Irrigation Infrastructure and Satellite-Measured Land Use Impacts: Evidence from the Senegal River Valley

Abdoulaye Cisse   Alain de Janvry   Joel Ferguson
Marco Gonzalez-Navarro   Samba Mbaye   Elisabeth Sadoulet
Mame Mor Anta Syll*

August 25 2024

Abstract

Increasing the extent of irrigation in sub-Saharan Africa is a potential approach to closing yield gaps and improving resilience to climate change. Using over 3,000 satellite images spanning more than 30 years, we show that development of irrigation infrastructure in the Senegal River Valley has led to a large increase in average cultivation rates: a sixfold increase from 4 to 24 percentage points in the non-rainy season and a tripling from 10 to 30 percentage points cultivated in the rainy season. In spite of this substantial increase in cultivation rates, we find widespread heterogeneity across projects, with 1/4 of total irrigated area remaining unused as of 2019. We provide farmer survey evidence that limited access to water remains a major constraint on production and that removing this constraint requires action beyond individual farmers.

JEL Codes: O13, O18, Q15, H54.

---

*Cisse: University of California, Berkeley, cisse@berkeley.edu, de Janvry: University of California, Berkeley, alain@berkeley.edu, Ferguson: Stanford University, joelferg@stanford.edu, Gonzalez-Navarro: University of California, Berkeley, marcog@berkeley.edu, Mbaye: Université Gaston Berger, Saint-Louis, Senegal, sambambayeugbcrdes@gmail.com, Sadoulet: University of California, Berkeley, esadoulet@berkeley.edu, Syll: Université Gaston Berger, Saint-Louis, Senegal, mamemorantasyll@gmail.com. The authors gratefully acknowledge financial support from the BASIS Feed the Future Innovation Lab for Markets, Risk and Resilience through the United States Agency for International Development (USAID) grant number A20-1825-S003. IRB approval was obtained from North Star Review Board under protocol NB300149. The contents are the responsibility of the authors and do not necessarily reflect the views of USAID or the US Government. We thank SAED for sharing their administrative data on irrigation projects and Khaly Fall and Ousmane Samaké from SAED for insightful information and discussions on the results. Authors are listed in alphabetical order.

1 Introduction

Food security in sub-Saharan Africa largely depends on an agricultural sector that lags behind the rest of the world in terms of level and stability of yields. Raising agricultural output is seen as a top priority in the face of challenges imposed by increasing demand due to rapid population growth — projected to reach 3.3 billion by the end of the century (UN, 2024). Climate change, too, threatens agricultural productivity and food security in sub-Saharan Africa (Lobell et al., 2008). In this context, one of the important strategies proposed to cope with food deficits is to increasingly shift from rainfed and flood recession agriculture to irrigated agriculture (Rosa et al., 2020).

In this paper, we document the heterogeneous consequences of the damming of the Senegal River in the Senegal River Valley (SRV) in the 1980s on expansion of cultivated land in Senegal using a long time-series of satellite imagery. Senegal provides a context that is uniquely amenable to analysis. First, the SRV is the site of one of the few existing large irrigation schemes in sub-Saharan Africa, though more such projects have been proposed recently (Higginbottom et al., 2021). Second, the Sahel, through which the Senegal River runs, is a region that has grappled with desertification since the 1970s (Hein and De Ridder, 2006), and so may provide a useful example for studying the role of irrigation in climate change adaptation. The limited background vegetation in the Sahel also makes detection of cultivated land using satellite imagery particularly easy. Finally, development of the valley within Senegal is administered by a single government agency, which has consistently recorded data on irrigation infrastructure development since at least the 1970s.

