The Productivity Effects of Irrigation Aid: Evidence from Border Difference-in-Differences in Sub-Saharan Africa

Nicole Robiro*

October, 2026

Click here for the latest version

This paper examines the microlevel effectiveness of irrigation-focused development aid in sub-Saharan Africa between 2005 and 2020 using georeferenced project data and high-resolution remote sensing. Employing a localized difference-in-differences strategy near country borders, the study finds evidence that irrigation aid raises the production value of cash crops, while staple crop yields remain largely unresponsive. This pattern indicates that the gains are concentrated in market-oriented production rather than extending broadly across crop categories. The study contributes to the literature on foreign aid and agricultural development by introducing a high-resolution empirical framework that leverages border matching and grid-level data to identify local treatment effects. This approach provides an alternative identification strategy to commonly used instrumental-variable and shift-share designs, expanding the scope for evaluating geographically targeted development interventions.

Keywords: Irrigation; Agricultural productivity; Cash crops; Aid effectiveness; Remote sensing; Sub-Saharan Africa

*Goethe University Frankfurt am Main, *robiro@econ.uni-frankfurt.de*.

1. Introduction

Agriculture remains a cornerstone of economic development in sub-Saharan Africa. The majority of Africa's extreme poor reside in rural areas and rely heavily on smallholder farming, making agricultural productivity central to poverty reduction and economic growth (World Bank, 2022, 2024). Yet despite its central role, agricultural productivity in the region lags far behind global standards, and less than six percent of cultivated land is irrigated (You et al., 2011), exposing smallholders to drought and intra-seasonal rainfall risk (Rockström et al., 2010; FAO, 2022). Irrigation is widely regarded as the needed adaptation technology because it decouples agricultural production from rainfall variability, enables multiple cropping, and increases attainable yields (Duflo & Pande, 2007; Jones et al., 2022; Dyer & Shapiro, 2023). However, with irrigation remaining scarce, dependence on rainfall is becoming increasingly costly due to climate change. Raising agricultural productivity is therefore not merely a sectoral concern but a central strategy for poverty reduction, food security, and economic development in the region (Diao et al., 2007; Christiaensen et al., 2011).

Whether irrigation reaches African farms is also a question about foreign aid. Public investment capacity for irrigation in most sub-Saharan African states is limited, and large-scale schemes have historically been developed and financed by governments with substantial support from bilateral and multilateral donors (Inocencio et al., 2007; Scheumann et al., 2017; Higginbottom et al., 2021). Donors, meanwhile, face an allocation problem under conditions of unusual scarcity. Aid budgets across major donor countries have decreased sharply in recent years. Therefore, donors have the need to know exactly how effective a constrained dollar is, based on previous outcomes. Surprisingly, little credible evidence exists on the question how effective irrigation aid is. This mainly results from two constraints. First, the aid-effectiveness literature has historically operated at the country level, where cross-national growth regressions yield fragile and contested results (Burnside & Dollar, 2000; Easterly et al., 2004; Rajan & Subramanian, 2008; Galiani et al., 2017). The more recent geocoded turn in this literature has moved the analysis to subnational units (Dreher & Lohmann, 2015; Isaksson & Kotsadam, 2018; Bluhm et al., 2025), but sector-specific causal evidence remains scarce. Second, the evidence that does exist on irrigation in developing countries comes from single-country studies such as dam construction in India (Duflo & Pande, 2007), groundwater in India (Sekhri, 2014), smallholder schemes in Mali (Dillon, 2011), or hill-

side irrigation in Rwanda (Jones et al., 2022). These studies are internally credible but capture idiosyncratic institutional settings and cannot tell a donor allocating a portfolio across multiple countries what the average return to irrigation aid is. The fundamental obstacle to answering that question directly is placement endogeneity: aid is deliberately targeted, flowing toward areas selected on agricultural potential, political salience, and donor priorities, so naive comparisons confound the effect of aid with the reasons for its placement.

This paper provides the first continent-wide analysis of the effect of irrigation-focused development aid on agricultural outcomes in sub-Saharan Africa. I combine geocoded aid projects from the Global Official Development Assistance Database (GODAD) with gridded agricultural outcomes from the Spatial Production Allocation Model (MapSPAM) across four waves (2005–2020). Identification comes from a border-discontinuity difference-in-differences design. Hereby, I compare cells within 25 kilometers of an irrigation aid project on the treated country's side of a national border against matched cells immediately across the border, before and after treatment, using coarsened exact matching (CEM) (Iacus et al., 2011) to balance geography and baseline development within border segments.

The border design addresses placement endogeneity through two complementary channels. First, African national borders, which were largely drawn in colonial-era negotiations with little regard for local conditions (Wesseling, 1996; Herbst, 2000; Englebert, 2009; Michalopoulos & Papaioannou, 2013, 2016), show homogeneous agro-ecological zones. Thus, cells on opposite sides share soil, climate, and topography but differ in exposure to national aid allocation. Hence, the border generates local comparisons in the spirit of geographic discontinuity designs. Second, and more importantly, the panel structure does not require that the two sides have to be similar in levels. Grid-cell fixed effects absorb every time-invariant difference across the border, national or local, shifting the entire identifying burden onto the assumption of parallel counterfactual trends. Most threats one might raise about African borders, such as institutional quality, market access, or ethnic composition, are level differences and are therefore absorbed. The inclusion of border-segment-by-year fixed effects ensures that treatment effects are identified from within-segment variation after accounting for all shocks shared by neighboring regions on both sides of the border in a given year. Consequently, the estimates are robust to local agro-climatic conditions, regional conflict dynamics, and market fluctuations that affect both sides of a border segment. What must hold is

that the timing of irrigation aid is not systematically related to diverging local trajectories. The institutional structure of aid delivery supports this: multi-year lags between project announcement, procurement, and commissioning, driven by donor budget cycles rather than local conditions, make it implausible that disbursement timing responds to short-run local agricultural dynamics. Additionally, I apply CEM to make grid cells more comparable, test baseline characteristics, and verify the parallel trends assumption directly with an event-study analysis.

I find that irrigation aid points to an increase in cash crop production value, while having no detectable effect on staple crop yields. Land-use analysis shows a decrease in cropping intensity for staple crops by 10%. However, no changes are present for the expansion of cultivated area, no changes in crop diversity, and no increase in forest loss, suggesting that irrigation primarily raises productivity on existing cultivated land rather than inducing land expansion or deforestation. The gains are concentrated in rural areas and depend on the local aid environment. These results suggest that irrigation investments primarily benefit market-oriented agricultural production rather than generating broad-based increases in staple crop productivity The results are robust to placebo tests, alternative bandwidths and treatment definitions.

The conceptual framework draws on theories of agricultural commercialization and induced innovation (Hayami & Ruttan, 1971). Following Schultz (1964) classic "poor-but-efficient" hypothesis, smallholders operate efficiently given their technology and resource constraints, implying that external investments which relax binding production constraints generate substantial productivity gains. In rainfed regions across sub-Saharan Africa, smallholder agriculture is constrained by rainfall dependence usually accounting only for one to two season within a year. Irrigation directly relaxes this constraint, elevating the marginal productivity of land and inducing demand for complementary inputs such as high-yielding seed varieties, inorganic fertilizers, and labor (Evenson & Gollin, 2003; Duflo et al., 2011). Technical progress or a decline in the cost schedule of a given commodity can raise returns to resources economy-wide if it diverts production from food cereals towards commodities facing more elastic demand in domestic and export markets (Hayami & Ruttan, 1971). Irrigation can enable such a shift: by reducing exposure to weather shocks, it makes high-value crops viable that require a more consistent water supply. These adjustments do not need to translate into an expansion of cultivated land. Instead, farmers may intensify production on existing plots, reallocate production toward more profitable cash crops, or adopt more productive cultivation practices while

leaving overall land use largely unchanged. In particular, I study whether irrigation changes land use, crop composition, cropping intensity, and whether its effectiveness depends on the surrounding aid environment. More broadly, the framework suggests that irrigation may induce a reallocation of land, labor, and other inputs toward higher-value cash crops rather than simply increasing production within the existing crop mix.

The paper contributes to the literature in three ways. First, it contributes to the literature on aid effectiveness with a particular focus on irrigation. Early work evaluated foreign aid primarily through cross-country growth regressions, yielding mixed conclusions regarding its impact on economic development (Burnside & Dollar, 2000; Hansen & Tarp, 2001; Easterly et al., 2004; Rajan & Subramanian, 2008). More recent research has shifted toward subnational analyses using geocoded aid information, allowing researchers to identify localized development impacts of aid projects (Dreher & Lohmann, 2015; Isaksson & Kotsadam, 2018; Bluhm et al., 2025). Within agriculture, however, the evidence remains comparatively limited. Existing studies examine the effectiveness of agricultural aid and technology-oriented interventions, including improved irrigation management, climate resilience, and technology adoption, but typically focus on individual projects or specific countries. Kaya et al. (2012, 2013) document heterogeneous effects of aid, with agricultural aid having positive impacts on growth and poverty reduction, while Ssozi et al. (2019) show substantial heterogeneity across aid categories, including negative substitution effects for some food-crop interventions. Recent evidence further demonstrates that donor-financed irrigation can improve agricultural productivity and resilience in specific settings such as Mali yet these evaluations remain project-specific and do not provide evidence on the broader effectiveness of irrigation aid across countries (BenYishay et al., 2023). By providing the first continent-wide evaluation of irrigation-focused development aid in sub-Saharan Africa, this paper bridges the aid-effectiveness literature with the literature on irrigation and agricultural development.

Second, the paper contributes to the literature on irrigation and agricultural transformation. Existing evidence demonstrates that irrigation can increase agricultural productivity, but almost all credible studies focus on individual countries or specific irrigation schemes, including large dams in India (Duflo & Pande, 2007), groundwater extraction (Sekhri, 2014), smallholder irrigation in Mali (Dillon, 2011), and hillside irrigation in Rwanda (Jones et al., 2022). While these studies offer strong internal validity, their findings are necessarily tied to particular institutional settings. At the same time, a growing literature documents that irrigation investments in sub-Saharan Africa frequently does

not meet expectations. Large-scale schemes deliver only a fraction of their planned irrigated area (Higginbottom et al., 2021), returns often fail to cover costs (Schuenemann et al., 2018), and performance has been linked to policy, institutional, and historical factors rather than to conditions inherent to the region (Bjornlund et al., 2020). This paper complements these two strands by estimating the average effects of donor-financed irrigation across a broad set of African countries, thereby providing evidence on whether irrigation investments translate into improved agricultural outcomes. Moreover, whereas previous work primarily evaluates productivity or yield effects, I examine the broader process of agricultural transformation by studying changes in crop composition, land use, cropping intensity, and environmental outcomes. In doing so, the paper speaks directly to recent work emphasizing the importance of understanding the adoption and consequences of agricultural technologies beyond yields alone.

Finally, the paper contributes methodologically to the empirical evaluation of geographically targeted development interventions. Because aid placement is inherently endogenous, recent studies have largely relied on instrumental-variable approaches to identify causal effects (Galiani et al., 2017; Dreher et al., 2021; Bluhm et al., 2025). Instead, I introduce a difference-in-differences design at the border that exploits comparisons between agroecologically similar regions located on opposite sides of national borders. Combined with high-resolution geospatial data, this approach offers a complementary empirical strategy for evaluating spatially targeted aid interventions without relying on external instruments. While the design is not intended to replace existing identification strategies, it demonstrates how natural geographic discontinuities can help overcome one of the central empirical challenges in the aid-effectiveness literature.

The paper proceeds as follows. Section 2 provides the context and institutional background. Section 3 elaborates on the data, the sample construction, and summary statistics. Section 4 outlines the empirical strategy, while Section 5 presents the results. Sections 6 and 7 present robustness checks and heterogeneity analyses. The final section concludes.

2. Context and Institutional Background

2.1. The Aid Delivery Pipeline

While a large agronomic literature documents the productivity effects of irrigation infrastructure, much less is known about whether irrigation projects financed through development aid generate measurable improvements in agricultural outcomes. Aid-funded irrigation projects differ from privately financed or government-funded irrigation schemes because project placement, implementation quality, maintenance arrangements, and complementary interventions are jointly determined within the aid process. Evaluating irrigation aid therefore provides evidence on the effectiveness of a major development policy instrument rather than on irrigation technology alone.

Unlike many policy interventions that become effective immediately after implementation, irrigation infrastructure is delivered through a lengthy administrative and construction process. The commitment date recorded in geocoded aid databases therefore marks the beginning of project implementation rather than the moment at which irrigation becomes available to farmers. Between project commitment and agricultural production lie several stages, including project preparation, procurement, construction, commissioning, and finally the gradual utilization of the completed irrigation scheme.

The first stage consists of project preparation and approval. Donor-financed infrastructure projects typically undergo extensive feasibility studies, technical appraisals, environmental assessments, and negotiations before construction can begin. Evidence from World Bank lending shows that the duration of project preparation is influenced primarily by donor-side political and administrative factors rather than recipient-country characteristics. For example, Kilby (2013) demonstrate that preparation periods are significantly shorter for politically aligned recipient countries, suggesting that the timing of project approval is shaped largely by donor decision-making rather than local economic conditions.

Project commitment is similarly only an imperfect indicator of implementation. Aid commitments frequently differ substantially from actual disbursements, and investment projects are typically financed over several years as construction progresses. Moreover, implementation delays may result from recipient countries' non-compliance with

aid conditions, administrative delays on the donor side, or changing economic and political circumstances, and geopolitical considerations affecting donor behavior.(Bulíř & Hamann, 2003; Bulíř & Hamann, 2008; Kilby, 2009). Consequently, the commitment date recorded in aid databases should not be interpreted as the date at which productive infrastructure becomes operational. These delays are particularly pronounced for irrigation infrastructure. Construction of dams, canals, pumping stations, and distribution networks requires substantial engineering works that often extend over many years. Studying large water infrastructure projects worldwide, Ansar et al. (2014) report an average construction period exceeding eight years, with the majority of projects experiencing schedule overruns.

The long implementation horizon has two important implications for empirical analysis. First, agricultural outcomes should not be expected to respond immediately following project commitment. Instead, productivity effects are likely to materialize only after infrastructure has been constructed and farmers begin adjusting their production decisions. This observation is consistent with the broader aid-effectiveness literature, which finds that even productive investment projects affect economic outcomes only after several years (Clemens et al., 2011). Second, because implementation is governed by donor approval procedures, procurement processes, and construction schedules, the precise timing of when irrigation becomes operational is influenced by administrative frictions that are largely unrelated to short-run changes in local agricultural performance. The empirical strategy developed in this paper builds directly on this implementation process by allowing sufficient time for irrigation infrastructure to become operational before evaluating its effects on agricultural production. This source in variation can provide exogeneity to some extent.

2.2. The Border Design in Sub-Saharan-Africa

The credibility of a border-based design depends on how borders came to be where they are, and sub-Saharan Africa is a unique setting. The continent's national boundaries were drawn in European capitals during the Scramble for Africa (1860–1905), at a time when the colonizing powers had barely settled the continent and possessed little knowledge of the geography or ethnic composition of the territories they were partitioning. The resulting borders cut across ecological and social space rather than following it. Alesina et al. (2011) show that African states are disproportionately "artificial," bounded

by straight-line borders, and Michalopoulos & Papaioannou (2016) demonstrate that, apart from the sheer land mass and water area of their homelands, partitioned and non-partitioned ethnic homelands are statistically similar across a wide range of geographic, economic, and historical characteristics.

