ii) substring match in either direction, (iii) fuzzy match via standard string-similarity scoring with a cutoff of 0.6, and (iv) manual override for the small set of cases where automatic matching fails. The procedure yields 713 successful ADM1 matches out of 859 GAUL polygons, a coverage rate of 83.0%. The unmatched units are concentrated in small island states (Cape Verde, the Comoros, the Seychelles, São Tomé and Príncipe, Mauritius) where small populations and complex administrative reorganisations create persistent naming mismatches, and in Madagascar and Kenya for the same structural reasons that motivate their exclusion from the IPC sample. Because these countries either receive negligible USAID disbursement (the island states) or are already excluded for IPC analysis (Madagascar, Kenya), the loss of GAUL→GADM matches does not affect our headline results in the buffer-level analysis, which uses the GODAD project-level coordinates directly and does not depend on the GAUL→GADM crosswalk.

A.3 Sectoral classification robustness: OECD-DAC standard vs. GODAD pre-aggregated variables

Throughout the paper we apply the OECD-DAC standard classification to the three-digit CRS purpose codes recorded at the project-record level in GODAD (social = 1xx, economic = 2xx, productive = 3xx, humanitarian = 5xx + 7xx). GODAD also distributes four pre-aggregated sectoral variables (*disb_soc_split*, *disb_eco_split*, *disb_prod_split*, *disb_other_split*) constructed under a project-specific aggregation of CRS purpose codes. The two aggregations apply different mappings to the underlying purpose codes and therefore produce numerically different sectoral exposures.

Table A2 reports the joint sectoral regression for the nightlight buffer outcome (1 km) and the two IPC food-security outcomes (Phase 3+ and Phase 4+) under both classifications. The structure of the regression and the fixed-effect specifications are identical to those used in the main-text Tables 4 (buffer, country-trends primary specification) and 10 (IPC, country × analysis-date fixed effects). The qualitative conclusions of the paper, that humanitarian-aid contributes to both outcomes, and productive-sector aid at the buffer

scale only, are preserved across the two classifications: under the GODAD pre-aggregated variables, the dominant signal on both outcomes loads on *disb_prod_split*, which contains the bulk of CRS codes 5xx and 7xx classified as humanitarian under the OECD-DAC standard, with a secondary signal on *disb_other_split* for the buffer outcome.

Table A2. Sectoral coefficients under alternative classifications: OECD-DAC standard vs. GODAD pre-aggregated variables

Dependent variable:	Buffer 1 km log(avg_rad + 0.01)	IPC Phase 3+ % pop. in Phase 3+	IPC Phase 4+ % pop. in Phase 4+
Panel A. OECD-DAC standard classification (first digit of three-digit CRS purpose code)			
Social (1xx)	+0.0006	−0.0003	+0.0006
	(0.0017)	(0.0007)	(0.0004)
Economic (2xx)	+0.0009	−0.0010	−0.0001
	(0.0030)	(0.0008)	(0.0002)
Productive (3xx)	−0.0043**	+0.0003	−0.0001
	(0.0021)	(0.0008)	(0.0003)
Humanitarian (5xx, 7xx)	−0.0051**	+0.0018**	+0.0011***
	(0.0023)	(0.0007)	(0.0004)
Panel B. GODAD pre-aggregated sectoral variables			
`disb_soc_split`	+0.0002	−0.0004	+0.0005
	(0.0014)	(0.0007)	(0.0003)
`disb_eco_split`	+0.0005	−0.0011	−0.0003
	(0.0025)	(0.0008)	(0.0003)
`disb_prod_split`	−0.0040**	+0.0019***	+0.0011***
	(0.0018)	(0.0007)	(0.0004)
`disb_other_split`	−0.0044**	+0.0012	+0.0004
	(0.0017)	(0.0008)	(0.0004)
Observations	153,888	2,795	2,795
Clusters	1,374 projects	388 IPC units	388 IPC units