In spite of its policy and academic interest, evaluating the role of irrigation infrastructure in most developing countries has proven challenging due to a lack of time series data on land use and agricultural output. Indeed, there are no high-resolution agricultural output or land use data for the SRV spanning the decades over which irrigation development has occurred. To fill this data gap, we turn to a collection of over 3,000 publicly available Landsat images collected between 1985 and 2019. To convert the imagery into useful data for studying agricultural outcomes, we calculate the normalized difference vegetation index (NDVI) of each 30m×30m pixel. A large body of remote sensing research documents the usefulness of NDVI in detecting biomass, with greater values indicative of denser and healthier plants. We take the maximum NDVI level over either the rainy season or the non-rainy (also called 'off-season'), and apply a simple threshold rule to infer cultivation. We validate this procedure by documenting that, in areas lacking irrigation, NDVI very rarely reaches above our chosen threshold.

We combine this newly developed cultivation dataset with administrative data provided by the Senegalese government on 871 distinct irrigation infrastructure projects completed between 1988 and 2019. These data record the geographic extent of projects and of subsections of projects utilized by different producer organizations or individual producers called UMVs (*Unités de Mise en Valeur*), the year of project completion, and some characteristics of projects such as whether they were publicly or privately developed, and the number of members by gender. By combining these datasets, we produce a UMV-level panel of cultivation rates by season spanning over three decades and covering 1,364 distinct UMVs.

Because we observe cultivation activity before and after irrigation project completion, we can estimate the effect of such projects using an event study design. There are two methodological challenges that must be accounted for. First, because irrigation projects are completed in different years, treatment timing is staggered, leading to potential issues with the traditional two-way fixed effects estimator in the presence of unrestricted treatment effect heterogeneity (Sun and Abraham, 2021). Second, the earlier half of the study period has a relatively large share of missing data due to lack of high-quality imagery, making our dataset an unbalanced panel.

We overcome both of these issues using the imputation event study estimator developed by Borusyak, Jaravel and Spiess (2024), which provides us with consistent estimates of a weighted average treatment effect on the treated under general treatment effect heterogeneity in panels which may be unbalanced. We posit a simple two-way fixed effect specification of agricultural land utilization in the absence of irrigation infrastructure and show that when this model is augmented with pre-event indicators there is no evidence of pre-project trends, supporting the parallel trends assumption necessary for causal identification. Once the model of utilization in the absence of irrigation is estimated, treatment effect estimates can be calculated by predicting untreated potential outcomes, subtracting them from observed treated potential outcomes, and averaging these differences across observations.

Our results provide strong evidence that cultivation rates increase substantially beginning in the first year after irrigation project completion, especially during the dry season. In the longer term, we find that increased cultivation rates are remarkably stable at around 20 percentage points above pre-irrigation levels for the first 20 years, and trend even higher from years 20 to 25. Considering that average cultivation rates prior to project completion were only 10 percent in the rainy season and 4 percent in the off-season across the sample, these are large and significant effects.

While our main results provide a positive picture of land cultivation patterns following introduction of irrigation infrastructure, this large estimated average impact masks significant heterogeneity across UMVs and over time. Irrigation has generally had an increasing impact on cultivation over time, with a sharp jump in year 2000, a year of major policy changes toward a greater role for the private sector, and project characteristics such as year of completion and sectoral origin of development affecting treatment effects.

We proceed to classify UMVs according to their cultivation performance post project completion. Exploiting our spatially and temporally granular dataset, for each UMV we calculate the average share of pixels that change cultivation status (from zero to one and vice versa) between year t and t-1. We classify UMVs as having "intermittent land use" if we observe cultivation rates that vary substantially across years (more than 10 percent annual average status change). The remaining UMVs are classified into "Low Use" (with an average cultivation rate below 33 percent), "Medium Use" (cultivation rate between 33 and 67 percent), and "High Use" (cultivation rate above 67 percent).

Using this classification, we observe that 18 percent of UMVs are Low Use and 36 percent Intermittent Use, so that more than half of the land opened to irrigation is underutilized in some form. Medium Use account for 31 percent, and only 14 percent of UMVs have High Use cultivation rates. This heterogeneity in cultivation rates implies existence of a large inefficiency in land use, as we do not find evidence that intermittent use is due to land rotation.