To answer the question on whether irrigation aid is effective nearby border regions, the treatment and control units should be comparable in agro-ecological conditions, which is the case for Africa. Borders were not drawn with agro-ecological reasoning, thus, grid cells facing each other across a national boundary typically share climate, soils, elevation, growing-season length, and to some extent farming traditions, crop portfolios, and land-use practices, due to the separation of the same ethnic group at the border. What differs discontinuously at the border is the national policy environment, including exposure to donor portfolios and the geocoded aid projects. A well-established literature exploits precisely this configuration, comparing outcomes within ethnic homelands split by national borders to isolate country-level treatments (Michalopoulos & Papaioannou, 2013). At the same time, McCauley & Posner (2015) mention that African border designs might offer weaker inferential leverage as expected. In particular, not every border segment is arbitrary. Some follow rivers or watersheds with direct agronomic relevance, populations sort across borders, and treatments diffuse spatially. However, the underlying border design is less demanding than the cross-sectional border comparisons this critique targets. This empirical strategy does not require that cross-border cells are identical in levels. In my setting, cell fixed effects absorb all time-invariant differences and border-segment-by-year fixed effects absorb any shocks common to both sides of a segment. The border's role is to supply comparison cells that share the local agro-ecological and cultural environment and are measured identically in the gridded outcome data on both sides. Furthermore, the border design does not assume that irrigation projects are randomly allocated across space. Development agencies may target areas with agricultural potential, existing infrastructure, or political importance. This cannot be ruled out with the border design, instead, I rely on the quasi-random arrival of aid timing. While the announcement of an irrigation aid project is not random, the construction of the irrigation site takes several years and usually takes longer than expected. Hence the actual treatment start cannot be anticipated by donors or governments and can be considered as quasi-random.[1]

[1]Section 4 elaborates more on this.

3. Data and Summary Statistics

3.1. Data

Here I describe the data used to test for the effectiveness of irrigation networks provided by donors on harvest yields.[2] A panel covering 4 years between 2005 and 2020 containing grid cells nearby country borders across entire sub-Saharan Africa forms the core of the data. Crop production data come from MapSPAM (Spatial Production Allocation Model) provided by the International Food Policy Research Institute (IFPRI) and is augmented with data on development aid projects from GODAD.

Bomprezzi et al. (2024) introduces GODAD, a geocoded database of development aid projects compiling project-level data from bilateral and multilateral donors, including major Western donors, multilateral organizations, and emerging donors such as China and India. For each project it records location, sector classification, financial commitments, duration, and textual descriptions. The full database spans from 1973 to 2023, though roughly 90 percent of projects fall in the post-2000 period, with observations concentrated between 2000 and 2020. Within Africa, aid spreads across many sectors, with agriculture being one of the dominant ones taking up a share of 13%. Irrigation forms a small component within this portfolio, with 896 projects across Africa and roughly five percent of agricultural projects in recent years.[3] For the analysis I restrict the sample to projects located near country borders, which forms the basis for the border difference-in-differences design. To ensure accurate spatial assignment, I keep only projects with a precision code of two or lower, corresponding to a maximum location error of about 20 kilometers. Irrigation projects are identified from textual information in project titles and descriptions.[4] When a project is implemented in multiple locations, each location is treated as a separate observation, since these represent distinct intervention sites.

MapSPAM provides crop-specific yields, production, and both harvested and physical areas for 42 crops for all countries in sub-Saharan Africa (IFPRI, 2024). It allocates administrative-unit crop statistics into 10km $\times$ 10km grid cells via a cross-entropy optimization that combines production statistics, cropland maps, agro-ecological suitability,

[2]For greater detail on data sources, the construction of key variables and the final sample see appendix A.1.

[3]More detailed information about aid allocation in Africa can be found in Appendix A.2

[4]For a detailed description of detecting irrigation projects from GODAD consider the section 3.2

population density, and crop prices, separating four production systems (irrigated high-input, and rainfed high-input, low-input, and subsistence). Observations are available for 2005, 2010, 2017, and 2020[5][6], which form the panel structure. The main outcome variables are crop yields, measured in kilograms per hectare (kg/ha), and the physical area under cultivation. For each 10 km × 10 km grid cell, I compute the mean value of the crop-specific yield raster, yielding the average estimated yield for each crop within the grid cell. In addition, I aggregate the crop-specific physical area raster to measure the total area allocated to each crop within the grid cell. Finally, I aggregate individual crops into two broader categories: staple crops[7] and cash crops[8] to construct measures of average yield and total cultivated area for each crop group.[9] To capture the economic value of agricultural production based on the population density, I construct an index-weighted production value measure by multiplying crop production (yield × harvested area) by the FAOSTAT (2025) producer price index dividing it by the local population. Because the producer price index is an index rather than a price level, this variable should be interpreted as a proxy for the relative value of agricultural production rather than a measure of farm revenues.[10][11]

One caveat of the data could be that it contains systematic measurement bias given that it is difficult to measure crop yields in specific areas across entire continents. MapSPAM's pixel values are model-based allocations scaled to match national FAO aggregates, not independent field measurements, and their accuracy depends on input statistics and cropland maps. Thus, the derivation of yields allocates crops using suitability, satellite cropland extent, population, market access, irrigation extent, and prior crop presence. One potential problem arising from the underlying irrigation extent could be that areas that receive irrigation networks such as those provided by the aid donors will yield higher outcomes than the corresponding control units, simply due to the algorithm. For this sample, however, this is not an issue given that the underlying irrigation extent relies on

[5]Data for 2005 are from IFPRI & IIASA (2016), for 2010 from IFPRI (2019), for 2017 from IFPRI (2020) and for 2020 from (IFPRI, 2024).

[6]Data for 2000 are also available, however, it differs from the remaining years in terms of crop availability and measurement. This year is therefore excluded.

[7]Staple crops include cassava, maize, millet, plantain, rice, sorghum, and yams.

[8]Cash crops include coffee, cocoa, cotton, groundnuts, oil palm, and tea.

[9]A more detailed overview of the MapSPAM methodology for harvest yield construction are provided in Appendix A.3.

[10]Real producer prices could not be used, as values were missing for more than one third of the countries in the sample.

[11]The FAOSTAT producer price index uses the average producer prices from 2014–2016 as its base period, with commodity weights based on average production values over the same years.

2005 data, which is prior to the first development aid being active in my database. Hence, all measures for rainfed and irrigated areas are determined prior to my study period. Similarly, the location of irrigation aid projects are not an input. AidData geocoded aid points are not in the MapSPAM pipeline, thus systematic measurement bias is highly unlikely to arise. With this regard, even though there might be measurement bias in the harvest yield estimates, it is likely that these measurement biases are mainly classical rather than systematic.[12]

Next to the two main data sources, I construct a set of control variables to capture baseline economic conditions and agro-ecological characteristics. Population density is taken from WorldPop (WorldPop, 2018). Travel time to the nearest city (measured in minutes) is computed using global road network data and friction surfaces (Nelson, 2008; Weiss et al., 2018). Nighttime lights are derived from satellite imagery (Li et al., 2023) and aggregated to the grid-cell level[13]. To proxy water availability and vegetation dynamics, I use indicators from FAO's WaPOR database (FAO WaPOR, 2025), including actual evapotranspiration, transpiration, interception, precipitation, total biomass, gross and net biomass water productivity. These measures capture water losses from soil and plant surfaces, overall vegetative growth, and the efficiency with which water is converted into biomass. Finally, data for slope and elevation were accessed via Digital Earth Africa Contributors (2021).[14]

3.2. Sample Construction

Detection of Irrigation Aid

I identify irrigation projects in the GODAD through a multilingual keyword classification applied to three text fields: Project title, short description, and long description, combined with the OECD-CRS sector classification. Project documentation appears in multiple languages (English, German, French, Spanish, and Portuguese) the search dictionary contains language-specific roots (e.g., irrig-, Bewässerung, périmètre irrigué, regadío, goutte-à-goutte).[15] Short terms prone to false positives (e.g., Spanish riego,

[12] Appendix A.4 develops this argument in detail.

[13] Hereby, the "majority" value among pixel points within a grid cell is used as the reference value

[14] Both of these variables come from U.S. Geological Survey (2000).

[15] Two text-processing problems required correction before matching. First, a subset of records suffers from character-encoding corruption (UTF-8 text read as Latin-1), so that accented characters in French and Portuguese entries appear as multi-byte artifacts; I repair these sequences prior to

French aspersion) are matched only at word boundaries, and drainage terminology (Entwässerung), which contains the German root for water but denotes the opposite intervention, does not enter the dictionary. Text-based detection is complemented by the OECD purpose category "Agricultural water resources" (CRS code 31140): this rescues projects whose titles are uninformative but whose sector coding identifies them as irrigation.

Detected projects are then classified by the centrality of irrigation to the intervention. A project is coded as main focus (focus 1) if an irrigation term appears in the project title or if it carries the CRS irrigation purpose code. Hereby, the title needs to emphasize an irrigation related project or sector coding indicates that water infrastructure constitutes the core budget component. Projects in which irrigation appears only in the description are coded as component (focus 2): these are typically broader food-security, market-development, or integrated rural programs that fund some irrigation among many activities (e.g., "Feed the Future Resilient Agricultural Markets Activity"). A third category (focus 3), captures projects whose primary purpose is unrelated to irrigation and captured by accident due to the mentioning of drinking-water supply, sanitation, schools, refugee response, or hydropower. These accidental matches are identified by a second keyword screen and are never treated as irrigation interventions. Of 2,480 detected irrigation project-locations in sub-Saharan Africa, 1,310 are main-focus, 784 component, and 386 incidental. My analysis restricts treatment to main-focus projects. Component projects deliver an unknown and likely small irrigation dose within a bundled intervention, so including them would downward bias the treatment definition and conflate the effect of irrigation aid with that of heterogeneous rural programming. Focus-2 projects are, however, kept as an identifier for irrigation project regions. This is done for the purpose of screening potential control cells, since a cell adjacent to any irrigation-related activity cannot be used as a clean control unit.

Finally, treatment assignment requires reliable coordinates. Each project-location carries a geocoding precision code; I define treatment using locations at precision 1 (exact coordinates) and 2 (within approximately 25 km). For the main sample I consider projects up to a precision code 2 and use precision-1-only assignment as a robustness check, since location error of the precision-2 magnitude (up to 20 kilometers) is similar to the treatment buffers and could lead to misclassifications of project regions. Of the

processing. Second, all remaining accents and umlauts are transliterated to plain ASCII so that, for example, German compounds of Bewässerung match regardless of encoding.

1,310 main-focus locations, 1,162 are precision 1 and 110 precision 2. After removing duplicate records of the same project at identical coordinates the matched estimation sample links treated cells to 182 distinct projects at 250 sites across 24 countries.

Identifying Treatment and Control Regions at the Border

The identified irrigation aid project sites are used to define the treatment and potential control sample. A grid cell is classified as treated if an irrigation project site is located within 25 kilometers of the cell and within the same country. The baseline 25-kilometer buffer is intended to approximate the spatial reach of irrigation investments, whose direct benefits are expected to be concentrated in nearby communities and to diminish with distance as access to irrigation infrastructure becomes more limited. Choosing a substantially smaller buffer would risk misclassifying locations that are effectively exposed to the intervention as untreated, while a substantially larger buffer would include locations with little or no direct exposure, thereby attenuating the estimated treatment effect. To demonstrate that the results are not driven by this choice, the robustness analysis considers alternative treatment buffers ranging from 5 to 55 kilometers. Cross-country misclassification of treatment is avoided by classifying a grid cell as treated only if the corresponding irrigation aid project is located within the same country according to the GODAD database.

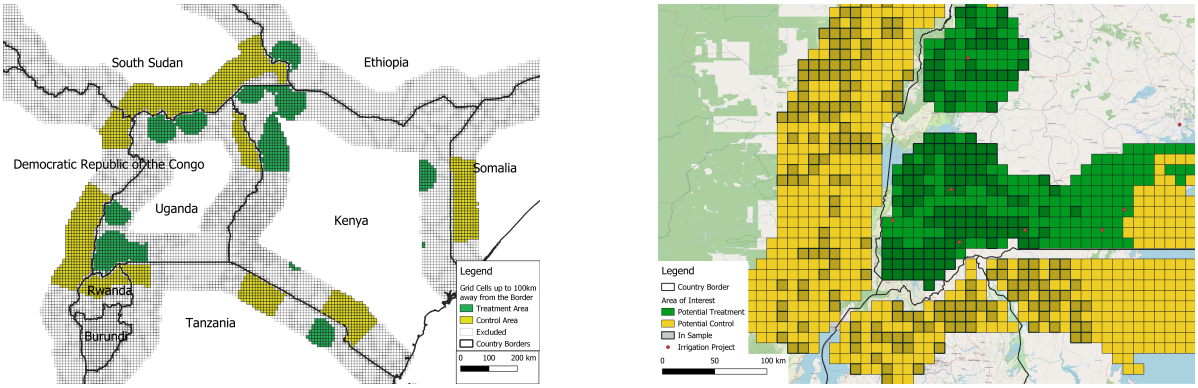
For control locations, only cells on the opposite side of the border are kept. To make sure that no irrigation aid within the country of the control grids is present, I consider a grid cell as control only if it is at least 100 kilometers away from the next irrigation site within the same country. Importantly, this screening is based on all detected irrigation projects regardless of their focus classification (focus 1 and 2) or geocoding precision. Hereby, I do not only use precision code 1 and 2 but extend the irrigation projects up to precision code 3 (project is located within the administrative unit), whereas the treatment definition relies exclusively on main-focus projects with precise coordinates (precision code 1 and 2). A project of a precision code 3 which is measured not precisely enough to define treatment can therefore still disqualify a nearby cell from serving as a control. Cross-border proximity to foreign irrigation projects does not disqualify a control cell.

All cross-country borders are divided into different border segments based on random dots allocated at the border. Grid cells are then allocated to the closes border segment available. A segment enters the estimation sample if exactly one side contains treated

cells and the opposite side provides clean control cells. If irrigation projects are present on both sides of a border segment, the segment is dropped because no suitable local control group is available. Similarly, a border segment does not enter the main sample if neither sides receives treatment. Since segment validity depends on the buffer width, the set of identifying segments varies across the distance specifications, ranging from 21 segments under the 5-kilometer definition to 71 under the 55-kilometer definition.

To improve comparability, the sample is restricted to grid cells within 100 kilometers of an international border. I choose this cutoff to reduce spatial heterogeneity and at the same time retain a sufficiently large sample for the analysis. To verify that results are not driven by this bandwidth choice, I vary the results in the robustness section. Additionally, cells that intersect the border itself are excluded, since they cannot be assigned unambiguously to one side. For a similar reason, cells close to a border (within 5 kilometers) are excluded as well to avoid misclassifications due to imprecise border definitions or cross-border farming.

Figure 1a illustrates the spatial distribution of treated and potential control grid cells for a subset of East Africa. Green cells indicate treated areas, yellow cells indicate potential controls, and unshaded cells are excluded based on the distance and border rules described above.



(a) Subsample: East Africa (b) After CEM Matching

Notes: The figure in panel a shows the spatial distribution of irrigation projects in East Africa. Green grid cells indicate treatment and yellow cells denote potential controls within the same border environment. Country boundaries are shown for reference; unshaded cells are excluded by the sample construction rules. Panel b illustrates the final grids that enter the sample (shaded grid cells) comprising the sample after CEM.

Figure 1: Illustration Subsample

Coarsened Exact Matching (CEM)

To establish the causal effect of irrigation development aid on crop yields a major threat is the non-random placement of donor-funded infrastructure. Governments and development agencies systematically target irrigation projects toward regions with specific, often unobserved, characteristics. On the one hand, donors may favor areas with high baseline agricultural potential and accessible water tables to maximize the return on investment. On the other hand, humanitarian objectives often direct resources toward marginalized, drought-prone environments with low baseline productivity. Simply comparing treated grid cells to untreated ones would bias the true treatment effect of irrigation with these confounding geographic and socio-economic factors. To overcome this selection bias, this study extends the the spatial border-discontinuity framework by applying Coarsened Exact Matching (CEM), restricting the comparison of treated units to highly comparable control units located on the immediate opposite side of the same border. CEM bounds the maximum level of imbalance by allowing to pre-determine meaningful, discrete bins for each covariate. Most importantly, by enforcing exact matching on the border segment identifiers, I guarantee that the spatial structure of the counterfactual is strictly preserved. This ensures that treated grid cells are evaluated exclusively against control cells that share the identical micro-climate, regional soil characteristics, and local topographical features of that specific stretch of the border.