Notes. Each column reports the four coefficients on $\log(\text{USAID}^{(s)})_i \times \text{Post}_t$ from a single joint regression with all four sectors entered simultaneously. Panel A constructs the sectoral exposures from the three-digit CRS purpose code attached to each project record in GODAD, applying the first-digit OECD-DAC standard mapping: social = 1xx, economic = 2xx, productive = 3xx, humanitarian = 5xx and 7xx. Panel B uses the four pre-aggregated sectoral variables distributed with GODAD. The buffer-outcome specification (column 1) includes project, country, and month fixed effects plus country-specific linear trends (the paper's primary buffer specification). The IPC-outcome specifications (columns 2 and 3) include IPC-unit and country × analysis-date fixed effects (the paper's IPC sectoral specification). Standard errors clustered at the project or IPC-unit level as appropriate. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.4 Additional summary statistics

Table A3. Summary statistics for the project-buffer panel at the four buffer radii not shown in Table 1

	N	Mean	SD	p25	p75
Panel: Project-buffer VIIRS panel (500 m radius)					
avg_rad (nW sr^{-1} cm^{-2})	153,888	7.17	18.02	0.27	4.30
log(avg_rad + 0.01)	153,888	0.00	1.99	−1.29	1.46
lit_count (pixels above 0.3 nW threshold)	153,888	1.64	1.01	0.50	2.31
Panel: Project-buffer VIIRS panel (5 km radius)					
avg_rad (nW sr^{-1} cm^{-2})	153,888	3.76	10.10	0.26	1.51
log(avg_rad + 0.01)	153,888	−0.34	1.64	−1.32	0.42
lit_count (pixels above 0.3 nW threshold)	153,888	195.72	125.38	66.58	305.96
Panel: Project-buffer VIIRS panel (10 km radius)					
avg_rad (nW sr^{-1} cm^{-2})	153,888	2.45	7.03	0.25	0.94
log(avg_rad + 0.01)	153,888	−0.52	1.44	−1.34	−0.05
lit_count (pixels above 0.3 nW threshold)	153,888	744.66	504.53	220.46	1236.89
Panel: Project-buffer VIIRS panel (25 km radius)					
avg_rad (nW sr^{-1} cm^{-2})	153,888	1.21	3.43	0.25	0.62
log(avg_rad + 0.01)	153,888	−0.76	1.13	−1.36	−0.46
lit_count (pixels above 0.3 nW threshold)	153,888	4281.17	3076.12	1140.79	7296.03

Notes. The project-buffer panel covers 1,374 USAID project locations × 112 monthly VIIRS composites, with one row per (location, radius, month). USAID disbursement statistics (constant across radii) are reported in Table 1.

A.5 IPC measurement-quality diagnostic

A natural concern about the IPC outcome is that the Integrated Food Security Phase Classification draws on multiple sources of evidence, e.g., WFP mobile vulnerability surveys, FEWS NET seasonal monitoring, partner needs assessments, some of which are partly funded by USAID. If the data collection underpinning IPC classifications has been degraded by the post-shock loss of USAID funding, the IPC measure could under-detect acute food insecurity in the post-shock period, biasing our estimates toward zero. We assess this concern with three diagnostics.

Table A4. IPC operational continuity, pre vs post-shock, by country USAID exposure

USAID exposure	N countries	N with post-shock analyses	Cadence (analyses/yr)		Coverage ratio (post/pre)
			Pre	Post	
Q1 (lowest)	7	4	1.39	0.86	1.02
Q2	7	5	1.39	1.14	1.20
Q3	6	5	1.46	1.33	1.13
Q4 (highest)	7	5	0.79	1.00	1.13
All 27 countries	27	19	1.26	1.07	1.10

Notes. Sample: 27 African countries with active IPC Technical Working Groups during the analysis window. Quartiles defined within the IPC sample by total USAID disbursement to the country during 2017–2020 (constant 2014 U.S. dollars). Pre-shock cadence is computed over the 2022–January 2025 window (3.08 years); post-shock cadence over February 2025 – January 2026 (1.00 year). The coverage ratio is the per-country average ratio of the number of IPC units classified per post-shock analysis to the corresponding pre-shock average. The country-level Pearson correlation between USAID exposure (in millions of U.S. dollars) and the post/pre cadence ratio is +0.19 ($p > 0.1$, $N = 22$ countries with non-zero pre-shock cadence); the Spearman counterpart is +0.27. *Sources:* authors' calculations from the IPC Acute Food Insecurity Country Data and from GODAD v1.0.