To explore the causes for such widespread under-utilization of irrigated land, we move from satellite data to our dedicated farmer survey. Our sample consists of farmers that were known in their community for having some land that was either intermittently used or continuously unused. We obtained information from farmers about the history of land use in their plots during the three off-seasons in the 2021-2023 period as well as the causes for non use.

We find that, in spite of all plots in our sample being in projects with completed irrigation infrastructure, lack of water availability was cited as the main constraint to land cultivation. This is the case for 61.6 percent of continuously unused land and 34.7 percent of intermittently used land across these three seasons. The second most important cause for lack of utilization is financial constraints, especially among intermittently used plots.

Further analysis shows that the water access problem causing non-use is not seen as something individual farmers can address on their own. In almost all cases, the solution is seen as requiring the intervention of actors external to the production unit, such as those insuring the maintenance of major canals that serve multiple UMVs.

This paper is closely related to the literature on irrigation development, land cultivation, and response to water availability. Duflo and Pande (2007) study the upstream and downstream impacts of dam construction, finding increased agricultural productivity and reduced sensitivity to weather shocks downstream, but increased poverty incidence upstream. Hornbeck and Keskin (2014) study agricultural production over the Ogallala Aquifer and nearby counties to explore the consequences of increased irrigation potential. They find that counties with a greater share of land over the aquifer increased the share of land devoted to water intensive crops over time, ultimately leading them to face similar weather-related risks as counties with less aquifer potential which specialized in more drought-tolerant crops. Blakeslee, Fishman and Srinivasan (2020) find that exogenous well failure in India induces households to transition out of agricultural production. We contribute to this literature by leveraging high-resolution data on project construction and land use to provide a detailed description of how agricultural production changes in response to irrigation availability both in the short and longer term.

Within this literature, the most closely related paper to ours is BenYishay et al. (2024) who study small scale irrigation program impacts on agricultural and nutritional outcomes in Mali. We utilize a similar event study design, using raw NDVI as an outcome in robustness checks, and obtain a 60% larger effect. One major difference between contexts is that in Mali irrigation extends and supplements water for agriculture in the rainy season, whereas in the Senegal River Valley there is a large effect on irrigation in the non-rainy season. In addition, our focus in this paper is on detecting land usage patterns and correlates of irrigation project success using high frequency observations afforded by satellite data, whereas BenYishay et al. (2024) are primarily focused on household level outcomes.

Another closely related paper is that of Asher et al. (2022) who study the long-run impacts of irrigation canals in India. As in our study, they combine information on the timing of irrigation infrastructure completion with a highly disaggregated dataset of outcomes, including cultivation rates. However, their analysis of agricultural outcomes focuses on a period long after irrigation canal completion, 30 years at the median. In contrast, we are able to dynamically trace out the effects of irrigation infrastructure over time and by year of completion. Put differently, Asher et al. (2022) study equilibrium outcomes, whereas we can study transition paths. Another notable distinction is the disaggregation level. Because we have cultivation measurements at the 30m×30m pixel-level, we can study within-project dynamics, such as rotations in crop cultivation.

Our use of satellite imagery builds on a large literature in remote sensing which seeks to map agricultural outcomes. Researchers have used a variety of vegetation indices as remote sensed proxies for plant biomass since the late 60s, with NDVI being the most widely used (Xue, Su et al., 2017). More recently, satellite imagery has been used to gain a better understanding of agriculture in low-resource settings where ground-truthed measurements may be sparse and of lower quality (Burke et al., 2021). Lobell et al. (2019) and Lobell et al. (2020) show that peak NDVI, the measure we use in this study to detect cultivation, is more strongly correlated with yields measured by full crop cuts than self-reported yields in the case of sorghum in Mali and maize in eastern Uganda. Because NDVI is consistently measured over time by sensors aboard Landsat satellites, we are able to map cultivation at a very high level of resolution for nearly the past four decades.