To prevent distortions from extreme topographic anomalies, slope and elevation are winsorized at the 5th and 95th percentiles. Because irrigation projects are initiated in different years across sub-Saharan Africa, I construct segment-specific onset windows. For each border segment, the baseline year is assigned to the closest available pre-treatment survey wave (2005, 2010, and 2017) preceding the introduction of any irrigation infrastructure along that segment. The data is then restricted to this unique pre-treatment window to keep baseline covariates. The CEM is done for elevation, slope and log of population density. Agroclimatic variables from FAO WaPOR (2025) such as evapotranspiration, precipitation and total biomass start from 2009 onwards and do not overlap with the first year of the panel, therefore I cannot consider it for the matching without major losses in the sample. To verify, whether my treatment and controls units are similar in agroecological conditions, I compare them based on these factors after the matching. Table 1 reports the matching statistics. The CEM process successfully matches 99.5 percent of the treated cells (4,927 out of 4,949) and 97.3 percent of the control cells (11,697 out of 12,015). The univariate imbalance across confounders improves with the CEM with all of the considered variables lying below the threshold of 0.2. This reduces model dependence and ensures that estimated treatment effects are not driven

by systematic differences in local characteristics. Figure 1b shows successfully matched treatment and control grids as shaded areas.

Table 1: Coarsened Exact Matching (CEM) Summary and Covariate Balance

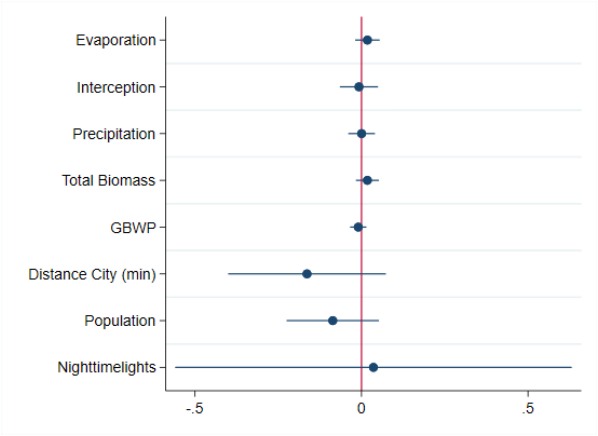
Panel A: Sample Summary				
	Control		Treated	
Total available grid cells	12,015		4,949	
Matched grid cells	11,697		4,927	
Unmatched grid cells (dropped)	318		22	
Panel B: Pre-matching Univariate Imbalance				
Covariate	L_1 Distance	Mean Diff.	Min Diff.	Max Diff.
Elevation	0.1510	-21.438	0	0
Slope	0.0797	29.178	0	.
log(Population)	0.1530	0.3278	1.341	.
Border Segment	0.1063	1.0252	0	0

Notes: This table presents the summary statistics for the CEM routine matching treated cells to control cells along border segments. Panel A documents the sample pruningPanel B reports the pre-match univariate imbalance metrics across the matching covariates. Continuous covariates (elevation, slope, and population) are coarsened into automated bins. All treatment and control units have to be matched at the corresponding border segment variable (*Border Segment*).

Before estimating the effects of irrigation on harvest yields, I provide descriptive evidence supporting the identifying assumptions. First, I examine whether the presence of irrigation projects is correlated with observable baseline characteristics. Specifically, I regress the treatment indicator on covariates, including nighttime light intensity, population density, mean precipitation, and total biomass production. These variables capture infrastructure, climatic conditions, and agricultural potential that may influence both project placement and productivity. The estimated coefficients are small and statistically insignificant across specifications, suggesting that, conditional on the border-based sample construction and CEM, treated and control cells are similar along key observable dimensions prior to treatment. Figure 2 summarizes these results.

3.3. Summary Statistics

Table 2 summarizes the composition of the final estimation sample. After applying all sample restrictions and the border segment-level CEM matching, the sample comprises 10,240 grid cells located in 71 border segments across 34 countries, observed over the four



Notes: This figure presents coefficient estimates from regressions of the treatment indicator on covariates to validate balance between treated and control grid cells. 95% confidence intervals are displayed.

Figure 2: Balance Check for Baseline Characteristics

panel waves, which yields a final sample size of 40,960 grid-year observations. Since the treatment definition varies with the distance buffer, the table reports the composition separately for each buffer width. Under the baseline definition of 25 kilometers, 667 grid cells are treated across 50 border segments, facing 5,555 potential control cells on the untreated side of the respective borders (an average of roughly 111 control cells per identifying border segment). The number of treated cells increases with the buffer width, from 26 cells in 21 segments at 5 kilometers to 2,697 cells in 71 segments at 55 kilometers, while the average control pool per border segment remains stable at around 110 cells.

The treatment traces back to 182 distinct irrigation projects implemented at 250 geocoded sites in 24 countries. These projects represent approximately USD 1.50 billion in disbursements, Where the implementation period is observed, the average project runs for roughly 4.2 years. Such a long delay in actual operation is considered in the analysis by including a five year lag.

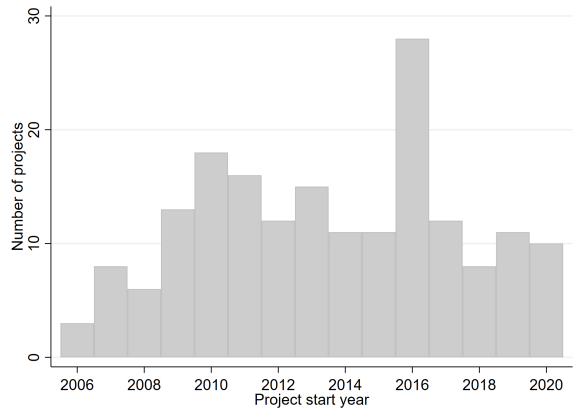
Figure 3 shows the distribution of project start years. Project onsets are spread continuously across the 2006–2020 window, rising through the late 2000s, peaking in 2016 with 28 new projects, and remaining at around ten projects per year thereafter.

Figure 4 documents the number of projects by donors. The projects in the sample are mainly financed by bilateral donors: Germany alone accounts for 48 percent of all projects, followed by the United States (13 percent), Spain (11 percent), and Belgium (8

Table 2: Sample composition by treatment buffer

	Treatment buffer					
	5km	15km	25km	35km	45km	55km
Treated cells	26	254	667	1,242	1,929	2,697
Border segments	21	39	50	56	65	71
Potential control cells	2,247	4,324	5,555	6,312	7,017	7,543
Avg. controls per segment	107	111	111	113	108	106
Main Sample: 25km Buffer						
Total grid cells			6,222			
Grid-year observations			24,888			
Panel years			2005; 2010; 2017; 2020			
Countries		Treated: 22			Total 33	

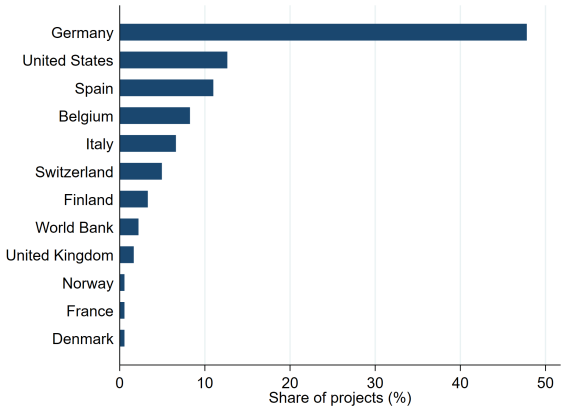
Notes: The table reports the composition of the CEM-matched estimation sample by treatment buffer. Treated cells are grid cells within the respective distance of a main-focus irrigation project site in the same country. Potential control cells are cells on the untreated side of border segments located at least 100km from any irrigation aid within their own country.



Notes: The figure shows the number of distinct irrigation projects in the matched sample by start year. Projects with onsets before 2005 are excluded from the sample.

Figure 3: Distribution of project start years

percent). The World Bank is the only multilateral donor and contributes four projects, which are, however, considerably larger and longer-running than the bilateral average, with a mean implementation period of nine years compared to roughly two to four years for most bilateral donors. The estimated treatment effects should therefore be interpreted primarily as the effect of bilateral, German-driven irrigation aid, and cannot be generalized to the broader universe of irrigation aid, in which multilateral donors such as the World Bank account for the largest share.



Notes: Share of distinct projects by donor. Multi-site projects are counted once.

Figure 4: Donor composition of matched irrigation projects

Treated cells are concentrated in Zambia, Benin, Niger, and Uganda, while several countries, including the Democratic Republic of the Congo, Chad, and Côte d'Ivoire, enter exclusively as control-side neighbors. This asymmetry is a feature of the border design: a country need not receive irrigation aid itself to contribute identifying variation, as long as it shares a border segment with a treated neighbor.[16]

4. Empirical Strategy

I estimate the effect of irrigation aid on agricultural outcomes using a difference-in-differences (DiD) framework at sub-Saharan country borders. The baseline specification is

$$Y_{itb} = \beta\big(\text{IrrigAid}_{ib} \times \text{Post}_{it}\big) + \delta_{tb} + \eta_i + \varepsilon_{itb}, \tag{1}$$

[16] Appendix A.5 reports the country composition of the baseline 25km sample.

where Y_{itb} denotes the outcome of interest (agricultural yields or harvest production value) for grid cell i in year t of a given border segment b. The main regressor of interest is $\text{IrrigAid}_{ib} \times \text{Post}_{it}$, which captures the average effect of exposure to irrigation aid after implementation.

The indicator IrrigAid_{ib} equals one for grid cells classified as treated based on proximity to an irrigation project (within 25km) and zero otherwise. Because aid is recorded in the GODAD database at the time of project approval or announcement, whereas implementation and disbursement typically occur over several years, I introduce a lag between project start and treatment in the outcome data. Specifically, I define $\text{Post}_{it} = 1\left[t \geq \text{onset}_i + 5\right]$, where onset_i is the start year of the earliest irrigation project within 25 kilometers of cell i. The five-year lag reflects the time needed for construction of irrigation aid: the average implementation period of the project sites in my sample is 4.2 years (Section 3.3). While a four year lag could be argued as well, the commissioning of the infrastructure and the first affected growing seasons add further time, which results in a five year lag. Section 6 shows that the results are not sensitive to this choice by varying the lag between one and eight years. Given that agricultural outcomes are observed at roughly five-year intervals (2005, 2010, 2017, 2020), the first post-treatment observation for a treated cell typically occurs in the second wave at least five years after project initiation.

Grid-cell fixed effects, η_i, absorb all time-invariant characteristics such as soil quality, long-run climate, and baseline infrastructure. Year fixed effects $\times$ border-segment fixed effects, δ_{tb}, capture common shocks to all locations, including regional weather anomalies and changes in global commodity prices. All regressions are weighted by the product of the CEM weights derived in Section 3.2 and a triangular kernel in the distance to the border, $w_i = \max\{0,\, 1 - |d_i|/110\}$, which places greater weight on cells close to the border where the identifying comparison is strongest; d_i denotes the distance of cell i to the border in kilometers and 110 the maximum distance in the grid of my sample. The error term ε_{itb} is clustered at the border-segment level.

The identifying assumption is that, conditional on the fixed effects and weights, treated and control cells within a border segment would have followed parallel outcome trajectories in the absence of irrigation aid. Importantly, the design does not require treated and control cells to be similar in levels. Any time-invariant difference across the border, e.g. in institutions, market access, land tenure, colonial legacy, or agro-ecological endowments, is absorbed by the grid-cell fixed effects, and any segment-wide shock in

a given wave is absorbed by the border-segment-by-year fixed effects. What remains is the requirement that the timing of irrigation aid must not be systematically related to diverging local trajectories of agricultural outcomes. The identification burden thus falls entirely on the exogeneity of treatment timing rather than of treatment location.

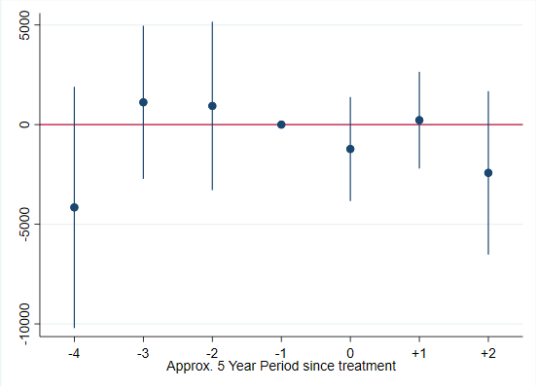
The exogeneity of treatment timing is supported by two considerations. First, the institutional structure of aid delivery usually takes longer then expected and the exact date of completion is unknown, especially for large infrastructure projects such as irrigation networks. Usually, it takes years between a project's announcement, procurement, construction, and the delivery of water, and this lag is governed by donor budget cycles and administrative procedures rather than by local harvest conditions. Thus a donor cannot time a canal to the moment when harvest yields are about to peak. Second, event-study estimates around treatment onset, reported in Section 5, show no systematic pre-treatment divergence between treated and control cells for the crop outcomes: the lead coefficients are small and statistically indistinguishable from zero.

Finally, problems can arise using borders as a cutoff. Concerns raised against African border designs are that borders were drawn arbitrarily but subsequently generated persistent differences in state capacity, ethnic favoritism, infrastructure provision, or land institutions. Nevertheless, all of these factors are level differences and are differenced out using fixed effects. The threats that survive are related to trends. There can be underlying factors that causes two sides of a border to diverge over the sample period even without irrigation aid. I take four such threats seriously. First, irrigation aid may co-locate with other development interventions (roads, agricultural extension, energy, or broader rural programs), so that the estimated effect would bundle irrigation with complementary aid; I address this by controlling for the effect of disbursements in other aid sectors within the same buffers. Second, conflict may appear asymmetrically across the border and simultaneously deter aid placement and depress agricultural activity. I show robustness on the subsample of segments without severe conflict, based on UCDP event data. Third, climatic trends, in particular drought incidence, may evolve differently across a border where it coincides with an agro-ecological gradient, and may itself attract irrigation aid. I examine this with the SPEI-based drought analysis. I examine this by restricting my sample to drought-prone areas using the SPEI-based drought index. Finally, national agricultural policies such as input-subsidy programs or price supports may evolve differentially on one side of the border. The segment-by-year fixed effects absorb these only insofar as they represent level shifts common to the segment.

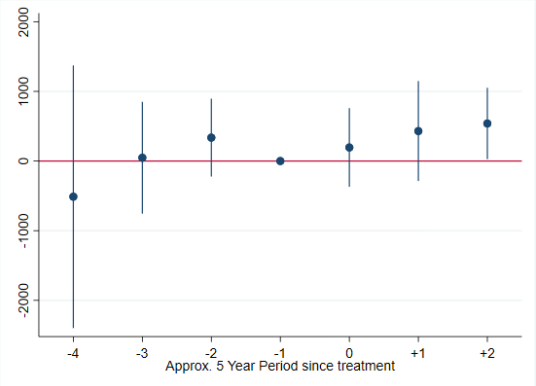
To address this, one could add country-by-year fixed effects. Including the entire set of fixed effects, however, absorbs nearly all of the variation in treatment. I therefore report these results in Table 14 in Appendix A.6.