Three patterns are visible in the table. First, IPC operational continuity is preserved across all four USAID-exposure quartiles. The mean post/pre cadence ratio is 0.86 (median 0.95) across the 19 countries with both pre- and post-shock analyses, approximately one analysis per year in both periods. Second, and most informative for the concern at issue, the cadence reduction (where it appears) is concentrated in the lowest, not the highest, USAID-exposure quartile. Q4 countries (the most USAID-exposed) have actually published *more* post-shock IPC analyses on a per-year basis than they did in 2022–2024; the country-level correlation between USAID exposure and the cadence ratio is positive (+0.19 Pearson, +0.27 Spearman), the opposite of the sign that would obtain if USAID-funded data-collection inputs were systematically failing in the most USAID-exposed countries. Third, the coverage per analysis (the number of IPC units classified) has been preserved or increased in all 19 countries with both pre- and post-shock analyses, with no country falling below 80% of its pre-shock coverage and several countries (Mozambique, Nigeria, the Central African Republic, Uganda) actually expanding their post-shock coverage.

These diagnostics cannot rule out degradation in the unobserved *quality* of the underlying survey inputs to IPC classifications. They do, however, reject the most direct empirical manifestations of such degradation: reduced cadence, reduced coverage, and concentration of any reduction in the most USAID-exposed countries. We note in addition that the principal data infrastructures feeding the IPC (WFP mobile vulnerability surveys, FEWS NET seasonal monitoring) have pre-existing financial diversification across multiple donors and are not USAID-dependent in their primary funding. To the extent that residual

measurement degradation cannot be ruled out, the direction of the resulting bias is toward zero (under-detection of acute cases), making the headline IPC estimates conservative lower bounds.

A.6 Leave-one-country-out: event study at the 1 km buffer

To verify that the post-shock pattern displayed in Figure 7 is not driven by any single country in the sample, we re-estimate the event study fifty times, each time dropping all USAID project locations from a single country, and record the re-centered coefficients on the first two complete post-shock semestral bins ($k = 0$: February–July 2025; $k = 1$: August 2025–January 2026). Figure A1 reports the leave-one-country-out (LOO) distribution of the two coefficients alongside the full-sample reference.

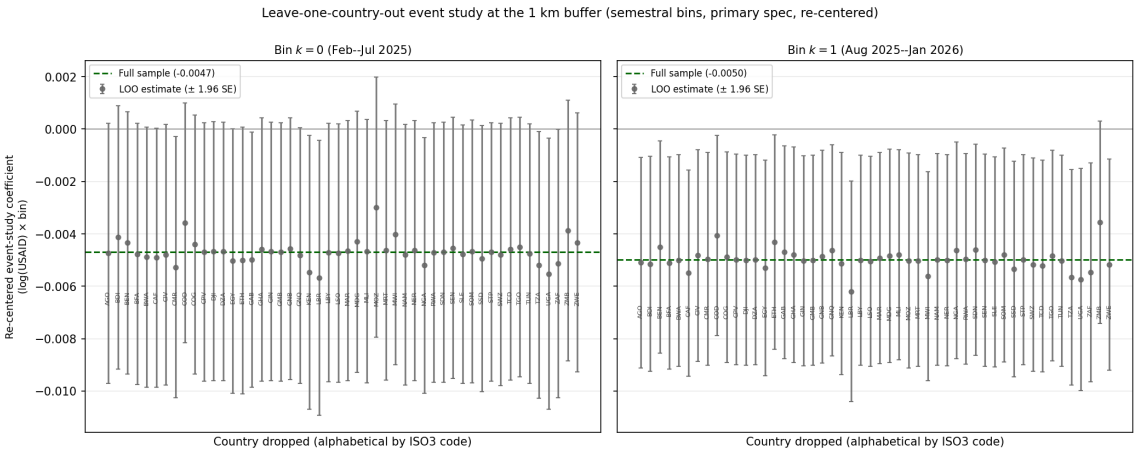

Figure A1. Leave-one-country-out event study at the 1 km buffer (semestral bins, primary spec, re-centered)

Notes. Each grey marker is the re-centered post-shock coefficient on the indicated bin from a separate regression that drops all USAID project locations in one of the fifty African countries in the sample. Whiskers are ±1.96 standard errors. The dashed green line is the full-sample re-centered coefficient. Countries are sorted alphabetically by ISO3 code along the horizontal axis. Primary specification (project FE, country FE, month FE, country-specific linear trends). *Source:* authors' calculations.