The approach we take to mapping cultivation is similar to recent work by Higginbottom, Adhikari and Foster (2023), who also study expansion of irrigation in the SRV using Landsat imagery. They train a random forest model to classify land as irrigated or not, using a dataset of images hand-labelled to outline areas that appear cultivated and find a large expansion of irrigated area, particularly since 2008. Our approach differs from theirs in two ways. First, we use administrative data on the location of irrigation projects, making the cultivation classification task much simpler as we do not need to discern agricultural land from other persistent vegetation such as stands of trees. Second, we apply our dataset to causally estimating the effect of irrigation project completion on land use, a step beyond the work of Higginbottom, Adhikari and Foster (2023) made possible by our access to administrative data on irrigation project completion.

More broadly, there is a growing body of economic research which uses remote sensing techniques to measure agricultural outcomes over large areas and long time scales in causal research designs. Wuepper et al. (2023) employ regression discontinuity at international borders around the world to study the effect of economic freedom, as measured by the Fraser Institute, on agricultural productivity using a global dataset of pixel-level observations of annual maximum enhanced vegetation index (EVI), a modified version of NDVI. The aforementioned study by Asher et al. (2022) also uses EVI as a proxy for agricultural productivity which is observable for every settlement in India. Murillo-Sandoval et al. (2021) develop a 31-year panel of land use based on Landsat imagery and using difference-in-differences find that conflict events in Colombia cause the conversion of forest to agriculture. Our study contributes to this literature by demonstrating how large collections of remotely sensed data can be combined with primary data collection to better support rigorous statistical analysis.

The rest of the paper is organized as follows. Section 2 presents the history and context for this paper. Section 3 presents our data. Section 4 introduces the econometric specification. In section 5, we present results on land use. Section 6 discusses constraints to land use. Concluding remarks are given in Section 7.

2 History and context of study

2.1 Context

The Senegal River Valley (SRV) is an agroecological region located in the northern and eastern part of Senegal, with the river itself being the boundary with Mauritania.[1] The region produces a significant share of the two major crops used for domestic consumption, rice and onion, and is more broadly specialized in cereals and horticulture.[2] The region receives minimal rainfall; most agriculture is consequently based on use of water from the river. Traditional agriculture took place in flooded areas, after water recession, benefiting from rich soils deposited by the river and residual humidity in the soil. Projects to control the water flow and transition towards an irrigation model were initiated as early as 1925, but only 6,000 ha were irrigated by the time of independence in 1960. Starting in the 1970s, and triggered by the famine of 1973, the Government of Senegal pivoted to investing in the development of irrigation agriculture.

A major turning point in this infrastructure development came with construction of two major dams that were completed in the 1980s: the Diama dam situated close to the river estuary in the lower valley, designed to stop saltwater intrusion into the river during the dry season, and the Manantali dam in the upper valley in Mali aimed at storing water from the Guinean affluents that could then be managed throughout the year. Prior to these dams, the Senegal River water flows were very irregular, both seasonally with huge flows in the rainy season producing flooding and a long dry season with little water, and across years with large variations in the extent and duration of flooding. Since the dams have been in operation, controlled floods have been introduced to allow traditional agriculture to take place while irrigation was progressively developed. Currently, about one third of agricultural land is in traditional flood agriculture and two thirds has access to irrigation canals.

A specialized agency, SAED (Société d'Aménagement et d'Exploitation des Terres du Delta du Fleuve Sénégal), has been the main government body charged with the development of the SRV. SAED oversees the construction of most canals, pumps, and drainage, and also their operation and maintenance. It divides its intervention territory into delegations. The delegations are broadly constructed to follow administrative units while also accounting for differences in hydraulic and agroecological characteristics. The SRV is thus divided into five delegations: Dagana, Podor,

[1] See Figure 1 for a map of Senegal.

[2] The SRV accounted for 29% of domestic rice production and 31% of domestic onion production in 2022 and 2018 respectively (Direction de l'Analyse, de la Prévision et des Statistiques agricoles, 2022; Agence Nationale de la Statistique et de la Démographie, 2018).