5. Results

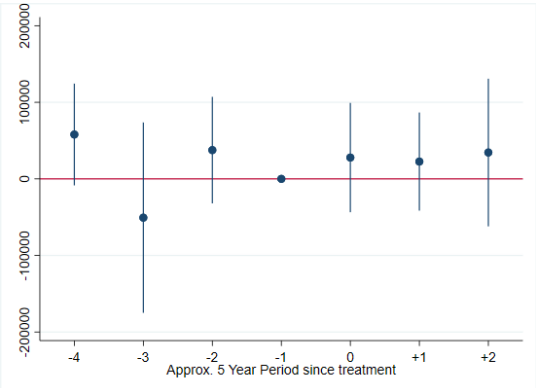
5.1. Event-Study



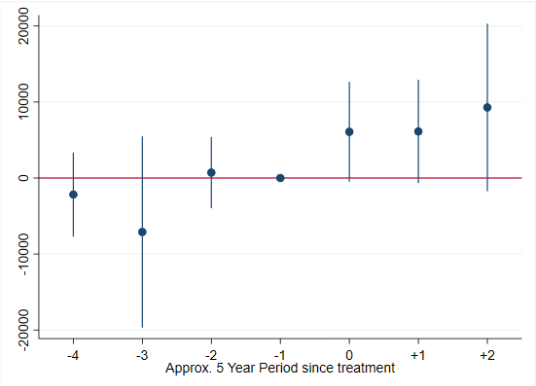
(a) Staple Crop Yield (kg/ha)



(b) Cash Crop Yield (kg/ha)



(c) Production Value (Staple Crops)

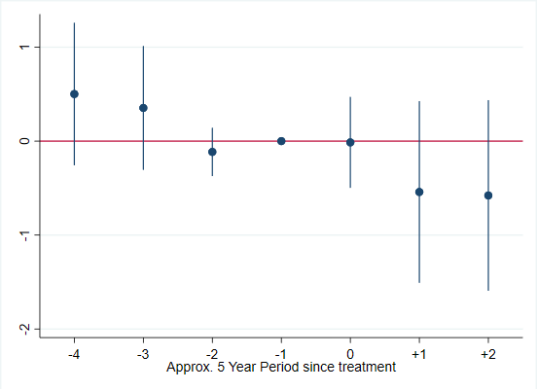


(d) Production Value (Cash Crops)

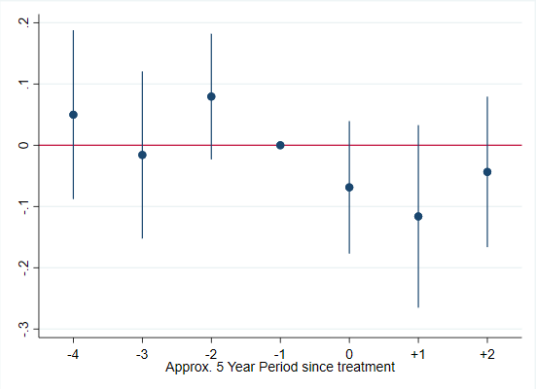
Notes: The figure reports event-study estimates of irrigation effects on the main agricultural outcomes of interest (95% confidence intervals). Panels (a) and (b) show effects on harvest yields measured in kilograms per hectare for staple and cash crops. Panels (c) and (d) report effects on the production value per capita generated from staple and cash crops. Coefficients are normalized relative to the panel period immediately preceding treatment.

Figure 5: Event-Study Plots - Main Outcomes

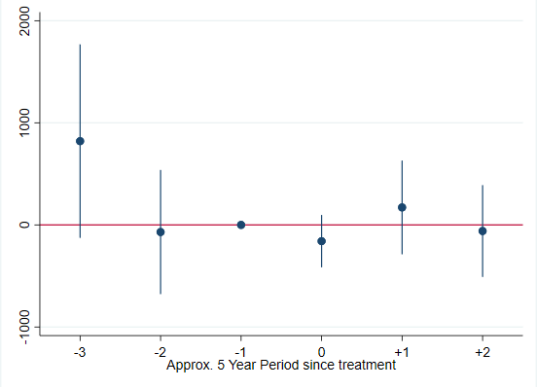
As a first result, I present event-study plots for the main outcome variables (Figure 5) as well as for a set of control variables (Figure 6). These plots illustrate the dynamic effects of irrigation projects before and after treatment. The event-study coefficients correspond to intervals of approximately five years, reflecting the structure of the panel, which covers the years 2005, 2010, 2017, and 2020. Consequently, the coefficient for period +1 captures the effect of an irrigation project that has been active for approximately five years, active meaning once the implementation begins rather than at project completion.



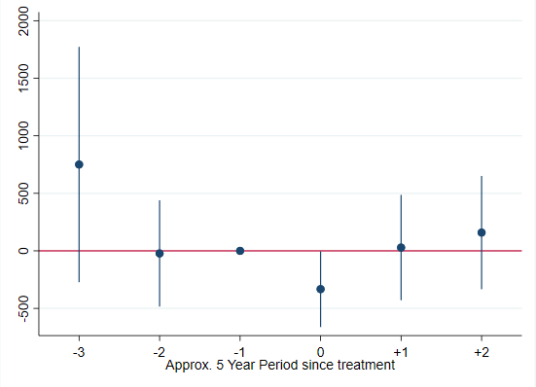
(a) Nighttime Lights



(b) Log of Total Population



(c) Precipitation



(d) Total Biomass

Notes: The figure reports event-study estimates of irrigation effects on agro-ecological and economic outcomes (95% confidence intervals). Panels (a) and (b) show effects on nighttime lights and population (logarithmized) respectively. Panels (c) and (d) report effects on precipitation and total biomass. Coefficients are normalized relative to the year immediately preceding treatment.

Figure 6: Event-Study Plots - Controls

Figure 5 Panel (a) shows the effects on staple crop yields measured in kilograms per hectare, while Panel (b) presents the corresponding effects for cash crops. Panels (c)

and (d) report the effects on agricultural production using the index-weighted production value for staple and cash crops, respectively. Across all four outcomes, the pre-treatment coefficients remain insignificant, supporting the parallel-trends assumption for both yields and production value. For staple crops, the estimates are relatively noisy and show no clear pattern after project implementation, neither for yield nor for production value. The estimates for cash crops, by contrast, display a more consistent upward pattern following treatment. Cash crop yield becomes statistically significant at the 5 percent level in period 2, roughly ten years after project implementation. For production value, the coefficients are consistently positive but do not reach significance at the 5 percent level, though all post-treatment coefficients are significant at the 10 percent level.

To demonstrate that the observed effects on cash crops are not driven by spurious correlations, Figure 6 presents event-study estimates for several control variables that should remain largely unaffected by irrigation aid. Specifically, I examine night-time lights, population, precipitation, and total biomass. None of these variables shows evidence of differential pre-trends, and there are no persistent treatment effects after irrigation projects become active. Although there is a small negative and statistically significant effect on total biomass in the first post-treatment period, this effect disappears over time.

This pattern sheds light on the mechanisms through which irrigation affects crop choice and agricultural outcomes, consistent with the conceptual framework in Section 1. The theory predicts that irrigation primarily facilitates a shift toward higher-value crops rather than uniform productivity gains across all crops. Taken together, the absence of pre-trends, the strong balance in baseline characteristics, and the stability of placebo outcomes support the validity of the empirical strategy.

5.2. Effects on Land Use

Table 3 establishes the baseline specification by gradually introducing fixed effects, sample restrictions, and weighting procedures. The dependent variable is the share of grid-cell area devoted to staple (Panel A) or cash crops (Panel B). Column (1) presents a baseline specification without fixed effects. The coefficients are positive, suggesting an increase in the share of cropped area, although the effects are not statistically significant. Column (2) incorporates grid cell fixed effects. The coefficients remain positive

Table 3: Identifying Main Specification

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Share staple crop area						
L5.Treated (25 km)	0.016	0.021**	0.002	0.002	0.000	-0.006
	(0.016)	(0.009)	(0.007)	(0.007)	(0.008)	(0.008)
Observations	32,432	32,432	32,432	31,884	31,648	31,648
R^2	0.00	0.73	0.78	0.78	0.77	0.77
Mean DV	0.048	0.048	0.048	0.048	0.048	0.048
Panel B: Share cash crop area						
L5.Treated (25 km)	0.011	0.013*	0.004	0.004	0.006	0.003
	(0.010)	(0.007)	(0.005)	(0.005)	(0.006)	(0.006)
Observations	32,432	32,432	32,432	31,884	31,648	31,648
R^2	0.00	0.63	0.72	0.72	0.71	0.70
Mean DV	0.015	0.015	0.015	0.015	0.015	0.015
Grid-cell FE		✓	✓	✓	✓	✓
Border × year FE			✓	✓	✓	✓
≤100 km border				✓	✓	✓
CEM					✓	✓
Border-distance						✓

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column reports the treatment effect for grid cells within 25 km of an irrigation aid project, lagged five years. The dependent variable is an indicator for the share of physical staple (Panel A) or cash (Panel B) crop area. Standard errors, clustered at the Border Segment, in parentheses. Column (2) and Column (3) include fixed effects. Column (4) restricts the sample to grid cells with a maximum distance of 100 kilometers to the border. Column (5) apply coarsened exact matching weights with exact matching on border segment. Column (6) additionally weights by distance to the border.

and become significant. However, once border segment × year fixed effects are included (Column (3)), the coefficients decrease in magnitude and become insignificant. Columns (4)–(6) then sequentially introduce the restrictions and weighting procedures used in the main specification. Column (4) restricts the sample to grid cells located within 100 kilometers of a national border. Column (5) additionally applies the CEM weights to improve the comparability of treated and control grid cells. Finally, column (6) incorporates distance-to-border weights, assigning greater weight to observations located closer to the border. Thus, the specification in column (6) represents the preferred identification strategy: it combines grid-cell fixed effects and border segment × year fixed effects with the restricted-border sample, CEM weights, and distance-to-border weights.

Across columns (3)–(6), the estimated coefficients remain small in magnitude. While the point estimates suggest a reduction in the share of land allocated to staple crops and an increase in the share allocated to cash crops, neither effect is statistically significant. The preferred specification in column (6) therefore provides no evidence that irrigation availability changes the share of land devoted to either staple or cash crops.

Table 4: Land Use

	Cropping Intensity		Diversity		Forest Loss
	(1) Staple	(2) Cash	(3) Staple	(4) Cash	(5) > 20% FC
L5.Treated (25 km)	-0.115**	-0.033	-0.067	-0.018	-116
	(0.054)	(0.021)	(0.043)	(0.044)	(238)
Observations	20,067	18,500	31,648	31,648	6,564
R^2	0.90	0.79	0.76	0.74	0.64
Mean DV	1.157	1.072	0.505	0.292	599
Grid Cell FE	✓	✓	✓	✓	✓
Year × Border FE	✓	✓	✓	✓	✓

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column reports the treatment effect for grid cells within 25 km of an irrigation aid project, lagged five years. Column (1) and Column (2) report the cropping intensity for staple and cash crops. Column (3) and Column (4) evaluate the extensive margin (presence) of staples and cash crops. Column 5 assesses environmental impacts via forest loss. Standard errors, clustered at the border segment, in parentheses. All specifications restrict the sample to grid cells within a 100 kilometer distance to the border, and apply CEM weights and distance to border weights.

In addition to the share of land allocated to crops, Table 4 examines potential effects on cropping intensity, crop diversity, and forest loss using the preferred specification. Cropping intensity for staple crops decreases by 0.115, corresponding to a decline of approximately 10 percent relative to the sample mean, and is statistically significant

at the 5 percent level. By contrast, the corresponding estimate for cash crops is also negative but smaller in magnitude and statistically insignificant. The crop diversity measures reported in columns (3) and (4) likewise suggest negative effects; however, none of the estimates are statistically distinguishable from zero.

To assess potential environmental consequences of irrigation expansion, column (5) examines forest loss. The estimate is small and statistically insignificant, providing no evidence that irrigation aid contributes to forest clearing. This finding is unsurprising, as most irrigation projects in the sample are located in areas with little forest cover, limiting the scope for irrigation-induced deforestation.

5.3. Effects on Harvest Yields

Table 5: Effects on Harvest Yields

	kilograms per hectare		Production Value (Index)	
	(1) Staple	(2) Cash	(3) Staple	(4) Cash
L5.Treated (25 km)	-1,542	-159	22,786	5,075**
	(1,016)	(234)	(18,299)	(2,078)
Observations	27,236	27,236	27,204	27,204
R^2	0.80	0.72	0.70	0.62
Mean DV	7,441	1,152	124,135	6,166
Gridcell FE	✓	✓	✓	✓
Year × Border FE	✓	✓	✓	✓

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column reports the treatment effect for grid cells within 25 km of an irrigation aid project, lagged five years. Column (1) and Column (2) report harvest yields in kg per hectare for staple and cash crops. Column (3) and Column (4) evaluate the production value generated per capita for staples and cash crops. Standard errors, clustered at the border segment, in parentheses. All specifications restrict the sample to grid cells within a 100 kilometer distance to the border, and apply CEM weights and distance to border weights.

Table 5 reports the central outcomes, distinguishing between harvest yields measured in kilogram per hectare and production value generated from harvest per capita. All specifications include the sample restrictions based on the distance to grid cells not exceeding 100 kilometers, as well as the CEM and distance-weighted specifications. While harvest yields measured in kilograms per hectare show negative point estimates after project implementation, these effects are statistically insignificant for both staple and cash crops. In contrast, the production value index indicates positive effects for both

crop types, with a statistically significant increase for cash crops at the 5 percent level. Relative to the sample mean, this corresponds to an increase of approximately 80 percent in cash-crop production value.

Taken together, these findings suggest that irrigation investments primarily affect agricultural intensification rather than expansion of cultivated land. While irrigation does not increase the area allocated to crops, it has the potential to raise the economic returns generated from existing agricultural land, particularly through cash crops.

6. Robustness Checks

Table 6: Robustness Checks

	Staple Yield (kg/ha)			Cash Crop Prod. Value (Index)		
	(1) Placebo	(2) Prec 1	(3) 10Y	(4) Placebo	(5) Prec 1	(6) 10Y
L5.Treated (25 km)	-427	-1,772	-1,542	5	5,173**	5,075**
	(609)	(1,132)	(1,016)	(5)	(2,316)	(2,078)
Observations	26,404	25,116	27,236	26,344	25,084	27,204
R^2	0.67	0.81	0.80	0.58	0.61	0.62
Mean DV	4,666	7,412	7,441	1,755	6,356	6,166
Grid Cell FE	✓	✓	✓	✓	✓	✓
Year × Border FE	✓	✓	✓	✓	✓	✓

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) and Column (4) report placebo regressions using border segments that never receive irrigation treatment and are located between 150 and 400 kilometres from the nearest irrigation project. Column (2) and Column (5) restrict the sample to border segments where all treated grid cells originate from Precision 1 geocoded aid projects.Column (3) and Column (6) restrict the sample to irrigation projects that have been active for approximately ten years. Standard errors, clustered at the border segment, are reported in parentheses.

To ensure that the results are not driven by the particular treatment definition or the composition of the sample, I carry out a series of robustness checks focusing on the two main outcomes identified in the baseline results: staple crop yields and cash crop production value. I assess the robustness of these findings to alternative sample definitions, treatment construction, bandwidth choices, treatment timing, and potential violations of the parallel-trends assumption.