The exercise yields the following result. For bin $k = 0$, all fifty LOO subsets produce a negative point estimate (range: -0.0057 to -0.0030; median -0.0047, identical to the full-sample point estimate). For bin $k = 1$, all fifty LOO subsets produce a negative point estimate (range: -0.0062 to -0.0036; median -0.0050). The smallest absolute magnitude in each bin is obtained when Mozambique (the largest single-country contributor to the headline DiD) and Zambia are dropped, respectively; the largest absolute magnitude when Liberia is dropped. No single country drives the post-shock pattern visible in Figure 7, and the post-shock level shift documented by the trend-break decomposition of Section 4.8 is robust to leave-one-country-out resampling at the event-study level.

A.7 Event study at the 5 km buffer

Figure A2 replicates the event study of Section 5.2 at the 5 km buffer, the smallest radius at which the headline difference-in-differences coefficient is statistically significant at the 1% level. The semestral binning, the primary specification (project FE + country FE + month FE + country-specific linear trends), and the visual treatment (re-centering on the pre-shock average, omission of the reference bin $k = -1$ and of the partial bin $k = +2$ from the plot) are the same as in the main-text event study at the 1 km buffer.

The pattern is informationally richer than at the 1 km buffer because the larger buffer averages over a greater number of pixels per unit, tightening the standard errors. The five plotted pre-shock coefficients ($k = -6, \ldots, -2$) scatter tightly around zero with confidence intervals all crossing zero, satisfying the pre-trend stability test more cleanly than at the 1 km buffer. The first complete post-shock semester ($k = 0$, February–July 2025) yields a re-centered coefficient of -0.0040 with a confidence interval that just touches zero; the second ($k = 1$, August 2025–January 2026) yields -0.0036 with a confidence interval that excludes zero at conventional levels, providing individual statistical significance for a post-shock bin that the 1 km event study lacks. The implied pre-to-post differential is broadly consistent with the headline 5 km coefficient of $\widehat{\beta} = -0.0038$ ($p = 0.008$) in Table 2. The pattern at the 5 km buffer thus visualises the same identifying variation captured by the main-text 1 km event study, with the post-shock level shift made more statistically visible by the tighter standard errors of the wider buffer.

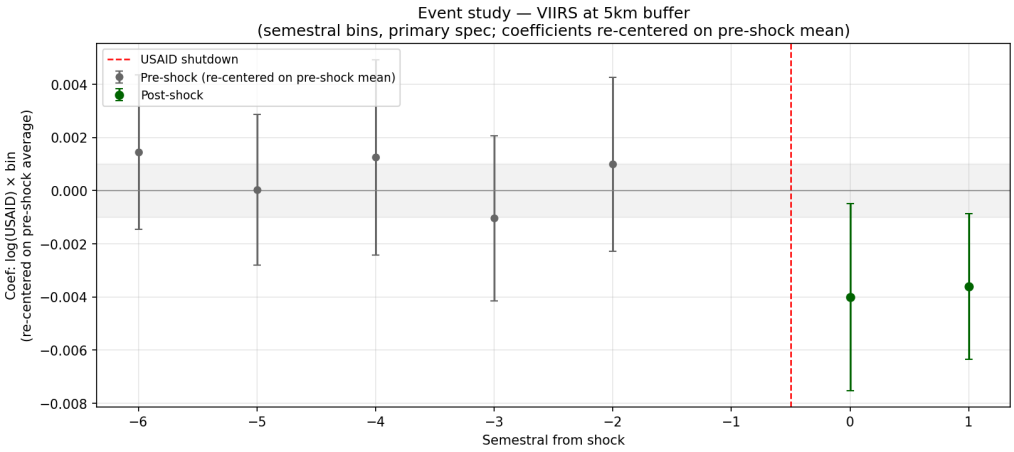

Figure A2. Event study at the 5 km buffer (semestral bins, re-centered on pre-shock mean)