Matam, Bakel, and Lac de Guiers. Our study focuses on the two largest delegations, Dagana and Podor.[3]

Throughout this study, we use two terms to refer to units of agricultural land: "project" and "*Unité de Mise en Valeur*" (UMV). A project covers all cultivated plots that share the same hydraulic infrastructures (e.g. adductors, irrigation perimeters, or drains) or the same water source (e.g. a pumping station, or a water intake on a large canal). In our analysis, we also compare public projects developed and funded by the State and private projects developed with private funds.[4]

A UMV corresponds to the land operated by a producer organization or independent private actors within a project. In our analysis sample, projects contain 1.6 UMVs on average: the majority of projects only have one UMV and most have 1-4 UMVs, although there are a few outlier projects with as many as 152 UMVs. UMVs contain on average 44.33 members: most UMVs have many producers, though some have only a few listed producers. The division of projects among UMVs simply corresponds to an organizational and property rights division, and does not translate into a hydraulic division.[5]

2.2 History of Land Development

Prior to 1970, irrigation infrastructure development in the SRV was not a priority for Senegalese policymakers. This changed after the 1973 famine. Large tracts of irrigated land, ranging from 500 to 2000 hectares, were then allocated to small scale farmers grouped into village-level organizations. In the late 1970s and 1980s, medium-sized tracts also managed by villagers were introduced. During this period, SAED financed and built the irrigation schemes, divided land tracts into blocks with feeder canals, while groups of 15–20 farmers were made responsible for water

[3]See Table A1 for the crop calendar and production volumes of the main crops in these two delegations.

[4]The public projects include the *Grands Aménagements* (GA) and *Aménagements Intermédiaires* (AI) which are projects of more and less than 400 hectares, respectively, that were managed by SAED before their progressive transfer to farmers' organisations, and the *Périmètres Irrigués Villageois* (PIV) which are projects of 20 to 30 hectares managed by village groups since their creation. The private projects include small private projects, known as *Périmètres Irrigués Privés* (PIP), which are privately managed by independent producers, and large private projects, which are managed by large producers.

[5]Producer organization designates both an organization in the strict sense (e.g a group of farmers organized in a cooperative), or companies and private producers involved in the management, maintenance, and development of the project. In our context, most organizations are structured and declared as *Groupement d'Intérêt Économique* (GIE), which are economic interest groups with a legal status whose members collectively conduct certain activities such as financing.

distribution and feeder canal maintenance within each block. The day-to-day management of the irrigation schemes was village-based, with a local committee in charge of management and operation (Bruckmann, 2018; Le Roy, 2000). Over these first two decades, irrigated land is estimated to have grown from 6,000 to 18,000 ha.

Since the early-1990s, there was a notable increase in private, intermediate-sized, developments (PIPs). This period of allowing and promoting private development has been much more successful in terms of developed area: irrigated land more than tripled to 65,000 ha by 1994 (Faye, Fofana and Bélières, 1995). Starting in 2000, larger private developments emerged, and there has been a focus on rehabilitating older projects using assistance from international organizations such as Germany's KfW, the World Bank, Japan's JICA, the United States' MCC (Millennium Challenge Account), the French bilateral aid agency AfD, and the African Bank for Development (Bourgoin and Diop, 2023). Total irrigated area now stands at 128,000 ha, out of a total estimated irrigable potential of 240,000 ha (SAED and JICA, 2019; MAER, 2015; USAID, 2017). While we discuss measurement in detail in subsection 3.2, we note here that using the broadest definition of cultivated land we can obtain, our estimates of satellite-measured cultivated land broadly match the levels and trends of these official estimates, as shown in the cultivated land time series in Figure 2, in spite of this graph excluding the smallest delegations.