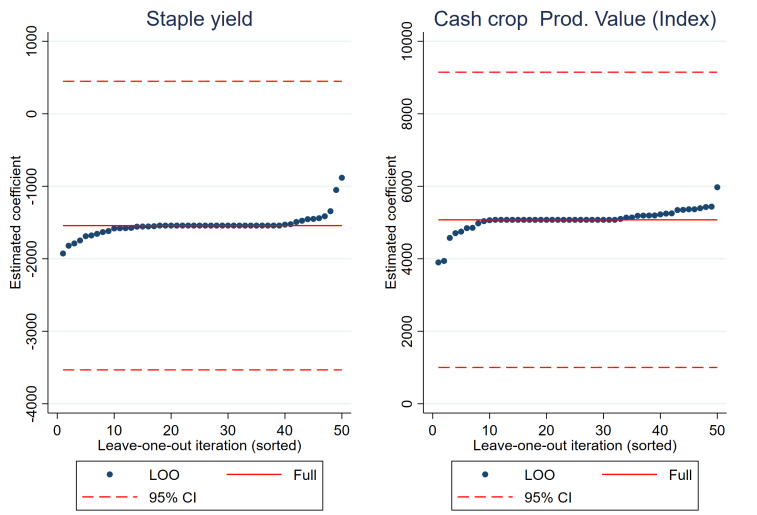
First, I examine whether the results could be explained by spurious correlations, measurement error in treatment assignment, or specific features of the sample construction.

To verify that the estimated treatment effects are reliable, I construct a placebo sample and test whether regions that never receive irrigation treatment also show no treatment effects. To avoid contamination from nearby irrigation projects, I restrict the placebo sample to segments sufficiently distant from treated areas while maintaining comparable border environments.[17] This procedure yields 40 placebo border segments across 20 countries, comprising 2,490 placebo-treated and 4,096 placebo-control grid cells. Second, I restrict my sample to irrigation aid that is precisely measured, by dropping all aid that has a precision code of 2. Finally, since the effectiveness of irrigation infrastructure may decline over time due to maintenance constraints, I restrict the treatment definition to projects that have been active for no longer than ten years. The placebo test addresses whether the estimated effects could reflect broader differences between border regions rather than irrigation exposure. The precision restriction reduces concerns that results are driven by measurement error in project locations, while the duration restriction examines whether the findings depend on projects with uncertain long-run functionality.

Column (1) and Column (4) of Table 6 report the results for the placebo sample. Neither staple yields nor the production value for cash crops show a significant effect. Column (2) and Column (5) report results for the more precisely measured irrigation sites. The effects remain stable and comparable to the baseline specification. Column (3) and Colum (6) show the ten year duration restriction of irrigation projects. Similarly, the coefficients and significance remains stable in comparison to the baseline.

To analyze whether no single border segment drives the estimates, I re-estimate the baseline specification by leaving out one of my border segments in each regression. Given that identification comes from variation across border segments, a concern is that the results may be driven by a small number of large irrigation projects or specific country pairs. Figure 7 plots the resulting coefficients against the full-sample estimate and its 95% confidence interval. For both outcomes, the leave-one-out estimates form a tight cluster around the full-sample coefficient and remain well within its confidence interval throughout, indicating that the results are not driven by any individual segment.

[17] The placebo sample is constructed from border segments satisfying three conditions: (i) no irrigation project is located within any grid cell of the segment, (ii) the segment is not included in the main estimation sample, and (iii) the nearest irrigation project in either bordering country is located between 150 and 400 kilometers away. The lower bound reduces concerns about spillovers from nearby irrigation projects, while the upper bound avoids selecting border segments that are geographically disconnected from the treated regions. Distances are calculated using grid-cell centroids.



Notes: The figure reports the baseline specifications by excluding one border segment at a time. The coefficients are sorted by magnitude against the full-sample estimate and its 95% confidence interval.

Figure 7: Coefficient stability: Leave-one-border-segment-out

Table 7: Sensitivity to Border Bandwidth

	Distance to border				
	≤70 km	≤80 km	≤90 km	≤100 km	≤110 km
	(1)	(2)	(3)	(4)	(5)
Panel A: Staple crop yield (kg/ha)					
L5.Treated (25 km)	-1,558	-1,507	-1,518	-1,542	-1,536
	(1,135)	(1,070)	(1,030)	(1,016)	(1,014)
Observations	20,876	23,404	25,624	27,236	27,656
R^2	0.81	0.81	0.80	0.80	0.80
Mean DV	7,608	7,544	7,483	7,441	7,423
Panel B: Cash crop production value (index)					
L5.Treated (25 km)	5,193**	5,301**	5,218**	5,075**	5,078**
	(2,147)	(2,087)	(2,068)	(2,078)	(2,077)
Observations	20,864	23,392	25,604	27,204	27,624
R^2	0.62	0.62	0.62	0.62	0.62
Mean DV	5,728	5,893	6,073	6,166	6,200

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification on the sample of grid cells within the indicated distance of the national border. The baseline sample corresponds to column (4), ≤100 km. All specifications include grid-cell and border-segment × year fixed effects and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.

The border design relies on comparisons between neighboring regions. The chosen bandwidth should therefore yield stable estimates when varying the bandwidth. Table 7 provides this sensistivity analysis and varies the maximum distance to the national border from 70 to 110 km. Panel A demonstrates the effects for staple crop yields while Panel B focuses on production value for cash crops. The cash crop estimate remains stable and significant at the five percent level for all specifications. The staple yield coefficient is similarly stable in magnitude and remains statistically insignificant throughout. The negative point estimate for staple yields is stable across bandwidth choices, although it remains statistically indistinguishable from zero.

Table 8: Sensitivity to Treatment Radius

	Distance to aid project					
	5 km	15 km	25 km	35 km	45 km	55 km
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Staple crop yield (kg/ha)						
L5.Treated	-2,187**	-1,831	-1,542	-1,392	-1,376	-1,474
	(781)	(1,097)	(1,016)	(1,071)	(1,003)	(1,052)
Observations	12,844	21,916	27,236	29,988	32,412	34,412
R^2	0.82	0.80	0.80	0.82	0.81	0.81
Mean DV	8,138	7,834	7,441	7,854	7,834	7,791
Panel B: Cash crop production value (index)						
L5.Treated	7,419**	6,017**	5,075**	5,198**	6,003**	6,886**
	(2,948)	(2,350)	(2,078)	(2,135)	(2,459)	(2,787)
Observations	12,828	21,884	27,204	29,956	32,380	34,380
R^2	0.56	0.57	0.62	0.62	0.63	0.63
Mean DV	7,468	6,148	6,166	6,861	6,666	6,684

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification with treatment being active for grid cells within the indicated distance of an irrigation aid project, lagged five years. The baseline corresponds to column (3), 25 kilometers. Rings are cumulative, so larger radii nest smaller ones. All specifications include grid-cell and border-segment $\times$ year fixed effects, include grid cells with a maximum distance of 100 kilometers to the border, and apply CEM $\times$ border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.

A related concern is that the spatial definition of treatment may capture areas that are too distant from irrigation projects to plausibly respond. Given that the aid database only provides a single geolocation of the project site but no clear defenition how large the affected region is or which exact areas receive treatment, I therefore vary the treatment radius. Table 8 varies the treatment radius from 5 to 55 km. Again Panel A measures the effect on staple crop yields and Panel B the cash crop production value. The estimate for cash crop is significant at every radius at the 5 percent significance level. The

estimates are stable. The persistence of the cash crop effect across alternative treatment radii suggests that the estimated gains are not an artifact of a particular spatial treatment definition. The larger negative and significant effect for staple yields within 5 km suggests that any reduction in staple yields is concentrated near irrigation infrastructure. Nevertheless, the significance disappears once larger radii are considered and attenuates monotonically as the ring widens. Thus, this finding needs to be interpreted with caution.

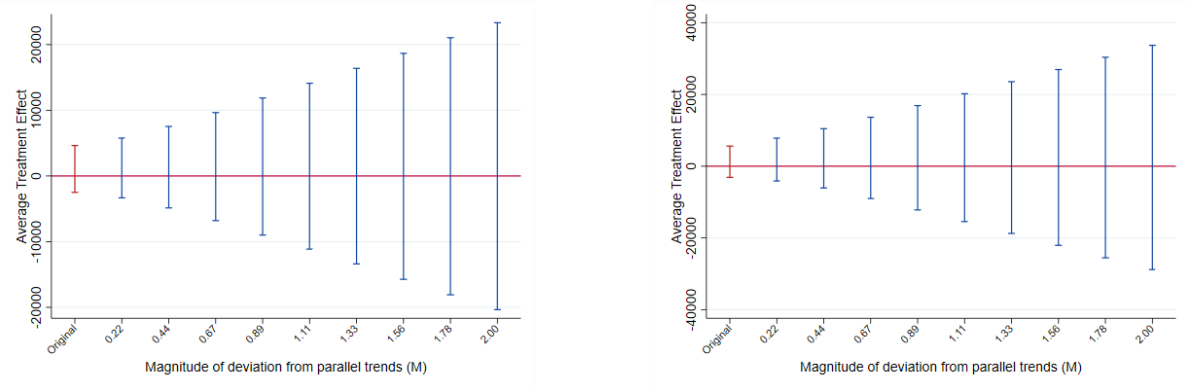
Table 9: Sensitivity to Treatment Lag

	Years since project commitment (minimum project age)						
	L1	L2	L3	L4	L5	L6	L7
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Staple crop yield (kg/ha)							
Treated (25 km)	-1,226	-1,179	-1,336	-1,274	-1,542	-1,641	-1,343
	(748)	(750)	(915)	(951)	(1,016)	(1,100)	(1,149)
Observations	27,236	27,236	27,236	27,236	27,236	27,236	27,236
R^2	0.80	0.80	0.80	0.80	0.80	0.80	0.80
Mean DV	7,454	7,435	7,450	7,445	7,441	7,453	7,489
Panel B: Cash crop production value (index)							
Treated (25 km)	7,213**	5,546***	5,796**	5,929**	5,075**	6,167**	5,451**
	(2,725)	(2,030)	(2,222)	(2,272)	(2,078)	(2,419)	(2,188)
Observations	27,204	27,204	27,204	27,204	27,204	27,204	27,204
R^2	0.62	0.62	0.62	0.62	0.62	0.62	0.62
Mean DV	6,218	6,198	6,181	6,179	6,166	6,151	6,136

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification varying the time lag (the time an irrigation site needs to start operating). The baseline corresponds to column (5), five years. All specifications include grid-cell and border-segment × year fixed effects, include grid cells with a maximum distance of 100 kilometers to the border, and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.

Finally, because irrigation infrastructure may require time before becoming operational, I examine whether the results depend on the assumed timing of treatment exposure. Table 9 varies the minimum time a project needs to start operating and illustrates treatment effects ranging from one to seven years. Cash crop estimates are significant across all seven horizons and have no systematic trend, suggesting that the effect is not driven by a particular choice of treatment lag. A slightly more significant effect appears for the second lag, but drops to 5% significance thereafter. The staple yield coefficient is negative at every lag, but remains insignificant throughout. This reduces concerns that the estimated effect reflects either premature measurement of projects that were not yet operational or the inclusion of projects whose effects had already weakened.

Across all four adjustments (leaving-one-segment-out, bandwidth, treatment area and treatment window) the cash crop production value effect is stable in magnitude and consistently significant, while the staple yield effect is stable in sign and magnitude but consistently imprecise. I therefore treat the production value gain as the robust finding. The effect on staple crops is consistently stable across all specifications but insignificant, which could be the case due to noise but also potential reallocation from staple crops to cash crops once irrigation is available.



(a) Staple Crop Yield (kg/ha) (b) Production Value (Cash Crops)

Notes: The figure reports sensitivity bounds based on the Honest DiD procedure of Rambachan & Roth (2023). The parameter M governs the magnitude of post-treatment deviations from parallel trends that are allowed relative to the largest deviation observed in the pre-treatment period. M=0 corresponds to exact parallel trends, while larger values permit progressively larger violations. The plotted intervals represent 95% robust confidence intervals for the treatment effect under the corresponding restriction on trend deviations.

Figure 8: Honest DiD

To assess whether the main conclusions are sensitive to potential violations of the parallel-trends assumption, I apply the sensitivity analysis proposed by Rambachan & Roth (2023). Figure 8 reports confidence intervals for the treatment effects on staple crop yields and cash crop production value under increasingly large deviations from parallel trends. The parameter M scales the allowed post-treatment deviation relative to the largest deviation observed in the pre-treatment period.

For both outcomes, the confidence intervals include zero under the assumption of exact parallel trends ($M = 0$) and remain consistent with zero as increasingly larger deviations are allowed. The estimated effects are therefore statistically indistinguishable from zero throughout the range of considered violations and the statistically detectable effects is not driven by a specific assumption of parallel pre-trends. However, the interpretation of this exercise requires caution because my panel contains only four periods,

providing a limited number of pre-treatment observations. Consequently, the maximum pre-treatment deviation used to calibrate M is estimated with considerable uncertainty. I therefore view this analysis as complementary evidence rather than a replacement for the event-study results, which provide the primary assessment of the identifying assumption.

7. Heterogeneity and Mechanisms

Table 10: Heterogeneity across Subsamples

	No conflict (1)	Rural (2)	Drought (3)	< disb. (4)	> disb. (5)	> yield (6)	< yield (7)
Panel A: Staple crop yield (kg/ha)							
L5.Treated	-1,836*	-1,323	-380	-112	-1,542	-1,257	-1,730*
	(1,046)	(1,174)	(865)	(913)	(1,016)	(2,225)	(896)
Observations	24,840	21,632	4,636	15,900	27,236	6,804	20,432
R^2	0.81	0.79	0.84	0.83	0.80	0.82	0.70
Mean DV	7,423	7,309	7,672	8,082	7,441	15,498	4,778
Panel B: Cash crop production value (index)							
L5.Treated	5,557**	4,951**	4,839	5,439*	5,075**	11,633**	2,154
	(2,276)	(2,331)	(2,962)	(2,664)	(2,078)	(3,948)	(2,050)
Observations	24,808	21,600	13,052	15,884	27,204	6,568	20,636
R^2	0.62	0.56	0.63	0.62	0.62	0.59	0.62
Mean DV	6,130	5,727	5,779	6,118	6,166	12,616	4,124

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification on the indicated subsample: cells outside regions with severe conflict (1), rural cells (2), drought-prone cells (3), cells in regions with below- (4) and above-median (5) cumulative irrigation disbursement and cells in high (6) and low (7) yield regions based on the year 2005. All specifications include grid-cell and border-segment × year fixed effects, include grid cells with a maximum distance of 100 kilometers to the border, and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.

To understand the channels through which irrigation affects agricultural outcomes, I examine heterogeneity along dimensions that may shape farmers' ability and incentives to respond to improved water availability: exposure to conflict, remoteness, climatic constraints, treatment intensity, and baseline agricultural productivity.

Table 10 illustrates the heterogeneity analysis. Column (1) excludes regions experiencing severe conflict.[18] Since only a small share of the sample is affected by severe conflict, the

[18]Severe conflicts are identified using UCDP conflict data (Pettersson & Davies, 2025). A severe conflict

estimated effect on cash crop production value remains stable. In contrast, staple yields become marginally significant at the 10 percent level once conflict-affected areas are excluded. This suggests that conflict-related disruptions may attenuate the estimated response in the full sample.

Column (2) focuses on rural areas, defined as locations requiring at least 90 minutes but no more than 360 minutes to reach the nearest city. In this subsample, the estimated effect on cash crop production value increases slightly with respect to the sample mean. Similarly, although not statistically significant, the negative estimate for staple crop yields becomes smaller, consistent with irrigation generating larger agricultural responses in locations where farming represents a more central economic activity.

Column (3) restricts the sample to areas that are affected by severe droughts, measured using the SPEI index. Roughly one-third of the sample belongs to the drought-prone areas in the sample, nevertheless, irrigation does not seem to show significant effects for these regions, leading to the conclusion that the effects are driven by areas that do not experience many drought periods.