Notes. Coefficients β_k from equation (3) at the 5 km buffer, with bins defined in calendar semesters relative to the USAID shock of January 2025. Each coefficient is re-centered by subtracting the average of the six pre-shock coefficients ($k = -6$ through $k = -1$, with $k = -1$ entered as zero by construction as the omitted reference of the regression); the resulting estimates represent deviations of each bin from the pre-shock average rather than from $k = -1$ alone. Whiskers are ± 1.96 standard errors. Bin $k = -1$, the omitted reference of the regression, is not plotted: it is identically zero in the unshifted regression by construction and is not an estimated quantity. Post-shock bin $k = +2$ (February–April 2026) is also excluded from the plot because it covers only three of the six calendar months of a complete semestral bin; its data continue to enter the regression through the unit and time fixed effects, and the re-centered coefficient is +0.0004 (SE 0.0027), statistically indistinguishable from zero. Primary specification (project FE, country FE, month FE, country-specific linear trends). Dashed vertical line: USAID shutdown date.

A.8 Trend-break decomposition under the primary specification

Table A5. Trend-break decomposition under the primary specification (country trends), 1 km buffer

	Primary spec (country trends)
β_1 (T × trend, pre-shock slope)	−0.00003
	(0.00003)
β_2 (T × trend × Post, slope break)	0.00063*
	(0.00035)
β_3 (T × Post, level shift)	−0.00743***
	(0.00283)
N	153,888

Notes. Trend-break specification on the 1 km buffer panel, equation (4). Dependent variable is $\log(\text{avg\_rad}_{i,t} + 0.01)$. Fixed effects: project (unit), country (iso3), month, plus country-specific linear trends. T denotes $\log(1 + \overline{D}_i^{2017\text{–}2020})$ and Post is the indicator for $t \geq$ February 2025. Coefficients β_1 and β_2 are slopes per month, hence of order 10^{-5}; we report five decimal places throughout this table to preserve readability of the pre-shock slope, departing from the four-decimal convention used in the main tables. Standard errors in parentheses clustered at the project-location level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.9 Leave-one-country-out tables

Table A6. Leave-one-country-out estimates for the buffer outcome (1 km, primary specification)

Dropped	$\widehat{\beta}$	SE	p	Dropped	$\widehat{\beta}$	SE	p
AGO	−0.0046**	(0.0020)	0.023	MAR	−0.0045**	(0.0020)	0.024
BDI	−0.0051**	(0.0021)	0.013	MDG	−0.0045**	(0.0020)	0.028
BEN	−0.0041**	(0.0020)	0.041	MLI	−0.0044**	(0.0020)	0.030
BFA	−0.0047**	(0.0020)	0.019	MOZ	−0.0038*	(0.0020)	0.052
BWA	−0.0047**	(0.0020)	0.019	MRT	−0.0046**	(0.0020)	0.021
CAF	−0.0049**	(0.0020)	0.014	MWI	−0.0047**	(0.0020)	0.022
CIV	−0.0045**	(0.0020)	0.026	NAM	−0.0047**	(0.0020)	0.020
CMR	−0.0047**	(0.0020)	0.022	NER	−0.0046**	(0.0020)	0.022
COD	−0.0034*	(0.0019)	0.070	NGA	−0.0045**	(0.0020)	0.026
COG	−0.0043**	(0.0020)	0.034	RWA	−0.0045**	(0.0020)	0.024
CPV	−0.0047**	(0.0020)	0.020	SDN	−0.0045**	(0.0020)	0.026
DJI	−0.0046**	(0.0020)	0.021	SEN	−0.0045**	(0.0020)	0.028
DZA	−0.0047**	(0.0020)	0.020	SLE	−0.0047**	(0.0020)	0.018
EGY	−0.0050**	(0.0021)	0.016	SOM	−0.0045**	(0.0020)	0.028
ETH	−0.0046**	(0.0020)	0.024	SSD	−0.0053***	(0.0020)	0.010
GAB	−0.0047**	(0.0020)	0.018	STP	—	(—)	—
GHA	−0.0046**	(0.0021)	0.024	SWZ	−0.0048**	(0.0020)	0.018
GIN	−0.0047**	(0.0020)	0.020	TCD	−0.0047**	(0.0020)	0.020
GMB	−0.0047**	(0.0020)	0.020	TGO	−0.0046**	(0.0020)	0.023
GNB	−0.0046**	(0.0020)	0.024	TUN	−0.0047**	(0.0020)	0.020
GNQ	−0.0046**	(0.0020)	0.023	TZA	−0.0054***	(0.0021)	0.010
KEN	−0.0049**	(0.0021)	0.021	UGA	−0.0057***	(0.0021)	0.007
LBR	−0.0058***	(0.0021)	0.007	ZAF	−0.0052**	(0.0021)	0.012
LBY	−0.0047**	(0.0020)	0.019	ZMB	−0.0040**	(0.0020)	0.049
LSO	−0.0047**	(0.0020)	0.018	ZWE	−0.0047**	(0.0020)	0.021
Full sample: $\widehat{\beta} = -0.0047^{**}$, SE = 0.0020, $p = 0.020$, $N = 153{,}888$							

Notes. Each row reports the estimate of $\widehat{\beta}$ on $\log(1 + \overline{D}_i^{2017\text{–}2020}) \times \text{Post}_t$ from the primary specification (project FE, country FE, month FE, country-specific linear trends) on the 1 km buffer panel, after dropping all observations for the indicated country. The dependent variable is $\log(\text{avg\_rad}_{i,t} + 0.01)$. Standard errors clustered at the project-location level. Coefficients are statistically significant at the 5% level in 47 of 50 subsets, and at the 10% level in 49 of 50 subsets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7. Leave-one-country-out estimates for the IPC outcome (Phase 4+, primary specification)

Dropped	$\widehat{\beta}$ (Phase 4+ × T × Post)	SE	p
BEN	0.0010***	(0.0004)	0.006
CAF	0.0010***	(0.0004)	0.008
CMR	0.0009***	(0.0004)	0.007
COD	0.0011***	(0.0004)	0.003
DJI	0.0010***	(0.0004)	0.006
GHA	0.0010***	(0.0004)	0.008
MLI	0.0010***	(0.0004)	0.007
MOZ	0.0009**	(0.0003)	0.011
MRT	0.0010***	(0.0004)	0.005
NER	0.0010***	(0.0004)	0.006
NGA	0.0009**	(0.0004)	0.015
SDN	0.0006***	(0.0002)	0.007
SLE	0.0009***	(0.0003)	0.007
SOM	0.0010***	(0.0004)	0.010
SSD	0.0009***	(0.0003)	0.008
TCD	0.0010**	(0.0004)	0.014
TGO	0.0010***	(0.0004)	0.006
UGA	0.0010***	(0.0003)	0.006
ZMB	0.0010***	(0.0003)	0.006
Full sample	0.0010***	(0.0003)	0.006

Notes. Each row reports the estimate on $\log(1 + \overline{D}_i^{2017\text{–}2020}) \times \text{Post}_t$ from the primary IPC specification (IPC unit FE, country FE, month FE, country-specific linear trends), after dropping all observations for the indicated country. The dependent variable is the share of population classified in IPC Phase 4+ during the current period of each analysis. Sample restricted to the 19 countries with both pre- and post-shock observations. Standard errors clustered at the IPC-unit level. Coefficients are statistically significant at the 5% level in 19 of 19 subsets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.10 Random-points placebo (buffer outcome)

A natural placebo for the buffer outcome is to verify that random African locations away from USAID activity do not exhibit the same post-shock pattern as USAID-exposed locations. We construct this test as a pooled difference-in-differences design: we generate 123 random points distributed across Africa, conditional on each being at least 50 km from any USAID project site, and we extract VIIRS radiance from January 2017 through April 2026 within 1 km and 5 km buffers around each random point using the same procedure (gas-flare mask, monthly composites, flare threshold) as for the USAID buffers. We then pool the random buffers (treated = 0) with the USAID buffers (treated = 1) and estimate

$$\log(\text{avg\_rad}_{i,t} + 0.01) \;=\; \beta\,(\text{treated}_i \times \text{Post}_t) \;+\; \alpha_i \;+\; \gamma_c \;+\; \gamma_t \;+\; \delta_c \cdot t \;+\; \varepsilon_{i,t},$$

restricting the sample to the 29 African countries in which both USAID project sites and random points are present. The coefficient β on $\text{treated}_i \times \text{Post}_t$ identifies the post-shock differential between USAID-exposed buffers and otherwise-comparable African buffers, conditional on country-specific linear trends and on month fixed effects that absorb any continent-wide common dynamic.