Column (4) and Column (5) distinguish between low- and high-disbursement areas, while Column (6) and Column (7) separate regions with low and high pre-treatment yields. In both cases, high and low are defined relative to the sample median. The coefficients for high and low disbursement regions are comparable and do not differ from the main specification. The similarity across disbursement groups suggests that the baseline effect is not driven exclusively by either small- or large-scale irrigation investments. For staple crops, the point estimate in low-disbursement areas is closer to zero than in high-disbursement areas, suggesting that any negative effect may be weaker where irrigation investments are smaller. However, neither estimate is statistically significant, and the difference between the two groups should therefore be interpreted with caution. One possible explanation is that larger irrigation investments induce stronger changes in land-use patterns, potentially including a greater reallocation away from staple production. This mechanism cannot be directly tested with the available data. Given that the main estimates for staple-crop yields, together with the robustness and sensitivity analyses, remain statistically indistinguishable from zero, these results should be interpreted as suggestive heterogeneity rather than evidence of a causal effect on staple-crop yields.

is defined as a border segment experiencing at least 1,000 fatalities over the observation period (2005 - 2020).

Finally, cash crop yields appear to increase more in areas that already exhibited relatively high agricultural productivity prior to treatment, while there is no significant effect for low yield areas for cash crops. In contrast, staple crop yields decline by approximately 26 percent in low-yield regions after irrigation projects become operational, an effect that is also significant at the 10 percent level, while there are no significant effects available for high yield regions. The stronger response of cash crop production value in initially productive areas, combined with the decline in staple yields in low-productivity regions, is consistent with irrigation inducing a reallocation of agricultural resources toward higher-value crops rather than uniformly increasing staple production.

Table 11: Co-located Aid and the Effect of Irrigation Aid

	Co-located aid category					
	Business	Industry	Agri.	Foodaid	Envir.	Econ. infra.
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Staple crop yield (kg/ha)						
L5.Treated (25 km)	-1,129	-1,679	-1,415	-1,584	-1,542	-1,580
	(868)	(1,017)	(1,058)	(1,034)	(1,045)	(1,016)
Co-located aid	675	-300	-375	-945	-1,052	1,034
	(748)	(713)	(342)	(1,062)	(701)	(744)
Treated × Co-located aid	-10,049**	2,378	-81	1,869*	-208	3,501***
	(4,293)	(1,488)	(1,015)	(1,090)	(1,131)	(1,012)
Observations	27,236	27,236	27,236	27,236	27,236	27,236
R^2	0.80	0.80	0.80	0.80	0.80	0.80
Mean DV	7,441	7,441	7,441	7,441	7,441	7,441
Panel B: Cash crop production value (index)						
L5.Treated (25 km)	5,401**	5,166**	5,729*	5,167**	5,393**	5,204**
	(2,141)	(2,147)	(2,941)	(2,121)	(2,167)	(2,072)
Co-located aid	895	481	539	429	348	6,608
	(880)	(1,578)	(1,163)	(1,540)	(1,726)	(5,048)
Treated × Co-located aid	-8,125***	-1,627	-1,299	-2,902	-5,116**	-7,178
	(2,525)	(2,894)	(2,770)	(1,895)	(2,527)	(8,692)
Observations	27,204	27,204	27,204	27,204	27,204	27,204
R^2	0.62	0.62	0.62	0.62	0.62	0.62
Mean DV	6,166	6,166	6,166	6,166	6,166	6,166

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column interacts the baseline treatment indicator with a dummy for the presence of co-located aid projects of the indicated category. All specifications include grid-cell and border-segment × year fixed effects, include grid cells with a maximum distance of 100 kilometers to the border, and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.

Africa receives aid projects across a wide range of sectors, and many irrigation projects are located close to other aid interventions. I examine whether the estimated effect

of irrigation survives while accounting for nearby projects not focusing on irrigation.[19] Table 11 interacts irrigation treatment with indicators for co-located aid by sector. The direct effect of irrigation on cash crop production value remains positive and statistically significant across specifications, suggesting that the main finding is not driven by other nearby aid interventions. Similarly, the null effects for staple crops are consistent when including co-located aid. However, strong increases in crop yields for staples appear when a project focusing on economic infrastructure is available nearby. The exception is aid flowing to the business sector. In areas where business-oriented aid is present, the positive effect on production value for cash crops reverses sharply. The interaction term is negative and statistically significant at the 1 percent level, more than offsetting the direct treatment effect. A similar pattern emerges for the corresponding interaction with staple crop yields. The presence of nearby business aid therefore appears to reverse the returns to irrigation aid, although this estimate should be interpreted with caution given that the location of complementary aid interventions is not randomly assigned. These findings suggest that the composition of the surrounding aid portfolio, and not only irrigation itself, shapes the agricultural response. From a policy perspective, this is an important result. Understanding which aid projects are already present in the surrounding area may have important implications for the overall effectiveness of newly implemented irrigation projects.

8. Conclusion

This study provides micro-level evidence on the effects of irrigation-focused development aid on agricultural outcomes in sub-Saharan Africa. Using a spatial border difference-in-differences design, I compare neighboring grid cells across national borders that share similar local agro-ecological conditions but differ in exposure to donor-funded irrigation projects. The results point to a tendency for irrigation aid to raise the production value of cash crops, suggesting that improved water access may support agricultural commercialization, while I find no detectable response in staple crop yields. Together, these findings suggest that irrigation investments primarily strengthen market-oriented agricultural production rather than generating broad-based increases in staple crop productivity. The absence of detectable effects on staple crops should not be interpreted as

[19] A nearby project is defined as an aid project located within 50 kilometers of an irrigation site

evidence of no impact, but rather as evidence that any effects are not precisely estimated in the available data.

Methodologically, the paper develops a localized difference-in-differences strategy that combines georeferenced aid data, satellite-based agricultural outcomes, and spatial variation near international borders to evaluate development interventions. Identification exploits the timing at which irrigation projects become operational, which is driven largely by donor procurement and construction cycles rather than by local agricultural conditions. While previous evaluations of aid often rely on aggregate outcomes or administrative units, the approach developed here exploits localized variation in project exposure and provides a framework for studying how geographically targeted interventions affect nearby populations and landscapes. The design is particularly useful in settings where detailed administrative or household panel data are limited.

The findings have several policy implications. They suggest that irrigation investments can contribute to agricultural transformation by enabling farmers to increase the value of production, particularly through cash crops. However, the results also highlight that irrigation should not automatically be viewed as a universal solution for raising staple crop productivity or food production. The effectiveness of irrigation depends on the surrounding agricultural and development environment, including access to complementary inputs, markets, and other interventions. For policymakers and donors, this emphasizes the importance of considering irrigation investments as part of broader agricultural development strategies rather than as standalone infrastructure projects.

Despite the novel approach, some limitations remain. First, the null effects for staple crops should be read as an absence of detectable effects rather than definitive evidence of no impact, since satellite-based yield measures may carry measurement error and the limited number of border segments reduces the precision with which smaller or heterogeneous effects can be detected. Second, the border design identifies the effect of irrigation aid in the specific environments surrounding international borders rather than for all irrigation projects across sub-Saharan Africa, and the estimation sample contains relatively fewer World Bank projects and a larger share of German development cooperation projects than the broader aid portfolio, which further limits external validity. Third, although the staggered rollout of irrigation projects would motivate a staggered difference-in-differences estimator (Callaway & Sant'Anna, 2021), the four survey waves and the multi-year lag until irrigation becomes operational leave too few post-treatment periods to support group-time effects or an event-study path, so I rely on the two-way

fixed-effects specification. The scope for bias from heterogeneous treatment timing is nonetheless limited, as identification comes overwhelmingly from comparisons between treated and never-treated units.

Several directions could extend this work. First, verifying the findings using within-country variation would strengthen external validity and address the cross-country differences that remain despite the border design. Second, linking the remote-sensing estimates to household surveys such as the Living Standards Measurement Study (LSMS) could provide complementary evidence on mechanisms, including crop choice, input use, market participation, and household welfare. Third, exploiting alternative geographic discontinuities, such as ethnic boundaries, would offer a further robustness check by comparing treated and untreated areas within more homogeneous social and institutional environments. Fourth, incorporating information on irrigation networks and water flows, rather than proximity to the project location alone, would allow future work to capture downstream effects. Finally, a more detailed classification of irrigation interventions, distinguishing scheme type, scale, management structure, and water-delivery system, and combined with project-level cost information, would help identify which investments generate the largest returns and inform cost-benefit assessments for donors, though such analysis is currently constrained by the limited availability of consistent project-level data.

References

Alesina, A., Easterly, W., & Matuszeski, J. (2011). Artificial states. *Journal of the European Economic Association*, 9(2), 246–277. https://doi.org/10.1111/j.1542-4774.2010.01009.x

Ansar, A., Flyvbjerg, B., Budzier, A., & Lunn, D. (2014). Should we build more large dams? the actual costs of hydropower megaproject development. *Energy Policy*, 69, 43–56. https://doi.org/https://doi.org/10.1016/j.enpol.2013.10.069

Beguería, S., Vicente-Serrano, S. M., Reig-Gracia, F., & Latorre Garcés, B. (2024). *Speibase: Standardised precipitation-evapotranspiration index database, global 0.5° gridded dataset, version 2.11.* Dataset, Spanish National Research Council (CSIC) & Estación Experimental de Aula Dei. Available at https://spei.csic.es/database.html (09.02.2026).

BenYishay, A., Goodman, S., Sayers, R., Singh, K., Walker, M., Rauschenbach, M., & Noltze, M. (2023). *Does irrigation strengthen climate resilience.*

Bjornlund, H., Bjornlund, V., & van Rooyen, A. (2022). Why food insecurity persists in sub-saharan africa: A review of existing evidence. *Food Security*, 14(4), 845–864. https://doi.org/10.1007/s12571-022-01256-1

Bjornlund, V., Bjornlund, H., & van Rooyen, A. F. (2020). Exploring the factors causing the poor performance of most irrigation schemes in post-independence sub-saharan africa. *International Journal of Water Resources Development*, 36(sup1), S54–S101. https://doi.org/10.1080/07900627.2020.1808448

Bluhm, R., Dreher, A., Fuchs, A., Parks, B. C., Strange, A. M., & Tierney, M. J. (2025). Connective financing: Chinese infrastructure projects and the diffusion of economic activity in developing countries. *Journal of Urban Economics*, 145, 103730. https://doi.org/https://doi.org/10.1016/j.jue.2024.103730

Bomprezzi, P., Dreher, A., Fuchs, A., Hailer, T., Kammerlander, A., Kaplan, L. C., Marchesi, S., Masi, T., Robert, C., & Unfried, K. (2024). Wedded to prosperity? spousal favoritism in foreign aid and regional development. Working Paper 10969, CESifo. https://doi.org/10.2139/ssrn.4748017

Bulíř, A. & Hamann, A. J. (2003). Aid volatility: an empirical assessment. *IMF Staff papers*, 50(1), 64–89.

Bulíř, A. & Hamann, A. J. (2008). Volatility of development aid: From the frying pan into the fire? *World Development*, 36(10), 2048–2066. https://doi.org/https://doi.org/10.1016/j.worlddev.2007.02.019. Special Section (pp. 2045-2102). The Volatility of Overseas Aid

Burnside, C. & Dollar, D. (2000). Aid, policies, and growth. *American Economic Review*, 90(4), 847–868. https://doi.org/10.1257/aer.90.4.847

Callaway, B. & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. https://doi.org/https://doi.org/10.1016/j.jeconom.2020.12.001. Themed Issue: Treatment Effect 1

Christiaensen, L., Demery, L., & Kuhl, J. (2011). The (evolving) role of agriculture in poverty reduction—an empirical perspective. *Journal of Development Economics*, 96(2), 239–254. https://doi.org/https://doi.org/10.1016/j.jdeveco.2010.10.006

Clemens, M. A., Radelet, S., Bhavnani, R. R., & Bazzi, S. (2011). Counting chickens when they hatch: Timing and the effects of aid on growth. *The Economic Journal*, 122(561), 590–617. https://doi.org/10.1111/j.1468-0297.2011.02482.x

Diao, X., Hazell, P. B., Resnick, D., & Thurlow, J. (2007). *The role of agriculture in development: Implications for Sub-Saharan Africa*, volume 153. Intl Food Policy Res Inst.

Digital Earth Africa Contributors (2021). *Digital elevation models notebook and data derivatives*. Jupyter Notebook / Dataset. https://docs.digitalearthafrica.org/en/latest/sandbox/notebooks/Datasets/Digital_Elevation_Models.html. Accessed: 2026-07-16

Dillon, A. (2011). The effect of irrigation on poverty reduction, asset accumulation, and informal insurance: Evidence from northern mali. *World Development*, 39(12), 2165–2175. https://doi.org/https://doi.org/10.1016/j.worlddev.2011.04.006

Dreher, A., Fuchs, A., Parks, B., Strange, A., & Tierney, M. J. (2021). Aid, china, and growth: Evidence from a new global development finance dataset. *American Economic Journal: Economic Policy*, 13(2), 135–74. https://doi.org/10.1257/pol.20180631

Dreher, A. & Lohmann, S. (2015). Aid and growth at the regional level. *Oxford Review of Economic Policy*, 31(3-4), 420–446. https://doi.org/10.1093/oxrep/grv026

Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American Economic Review*, 101(6), 2350–90. https://doi.org/10.1257/aer.101.6.2350

Duflo, E. & Pande, R. (2007). Dams*. *The Quarterly Journal of Economics*, 122(2), 601–646. https://doi.org/10.1162/qjec.122.2.601

Dyer, J. & Shapiro, J. (2023). Pumps, prosperity and household power: Experimental evidence on irrigation pumps and smallholder farmers in kenya. *Journal of Development Economics*, 163, 103034. https://doi.org/https://doi.org/10.1016/j.jdeveco.2022.103034

Easterly, W., Levine, R., & Roodman, D. (2004). Aid, policies, and growth: Comment. *American Economic Review*, 94(3), 774–780. https://doi.org/10.1257/0002828041464560

Englebert, P. (2009). *Africa: Unity, sovereignty, sorrow*. Lynne Rienner Publishers.

Evenson, R. E. & Gollin, D. (2003). Assessing the impact of the green revolution, 1960 to 2000. *Science*, 300(5620), 758–762. https://doi.org/10.1126/science.1078710

FAO (2022). *Food and agriculture organization of the united nations , increasing agricultural productivity in africa: Can sti help africa to make a quantum leap in agricultural productivity?* https://www.fao.org/science-technology-and-innovation/increasing-agricultural-productivity-in-africa-can-sti-help-africa-to-make-a-quantum-leap-in-agricultural-productivity/en. Accessed: 2025-01-22.

FAO WaPOR (2025). *FAO WaPOR: Water Productivity through Open access of Remotely sensed derived data*. https://data.apps.fao.org/wapor/. Accessed: 2025-01-18

FAOSTAT (2025). *FAOSTAT Producer Prices – Annual*. https://www.fao.org/faostat/en/#data/PP. Accessed: 2024-10-23

Frankema, E., Williamson, J., & Woltjer, P. (2018). An economic rationale for the west african scramble? the commercial transition and the commodity price boom of

1835–1885. *The Journal of Economic History*, 78(1), 231–267. https://doi.org/10.1017/S0022050718000128

Galiani, S., Knack, S., Xu, L., et al. (2017). The effect of aid on growth: evidence from a quasi-experiment. *Journal of Economic Growth*, 22(1), 1–33. https://doi.org/10.1007/s10887-016-9137-4

Hansen, H. & Tarp, F. (2001). Aid and growth regressions. *Journal of development Economics*, 64(2), 547–570.

Hayami, Y. & Ruttan, V. W. (1971). *Agricultural Development: An International Perspective* (revised and expanded ed.). Johns Hopkins University Press.

Herbst, J. (2000). *States and power in Africa: Comparative lessons in authority and control: Comparative lessons in authority and control*. Princeton University Press.

Higginbottom, T. P., Adhikari, R., Dimova, R., Redicker, S., & Foster, T. (2021). Performance of large-scale irrigation projects in sub-saharan africa. *Nature Sustainability*, 4(6), 501–508.