Table A8 reports the result at both buffer radii. At the 1 km buffer, the USAID-vs-random differential is negative and statistically significant at the 5% level. At the 5 km buffer the differential is negative but not statistically distinguishable from zero, consistent with the spatial-decay pattern documented in the main text. The 1 km result confirms the USAID-specific identification of the post-shock buffer-level decline: even relative to a control group of otherwise-comparable African locations, USAID buffers experience an additional decline that cannot be attributed to a generic continent-wide post-shock pattern.

Table A8. Random-points placebo: pooled DiD comparing USAID and random buffers

	(1) 1 km buffer	(2) 5 km buffer
$\text{treated}_i \times \text{Post}_t$	-0.0503^{**}	-0.0256
	(0.0223)	(0.0176)
N (USAID buffers)	103,376	103,376
N (random buffers)	13,440	13,440
Countries	29	29

Notes. Dependent variable is $\log(\text{avg\_rad}_{i,t} + 0.01)$. The sample pools 1,374 USAID project locations ($\text{treated}_i = 1$) with the 123 random African points ($\text{treated}_i = 0$) located at least 50 km from any USAID site. Sample restricted to the 29 African countries in which both groups are present. All specifications include unit fixed effects, country fixed effects, month fixed effects, and country-specific linear trends. Standard errors clustered at the unit (project or random point) level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We also examined a donor-substitution placebo at the ADM1 level for non-US donors, using the ADM1-level disbursement data for France, the United Kingdom, Germany, the World Bank, and Chinese commitments that GODAD makes available for all 23 donors in its panel. France ($\widehat{\beta} = -0.0006$, $p = 0.45$) and Chinese commitments ($\widehat{\beta} = -0.0003$, $p = 0.71$) yield null estimates; the United Kingdom is marginally significant at the 10% level ($\widehat{\beta} = -0.0017$, $p = 0.10$). Germany and the World Bank instead produce significant negative coefficients ($\widehat{\beta} = -0.0031$, $p < 0.001$, and $\widehat{\beta} = -0.0047$, $p < 0.001$, respectively). The Germany result is consistent with the contemporaneous 8% reduction in German bilateral ODA announced in the 2025 federal budget (Donor Tracker (SEEK Development), 2024); the World Bank result is consistent with the substantial U.S. presence in multilateral funding (approximately 17% of IBRD financial commitments and 19.6% of cumulative IDA contributions; Congressional Research Service, 2024), which makes World Bank

operations partly co-dependent on U.S. policy. The random-points placebo presented here avoids these donor-interdependence complications by comparing USAID buffers directly to non-treated African locations rather than to other-donor disbursements.

A.11 Time-shifted placebo (IPC outcome)

A complementary placebo for the IPC outcome replaces the actual shock date (28 January 2025) with a fictitious shock date two years earlier (1 February 2023) and re-estimates the baseline DiD on a sample restricted to the pre-shock period (January 2021 through January 2025), so that no portion of the actual post-shock period contaminates the test. The estimation uses the country-trends specification (IPC unit fixed effects, country fixed effects, analysis-date fixed effects, country-specific linear trends). We estimate the four placebo regressions corresponding to the four baseline specifications of Table 9: Phase 3+ and Phase 4+ outcomes, total-USAID and humanitarian-sector treatments.

Table A9 reports the placebo coefficients side by side with the actual-shock coefficients estimated on the full sample under the same country-trends specification. The four placebo coefficients are all statistically indistinguishable from zero at conventional levels and are several times smaller in magnitude than the corresponding actual-shock coefficients, three of which are statistically significant at the 5% level or better. The contrast between the two columns is consistent with the post-shock effect on acute food insecurity being specific to the timing of the actual shock rather than the artefact of a pre-existing differential trend in USAID-exposed IPC units.