Iacus, S. M., King, G., & Porro, G. (2011). Multivariate matching methods that are monotonic imbalance bounding. *Journal of the American Statistical Association*, 106(493), 345–361. https://doi.org/10.1198/jasa.2011.tm09599

IFPRI (2019). *International food policy research institute , global spatially-disaggregated crop production statistics data for 2010 version 2.0*. Harvard Dataverse. https://doi.org/10.7910/DVN/PRFF8V. V4

IFPRI (2020). *International food policy research institute , spatially-disaggregated crop production statistics data in africa south of the sahara for 2017*. Harvard Dataverse. https://doi.org/10.7910/DVN/FSSKBW. V2

IFPRI (2024). *International food policy research institute , global spatially-disaggregated crop production statistics data for 2020 version 1.0.0*. Harvard Dataverse. https://doi.org/10.7910/DVN/SWPENT. V1

IFPRI & IIASA (2016). *International food policy research institute and international institute for applied systems analysis , global spatially-disaggregated crop production statistics data for 2005 version 3.2*. Harvard Dataverse. https://doi.org/10.7910/DVN/DHXBJX. V9

Inocencio, A. B. et al. (2007). *Costs and performance of irrigation projects: A comparison of sub-Saharan Africa and other developing regions*, volume 109. IWMI.

Isaksson, A.-S. & Kotsadam, A. (2018). Chinese aid and local corruption. *Journal of Public Economics*, 159, 146–159. https://doi.org/https://doi.org/10.1016/j.jpubeco.2018.01.002

Johnson, R. (2005). Agriculture and trade in sub-saharan africa: A review. *Outlook on Agriculture*, 34(2), 71–76. https://doi.org/10.5367/0000000054224382

Jones, M., Kondylis, F., Loeser, J., & Magruder, J. (2022). Factor market failures and the adoption of irrigation in rwanda. *American Economic Review*, 112(7), 2316–52. https://doi.org/10.1257/aer.20210059

Kaya, O., Kaya, I., & Gunter, L. (2012). Development aid to agriculture and economic growth. *Review of Development Economics*, 16(2), 230–242. https://doi.org/https://doi.org/10.1111/j.1467-9361.2012.00658.x

Kaya, O., Kaya, I., & Gunter, L. (2013). Foreign aid and the quest for poverty reduction: Is aid to agriculture effective? *Journal of Agricultural Economics*, 64(3), 583–596. https://doi.org/https://doi.org/10.1111/1477-9552.12023

Kilby, C. (2009). The political economy of conditionality: An empirical analysis of world bank loan disbursements. *Journal of Development Economics*, 89(1), 51–61. https://doi.org/https://doi.org/10.1016/j.jdeveco.2008.06.014

Kilby, C. (2013). The political economy of project preparation: An empirical analysis of world bank projects. *Journal of Development Economics*, 105, 211–225. https://doi.org/https://doi.org/10.1016/j.jdeveco.2013.07.011

Li, X., Zhou, Y., Zhao, M., & Zhao, X. (2023). *Harmonization of dmsp and viirs nighttime light data from 1992–2021 at the global scale*. https://doi.org/10.6084/m9.figshare.9828827.v8

McCauley, J. F. & Posner, D. N. (2015). African borders as sources of natural experiments promise and pitfalls. *Political Science Research and Methods*, 3(2), 409–418. https://doi.org/10.1017/psrm.2014.37

Michalopoulos, S. & Papaioannou, E. (2013). National institutions and subnational development in africa *. *The Quarterly Journal of Economics*, 129(1), 151–213. https://doi.org/10.1093/qje/qjt029

Michalopoulos, S. & Papaioannou, E. (2016). The long-run effects of the scramble for africa. *American Economic Review*, 106(7), 1802–48. https://doi.org/10.1257/aer.20131311

Nelson, A. (2008). *Estimated travel time to the nearest city of 50,000 or more people in year 2000*. Available at https://forobs.jrc.ec.europa.eu/gam/download. Accessed: 25/07/2024.

Pettersson, T. & Davies, S. (2025). *Ucdp georeferenced event dataset (ged) global version 25.1*.

Rajan, R. G. & Subramanian, A. (2008). Aid and growth: What does the cross-country evidence really show? *The Review of Economics and Statistics*, 90(4), 643–665. https://doi.org/10.1162/rest.90.4.643

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

Rockström, J., Karlberg, L., Wani, S. P., Barron, J., Hatibu, N., Oweis, T., Bruggeman, A., Farahani, J., & Qiang, Z. (2010). Managing water in rainfed agriculture—the need for a paradigm shift. *Agricultural Water Management*, 97(4), 543–550. https://doi.org/https://doi.org/10.1016/j.agwat.2009.09.009

Scheumann, W., Houdret, A., & Brüntrup, M. (2017). Unlocking the irrigation potential in sub-Saharan Africa: Are public-private partnerships the way forward? Briefing Paper 7/2017, German Development Institute / Deutsches Institut für Entwicklungspolitik (DIE), Bonn. https://www.idos-research.de/fileadmin/migratedNewsAssets/Files/BP__7.2017.pdf

Schuenemann, F., Thurlow, J., Meyer, S., Robertson, R., & Rodrigues, J. (2018). Evaluating irrigation investments in malawi: economy-wide impacts under uncertainty and labor constraints. *Agricultural Economics*, 49(2), 237–250. https://doi.org/https://doi.org/10.1111/agec.12412

Schultz, T. W. (1964). *Transforming Traditional Agriculture*. Yale University Press.

Sekhri, S. (2014). Wells, water, and welfare: The impact of access to groundwater on rural poverty and conflict. *American Economic Journal: Applied Economics*, 6(3), 76–102. https://doi.org/10.1257/app.6.3.76

Ssozi, J., Asongu, S., & Amavilah, V. H. (2019). The effectiveness of development aid for agriculture in sub-saharan africa. *Journal of Economic Studies*, 46(2), 284–305. https://doi.org/10.1108/JES-11-2017-0324

U.S. Geological Survey (2000). *Shuttle radar topography mission (srtm) 1 arc-second global*. https://doi.org/10.5066/F7PR7TFT

Weiss, D. J., Nelson, A., Gibson, H. S., Temperley, W., Peedell, S., Lieber, A., Hancher, M., Poyart, E., Belchior, S., Fullman, N., Mappin, B., Dalrymple, U., Rozier, J., Lucas, T. C. D., Howes, R. E., Tusting, L. S., Kang, S. Y., Cameron, E., Bisanzio, D., Battle, K. E., Bhatt, S., & Gething, P. W. (2018). A global map of travel time to cities to assess inequalities in accessibility in 2015. *Nature*. https://www.resourcewatch.org. Accessed through Resource Watch, (25.07.2024).

Wesseling, H. L. (1996). Divide and rule: The partition of africa, 1880-1914.

World Bank (2022). *Poverty and Shared Prosperity 2022: Correcting Course*. World Bank. https://www.worldbank.org/en/publication/poverty-and-shared-prosperity. Accessed: 2025-01-12

World Bank (2024). *Measuring Smallholder Farmers' Productivity in Sub-Saharan Africa*. Number 7150 in Policy Research Working Paper. World Bank. https://openknowledge.worldbank.org/handle/10986/40230. World Bank Policy Research Working Paper

WorldPop (2018). *(school of geography and environmental science, university of southampton; department of geography and geosciences, university of louisville; département de géographie, université de namur) and center for international earth science information network (ciesin), columbia university, global high resolution population denominators project*. https://doi.org/10.5258/SOTON/WP00647. Funded by The Bill and Melinda Gates Foundation (OPP1134076)

You, L., Ringler, C., Wood-Sichra, U., Robertson, R., Wood, S., Zhu, T., Nelson, G., Guo, Z., & Sun, Y. (2011). What is the irrigation potential for africa? a combined

biophysical and socioeconomic approach. *Food Policy*, 36(6), 770–782. https://doi.org/https://doi.org/10.1016/j.foodpol.2011.09.001

A. Data

A.1. Variable Definitions

Outcome Variables

All agricultural outcomes are constructed from the Spatial Production Allocation Model (MapSPAM) waves 2005, 2010, 2017, and 2020, aggregated to the analysis grid. Crops are classified into staple crops (cassava, maize, rice, yams, sorghum, millet, and plantain) and cash crops (cotton, cocoa, coffee, tea, groundnut, and oil palm). The staple category comprises the principal calorie sources in sub-Saharan African food balance sheets, while the cash category corresponds to the region's traditional export crops (Johnson, 2005; Frankema et al., 2018; Bjornlund et al., 2022). Groundnut and oil palm are partly consumed domestically; their assignment to the cash category follows their historical role as marketed export commodities.

Yields. Crop yields are measured in kilograms per hectare and aggregated across the crops within each category. To limit the influence of extreme values generated by the downscaling procedure, yields are winsorized at the 5th and 95th percentile within each panel wave.

Physical area. Physical crop area measures the area within a grid cell on which a crop category is cultivated, expressed as a share of the cell and bounded at one. Because physical area does not count repeated harvests, it captures the extensive margin of cultivation. Binary indicators equal one if any positive area of the respective category is cultivated in the cell.

Production value (index). A production value is constructed by multiplying the productivity of a given grid cell (harvest yields $\times$ harvest area) with the producer price index from FAOSTAT (FAOSTAT, 2025) for each particular crop and country. The base period of the indices is the average producer prices of crop products in 2014, 2015 and 2016. This is done individually for each crop and then aggregated to staples and cash crops. To get a measure based on the local population density, the production value index is divided by the population densidy. All value-based outcomes are winsorized at the 5th and 95th percentile within waves.

Crop diversity. Crop diversity is measured by an abundance-based Simpson index, $D = 1 - \sum_c p_c^2$, where p_c is the share of crop c in the cell's total physical crop area. The index is computed over all crops as well as separately within the staple and cash categories and ranges from zero (monoculture) to values approaching one (maximally even allocation).

Forest loss. For the land-use analysis, annual forest-cover loss from **?** is aggregated to the panel structure by summing loss pixels over the years between consecutive waves (for example, the 2010 observation comprises losses over 2005–2009). Initial forest cover is measured as median tree cover in 2000.

Treatment Variables

Treatment. A grid cell is treated if a main-focus irrigation project site with precise geocoding is located within 25 kilometers of the cell and within the same country. Treatment enters the estimation with a five-year lag: Thus treatment equals one from five years after the earliest project start within the buffer, reflecting the multi-year gap between project announcement and operational irrigation infrastructure. A cell first treated in 2005 is therefore considered treated from the 2010 wave onward. Robustness checks vary the lag between one and eight years and the buffer between 5 and 55 kilometers; because the buffers are nested, wider definitions add progressively more distant cells to the treatment group.

Disbursement intensity. The intensive margin is measured by cumulative irrigation-aid disbursements within the buffer. Disbursements are recorded in the project's start year and accumulated over time, so that the stock increases whenever an additional project becomes active within the buffer. Project-level amounts are split evenly across a project's geocoded sites to avoid double-counting multi-location programs.

Alternative aid categories. To assess whether effects are specific to irrigation rather than to aid presence in general, cumulative disbursements are constructed analogously for other aid sectors (among others education, health, water supply, transport, energy, business, industry, other agriculture, environment, and economic infrastructure) from the same geocoded database GODAD, and sector-specific presence indicators are defined as positive cumulative disbursement within the buffer.

Control Variables

Nighttime lights. Nighttime light intensity is taken from the harmonized DMSP-consistent night-light series (**?**), ensuring comparability across the DMSP and VIIRS eras. The indicator entering the analysis is based on the modal digital-number value within the cell.

Population. Population counts are 1km WorldPop (WorldPop, 2018) rasters aggregated to the grid for each panel wave; population density enters in logs.

Roads. Road density is measured as the travel time needed from the corresponding grid cell to the nearest city for 2000 and 2015 (**?**Weiss et al., 2018).

Heterogeneity Variables

Conflict. Conflict exposure is constructed from UCDP (Pettersson & Davies, 2025) georeferenced event data. A border segment is classified as a severe-conflict region if any cell in the segment records high-fatality events during the sample period; the no-conflict subsample excludes these segments.

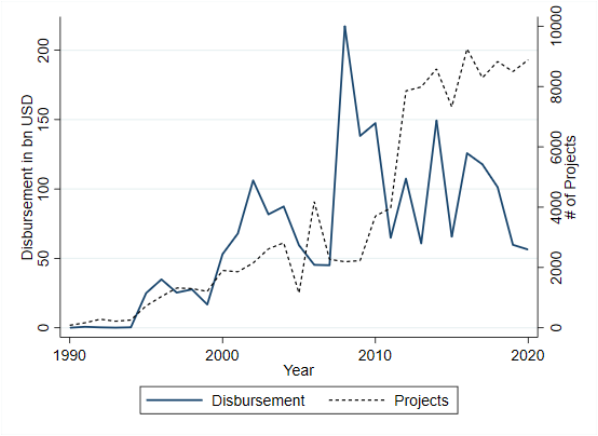
Drought. Drought exposure is based on the Standardized Precipitation-Evapotranspiration Index (SPEI) (Beguería et al., 2024). A treated cell is classified as drought-affected if it records SPEI drought conditions, and the drought subsample comprises border segments containing at least one such cell; the number of drought periods enters as an additional control.

Rurality. A treated cell is classified as rural if its travel time to the nearest city (2017) lies between 90 and 360 minutes; border segments containing such cells constitute the rural subsample.

Disbursement intensity split. High- and low-disbursement subsamples split treated cells at the median of cumulative irrigation disbursements within the 25km buffer.

A.2. Aid Distribution across Africa

For a general overview of aid distribution within Africa this section in the appendix provides additional information obtained from the GODAD database. In total there are 113,140 development aid projects across various sectors, with agriculture being one of the dominant sectors. The number of projects rises steeply over the period, with the sharpest jump around 2012, reaching its peak in 2016. Average disbursement per



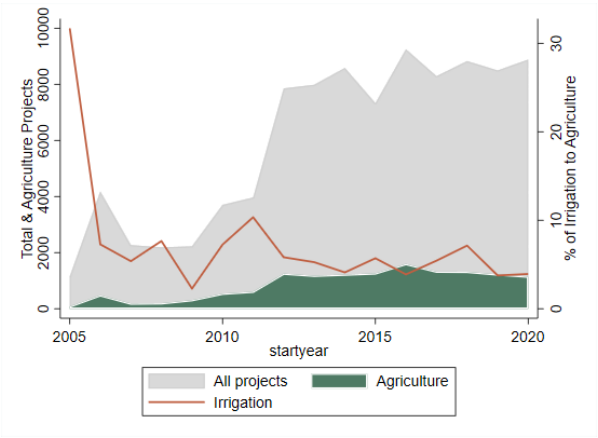
Notes: The figure illustrates the evolution of aid disbursement (solid line) and the amount of aid projects (dashed line) over time in Africa.

Figure 9: Evolution of Aid in Africa

project declines in parallel, from roughly USD 100 million in the peak years of the 2000s to about USD 6 million by the end of the sample. Figure [9] plots annual disbursement against project counts.

The number of active donors grows from 11 in 1990 to 19 donors in 2020. Donor activity is highly concentrated: the five largest financiers, the World Bank, the United States, Germany, Spain, and Italy, together account for roughly sixty-three percent of all projects. Non-DAC financiers such as China and India appear only marginally in the data, at under two percent of projects each. This most likely reflects the database's more limited coverage of non-Western donors rather than their true footprint on the continent, and the general-aid figures reported here should be read with that caveat in mind.