Table A9. Time-shifted placebo: IPC DiD with fictitious shock at 1 February 2023

		Placebo shock (Feb 2023)		Actual shock (Jan 2025)	
Outcome	Treatment	$\widehat{\beta}$	p	$\widehat{\beta}$	p
Phase 3+	log(USAID total)	0.0009	0.103	0.0008	0.261
Phase 3+	log(USAID humanitarian)	0.0006	0.167	0.0017**	0.014
Phase 4+	log(USAID total)	0.0003	0.238	0.0010***	0.006
Phase 4+	log(USAID humanitarian)	0.0002	0.345	0.0010***	0.004

Notes. The placebo specification restricts the sample to January 2021 – January 2025 and replaces the actual shock date (28 January 2025) with a fictitious shock at 1 February 2023. The humanitarian-sector treatment is the log of USAID disbursement classified under the OECD-DAC standard humanitarian category (CRS codes 5xx food assistance and 7xx emergency response, including BHA and Food for Peace operations). Both columns use the country-trends specification: IPC unit fixed effects, country fixed effects, analysis-date fixed effects, and country-specific linear trends. The actual-shock column uses the full sample (January 2017 – January 2026). Standard errors clustered at the IPC unit level. Placebo $N = 1{,}769$; actual-shock $N = 2{,}800$. None of the four placebo coefficients is statistically distinguishable from zero at conventional levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.12 Magnitude calculations

This appendix reports the detailed IQR-based magnitude calculations that underlie the order-of-magnitude implications referenced in Section 7. The two paragraphs that follow translate the headline coefficients into substantively interpretable units (percentage changes in nighttime light radiance; percentage points of population in Phase 4+ food insecurity; absolute population counts).

Local economic activity. At the 1 km buffer, the country-trends specification yields a coefficient of $\widehat{\beta} = -0.0047$ on the interaction of log USAID disbursement with the post-shock dummy. The within-sample interquartile range of $\log(1+\overline{D}_i)$ at the project-location level is approximately 4.4 log units (from 10.1 at the 25th percentile to 14.4 at the 75th). The differential post-shock change in log(avg_rad+0.01) between a project at the 75th and a project at the 25th percentile of pre-shock exposure is therefore approximately -0.020 log units, equivalent to a 2.0% reduction in nighttime light radiance within the 1 km buffer. For a project at the 90th percentile of pre-shock exposure (cumulative disbursement of approximately $15 million) relative to the median (approximately $119,000), the implied differential effect is approximately 2.3% on local nightlight intensity within the immediate vicinity of the project site. Aggregated across the 1,374 USAID project locations in our sample, the implied loss of measured nighttime activity in the immediate vicinities of these sites is on the order of a few percentage points of pre-shock baseline radiance, with cumulative reductions concentrated at the most exposed project locations.

Food security. The humanitarian-sector treatment at the IPC analytical unit level produces a coefficient on Phase 4+ of $\widehat{\beta}^{\text{hum}} = +0.0010$. Multiplied across the within-sample interquartile range of log humanitarian-sector disbursement (approximately 11.5 log units in the IPC sample, reflecting that the 25th percentile corresponds to IPC units with no recorded USAID humanitarian disbursement while the 75th percentile corresponds to units with approximately $96,000 in annualised humanitarian disbursements 2017–2020), the implied differential effect in Phase 4+ population share between an IPC unit at the 75th and a unit at the 25th percentile of humanitarian-sector exposure is on the order of 1.2 percentage points. Translated into population terms in a typical African IPC unit (with average population on the order of one to two million people), this represents tens of thousands of additional individuals classified in Emergency or worse food insecurity for each high-exposure IPC unit, with the largest absolute populations affected concentrated in the most populous and most-exposed countries (the Democratic Republic of the Congo, Ethiopia, Nigeria, Sudan, and Mozambique). Summed across the high-exposure IPC units in our 19-country pre-and-post-shock sample, the implied additional burden of Phase 4+ food insecurity attributable to the USAID shock is on the order of one to several million

individuals, with a substantial part of this aggregate concentrated in the conflict-affected and chronically food-insecure subset of countries that contributes most of the leverage in the leave-one-country-out exercises.