Aid is distributed across a wide range of sectors, but agriculture is the most prominent on both margins. By project count it is the single largest sector at roughly fifteen percent, marginally ahead of health. By cumulative disbursement it again leads, at roughly USD 480 billion, ahead of transport infrastructure, humanitarian assistance, energy, and health. The prominence of agriculture is not static: its share of projects rose from around seven percent in the mid-2000s to a peak near seventeen percent in the mid-2010s, before easing to about thirteen percent by 2020. Irrigation forms a small and clearly delineated component of this agricultural portfolio. There are 896 irrigation projects across the African sample. In more recent years these make up roughly five percent of agricultural projects. Figure [10] shows this nesting directly, with agriculture



Notes: The figure illustrates the total amount of all aid projects, and agricultural aid projects. The solid line illustrates the share of irrigation aid projects in relation to agricultural projects

Figure 10: Overivew Aid Projects

as a share of all projects and irrigation as the narrow band within it. Irrigation aid is therefore a minor but well-defined part of the largest aid sector on the continent, which motivates the focused evaluation in the main text.

A.3. Harvest Yield Database

The Spatial Production Allocation Model (SPAM, also known as MapSPAM) is a globally consistent, spatially-disaggregated dataset that allocates crop production statistics from administrative units into gridded pixels ($\approx$ 10×10 km / 5-minute cells) using a constrained optimization (cross-entropy) approach. The methodology integrates various spatially explicit input data, including crop production statistics, land cover classifications, crop suitability assessments, population density, and expert knowledge. Beginning with administrative units containing available production statistics, SPAM refines estimates using high-resolution cropland maps and agro-ecological suitability models from FAO/IIASA. The model accounts for four production systems(irrigated high-input, rainfed high-input, rainfed low-input, and rainfed subsistence) given that suitability and attainable yields differ by system. It applies constraints based on agricultural land availability, irrigation extent, and suitability indices to allocate crop production optimally at a fine spatial scale. To improve accuracy, SPAM incorporates cropping intensity adjustments, reconciles outputs with FAO national statistics, and factors in economic incentives using international crop prices. The final outputs include estimates of physical area, harvested area, yield and production for 42 crops across different production

systems. These outputs are computed per pixel and then scaled such that national aggregates align to FAO statistical averages used for the SPAM version.

The physical area indicates the actual area where a crop is grown, thus, the crop distribution within a pixel. Therefore, the sum of pixels of a given crop for the physical area cannot exceed the pixel area. This variable can be used to detect where cropland exists, and to measure cropland expansion over time. However, it does not count how often production was harvested from it. To detect multiple cropping, the measure harvested area can be used. It is also measured in hectare, and the harvested area should be at least as large as the physical area. It can exceed the value of physical area if multiple cropping takes place of one crop on the same plot. To measure cropping-intensity, the harvest area can be used. Yield is a measure of productivity, mainly the amount of production per harvested area. It is measured in kilograms per hectare. Given that there are four different production types available for the analysis, the total amount of yield is not only the sum of all the production systems, but the weighted average of the four yields. Finally, the production measure indicates the production for each system and crop, and is calculated by multiplying the area harvested times yield.

Data on the different crops are available for the years 2000, 2005, 2010, 2017 (sub-Saharan Africa only), and 2020. Although data is available for the year 2000, the analysis starts in 2005. This is due to the fact that the SPAM methodology is consistent for the years starting from 2005. SPAM 2000 differs substantially in its input data and crop coverage. Only major crops are defined, leading to missing information, especially for the cash crop classification.

It should be noted that SPAM scales outputs to conform to national FAO averages used in that release. This means pixel values are model-based allocations, consistent with national statistics, not independent field measurements. Additionally, the accuracy of the data highly depends on the quality of input statistics and cropland maps. Therefore, the obtained database currently provides the "best-available model estimates" but still can contain uncertainties.

A.4. Measurement Error in MapSPAM Outcomes

Because MapSPAM outcomes are model-based allocations rather than direct observations, measurement error is unavoidable. Three arguments imply that this error is

classical with respect to the treatment and therefore attenuates rather than biases the estimates.

First, the treatment plays no role in the allocation model. The cross-entropy procedure allocates administrative crop statistics to pixels using agro-ecological suitability, satellite-derived cropland extent, population density, market access, irrigation extent, and prior crop presence. The location of aid projects is not among these inputs: geocoded aid data enter the MapSPAM pipeline at no stage. For the measurement error to be correlated with treatment, the error-generating process would need to respond to the location of aid projects, which it structurally cannot. Moreover, the irrigation-extent layer used in the allocation is anchored to its 2005 vintage and held essentially constant across releases; since 2005 is the baseline year and all projects in my sample start thereafter, even this input predates treatment.

Second, the panel design differences out the level of the error. Even if the downscaling persistently over- or under-allocates production to a particular cell, the grid-cell fixed effects absorb any time-invariant component of the error. A potential bias in a persistently misallocated cell is therefore accounted for. What remains is the time-varying, idiosyncratic component, thus a classical error.

Third, the border design differences out the regional error structure. MapSPAM allocates within administrative units and benchmarks pixel values to national FAO totals, so the systematic component of the error is organized at the country and administrative-unit level. Treated and control cells within a border segment draw on the same satellite, suitability, and market-access layers, and the segment-by-year fixed effects absorb allocation shocks common to a segment in a given wave. Country- or region-level allocation bias therefore hits treated and control cells symmetrically and nets out in the difference. The design is built to remove geography- and boundary-correlated confounds, and the error structure of MapSPAM is correlated with exactly these dimensions.

A.5. Countries in Sample

Table 12: Sample Composition by Country

Country	**Grid Cells**	**Treated Cells**	**Border Segments**
Zambia	478	130	7
Benin	324	80	6
Niger	291	72	5
Uganda	181	68	2
Ghana	134	40	2
Ethiopia	357	33	5
Mozambique	122	30	3
Somalia	191	29	3
Zimbabwe	702	29	7
Nigeria	838	21	6
Liberia	44	19	1
Cameroon	100	18	4
Kenya	203	17	4
United Republic of Tanzania	350	17	4
Eritrea	136	14	2
South Sudan	239	13	4
Senegal	48	12	1
Sudan	261	8	5
Burkina Faso	117	6	3
Guinea-Bissau	27	6	2
Namibia	26	4	1
Mali	284	1	3
Democratic Republic of the Congo	388	0	4
Central African Republic	215	0	2
Botswana	238	0	3
South Africa	214	0	2
Côte d'Ivoire	361	0	4
Rwanda	21	0	1
Chad	604	0	3
Togo	214	0	3
Guinea	167	0	2
Eswatini	39	0	1
Angola	118	0	1
Total	**8,432**	**625**	**101**

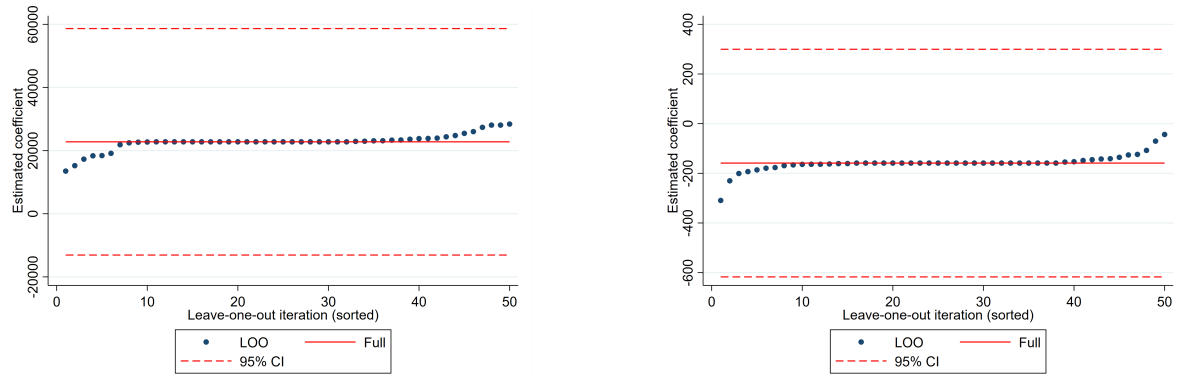
Notes: This table reports the geographical distribution of the sample across the 33 countries. It displays the total number of grid cells, the subset of treated cells containing irrigation development aid, and the active border segments per country.

A.6. Results using Staple Crop Production Value and Cash Crop Yields

Table 13: Robustness Checks

	Cash Yield (kg/ha)			Staple Crop Prod. Value (Index)		
	(1) Placebo	(2) Prec 1	(3) 10Y	(4) Placebo	(5) Prec 1	(6) 10Y
L5.Treated (25 km)	-123	-143	-159	42	15,721	22,786
	(143)	(266)	(234)	(55)	(17,767)	(18,299)
Observations	26,404	25,116	27,236	26,344	25,084	27,204
R^2	0.61	0.71	0.72	0.56	0.70	0.70
Mean DV	.	1,180	1,152	.	128,967	124,135
Grid Cell FE	✓	✓	✓	✓	✓	✓
Year × Border FE	✓	✓	✓	✓	✓	✓

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) and Column (4) report placebo regressions using border segments that never receive irrigation treatment and are located between 150 and 400 kilometres from the nearest irrigation project. Column (2) and Column (5) restrict the sample to border segments where all treated grid cells originate from Precision 1 geocoded aid projects.Column (3) and Column (6) restrict the sample to irrigation projects that have been active for approximately ten years. Standard errors, clustered at the border segment, are reported in parentheses.



(a) Staple Production Value (Index) (b) Cash Crop Yield

Notes: The figure reports the baseline specifications by excluding one border segment at a time. The coefficients are sorted by magnitude against the full-sample estimate and its 95% confidence interval.

Figure 11: Coefficient stability: Leaving-one-border-segment-out

Table 14: Robustness Checks

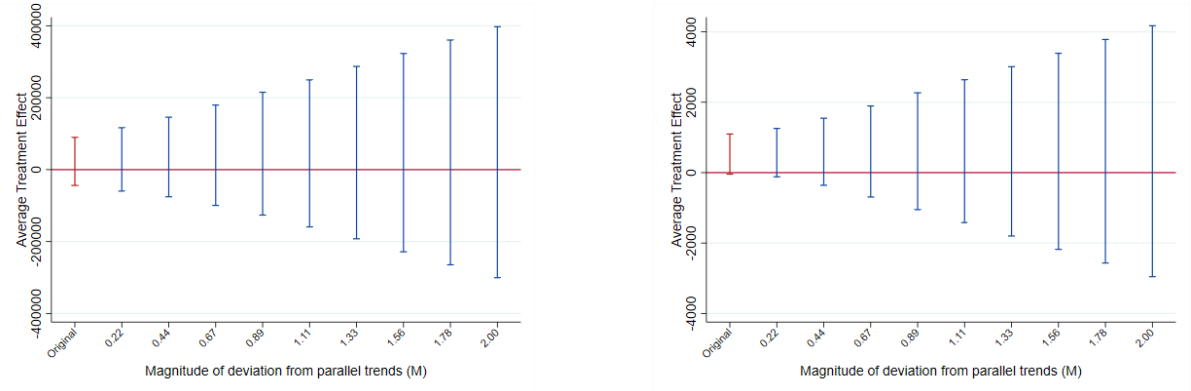
	Yield (kg/ha)		Prod. Value (Index)	
	(1) Staple	(2) Cash	(3) Staple	(4) Cash
L5.Treated (25 km)	417	37	1,634	13,974
	(677)	(164)	(1,151)	(21,025)
Observations	27,148	27,148	27,120	27,120
R^2	0.84	0.76	0.71	0.75
Mean DV	7,410	1,149	6,148	122,858
Grid Cell FE	✓	✓	✓	✓
Year × Border FE	✓	✓	✓	✓
Year × Country FE	✓	✓	✓	✓

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table includes adds additional fixed effects. All specifications contain next to grid cells and border-segment × year fixed effects, country × year fixed effects. Standard errors are reported in parentheses.

Table 15: Sensitivity to Border Bandwidth

	Distance to nearest border				
	≤70 km	≤80 km	≤90 km	≤100 km	≤110 km
	(1)	(2)	(3)	(4)	(5)
Panel A: Cash crop yield (kg/ha)					
L5.Treated (25 km)	-212	-186	-166	-159	-158
	(254)	(246)	(238)	(234)	(233)
Observations	20,876	23,404	25,624	27,236	27,656
R^2	0.72	0.72	0.72	0.72	0.72
Mean DV	1,178	1,169	1,160	1,152	1,149
Panel B: Staple crop production value (index)					
L5.Treated (25 km)	29,477	26,779	24,071	22,786	22,897
	(19,234)	(18,471)	(18,408)	(18,299)	(18,237)
Observations	20,864	23,392	25,604	27,204	27,624
R^2	0.70	0.70	0.70	0.70	0.70
Mean DV	123,451	123,155	123,908	124,135	124,499

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification on the sample of grid cells within the indicated distance of the national border. The baseline sample corresponds to column (4), ≤100 km. All specifications include grid-cell and border-segment × year fixed effects and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.



(a) Staple Production Value (b) Cash Crop Yields (kg/ha)

Notes: The figure reports sensitivity bounds based on the HonestDiD procedure of Rambachan & Roth (2023). The parameter M governs the magnitude of post-treatment deviations from parallel trends that are allowed relative to the largest deviation observed in the pre-treatment period. M=0 corresponds to exact parallel trends, while larger values permit progressively larger violations. The plotted intervals represent 95% robust confidence intervals for the treatment effect under the corresponding restriction on trend deviations. If an interval excludes zero, the estimated treatment effect remains statistically distinguishable from zero under that degree of deviation from parallel trends.

Figure 12: Honest DiD

Table 16: Sensitivity to Treatment Radius

	Distance to aid project					
	5 km	15 km	25 km	35 km	45 km	55 km
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Cash crop yield (kg/ha)						
L5.Treated	-524	-153	-159	-128	-166	-187
	(453)	(263)	(234)	(205)	(195)	(202)
Observations	12,844	21,916	27,236	29,988	32,412	34,412
R^2	0.75	0.75	0.72	0.75	0.75	0.75
Mean DV	1,317	1,214	1,152	1,255	1,218	1,210
Panel B: Staple crop production value (index)						
L5.Treated	42,946	34,795*	22,786	17,370	21,522	23,740
	(31,315)	(19,124)	(18,299)	(18,490)	(20,051)	(21,776)
Observations	12,828	21,884	27,204	29,956	32,380	34,380
R^2	0.68	0.70	0.70	0.70	0.70	0.71
Mean DV	139,279	123,984	124,135	130,355	129,205	131,163

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification with treatment being active for grid cells within the indicated distance of an irrigation aid project, lagged five years. The baseline corresponds to column (3), 25 kilometers. Rings are cumulative, so larger radii nest smaller ones. All specifications include grid-cell and border-segment × year fixed effects, include grid cells with a maximum distance of 100 kilometers to the border, and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.

Table 17: Sensitivity to Treatment Lag

	Years since project commitment (minimum project age)						
	L1	L2	L3	L4	L5	L6	L7
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Cash crop yield (kg/ha)							
Treated (25 km)	23	-41	-64	-119	-159	-302	-162
	(184)	(178)	(198)	(217)	(234)	(195)	(156)
Observations	27,236	27,236	27,236	27,236	27,236	27,236	27,236
R^2	0.72	0.72	0.72	0.72	0.72	0.72	0.72
Mean DV	1,152	1,153	1,151	1,152	1,152	1,153	1,156
Panel B: Staple crop production value (index)							
Treated (25 km)	29,640	11,068	15,242	24,013	22,786	36,802	59,210*
	(24,153)	(19,800)	(21,781)	(22,065)	(18,299)	(25,169)	(30,534)
Observations	27,204	27,204	27,204	27,204	27,204	27,204	27,204
R^2	0.70	0.70	0.70	0.70	0.70	0.70	0.70
Mean DV	124,610	124,188	124,094	123,981	124,135	124,293	124,385

Notes: Significance levels are defined as * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column re-estimates the baseline specification varying the time lag (the time an irrigation site needs to start operating). The baseline corresponds to column (5), five years. All specifications include grid-cell and border-segment × year fixed effects, include grid cells with a maximum distance of 100 kilometers to the border, and apply CEM × border-distance weights. Standard errors, clustered at the border-segment level, in parentheses.