2.3 History of Agricultural Policies

Between 1980 and 2000, the overarching objective of Senegal agricultural policies was to reorient production from the traditional colonial exporting of groundnuts (peanuts) towards an import substitution strategy centered around the domestic production of cereals that could increase national food security and reduce the trade deficit. Multiple programs were enacted, and the development of irrigation infrastructure in the SRV was a major component. This strategy was accompanied by strong involvement of the government in farming activities through initiatives such as the *Plan de Redressement Économique et Financier* (PREF), the *Nouvelle Politique Agricole* (NPA), and the *Programme d'Ajustement du Secteur Agricole* (PASA). This strategy, while effective at expanding land under irrigation, eventually ran into the limits of state-owned rural management structures that were poorly run (Ministère du Développement Rural, 1984, République du Sénégal, 1995).

Following the change in political regime in March 2000, Senegal embarked on a wave of policy reforms focused on boosting agriculture productivity starting with the *Loi d'Orientation Agro Sylvo Pastorale* (LOASP). One key change consisted of new policies to incentivize private investment in rural areas through easier land access for national private producers (including agribusinesses) and an expansion of subsidies for agricultural inputs (Seck, 2016).

The post-2000 period was also marked by a shift in attention from pure expansion of the irrigation network to prioritization of rehabilitation work and improvement of older canals. These include the Great Offensive for Agriculture, Food, and Abundance (GOANA), and the Program for Accelerated Agricultural Growth (PRACAS). GOANA earmarked 35,000 hectares in the Senegal River Valley for hydro-agricultural scheme rehabilitation, provided financial assistance for pumps, and other essential irrigation equipment (IPAR, 2015). PRACAS continued the focus on rice and in particular the rehabilitation of hydro-agricultural schemes to improve water management.

In addition, this period saw the emergence of improved access to credit through private sector financing. As the agricultural sector became more efficient through the development of value chains, the volume of agricultural credit in SRV increased compared to the 1990s, and with it the amount of cultivated area (Le Roy, 2011).

These policies generated immediate impacts: Senegal's paddy production increased from 124 million tons in 1999 to 240 million tons in 2000 with a doubling of area planted (USDA, 2000).

Overall, agriculture in the new century in the SRV has been characterized by more flexible access to land by independent farmers, improved financial conditions, and a more consistent focus by government on rehabilitation of water canals.

3 Data

Our analysis relies on three data sets. The first is administrative data from SAED on UMVs in the delegations of Dagana and Podor. These data inform the geographic boundaries of UMVs as well as the characteristics of associated projects and producer organizations. The second is Landsat satellite imagery spanning over three decades that we use to obtain granular high frequency measures of land cultivation by season. The third is a phone survey of farmers in irrigated UMVs in Podor that community leaders determined to have low or intermittent cultivation. We describe each of these three data sources in what follows.

3.1 SAED Administrative Data

The sample for our analysis of the effects of irrigation infrastructure development on agricultural outcomes is built using a collection of data sets provided by SAED detailing irrigation projects as of 2020. The first data set of this collection is a shapefile delimiting the boundaries of UMVs and the associated project, UMV, and producer organization identifiers.

The project and UMV identifiers contained in the shapefile are used to merge in characteristics of UMVs from three separate files: a project characteristics file, a producer organization file, and a UMV characteristics file. The project characteristics file contains information on year of completion of each irrigation project, whether it was publicly or privately developed, and which sector (SAED's sub-delegation administrative unit) it belongs to. The producer organization file contains basic information such as the number of members and the type of organization (e.g. a women's organization, a village organization, etc.). The UMV characteristics file contains self-reported information on the amount of land within the UMV that has been abandoned or added since creation of the project.

We limit the analysis sample to projects constructed in 1988 at the earliest to allow for assessment of pre-trends, given that our satellite imagery data begin in 1985. UMVs are excluded from the analysis sample if any of the following are true: (i) they cannot be georeferenced via the shapefile (1513 UMVs), (ii) their construction date is not available in the project characteristics dataset (110 UMVs), or (iii) they are associated with a project that was constructed before 1988 (945).

In total, our sample consists of 1,364 UMVs from 860 different irrigation projects. The UMVs in our sample cover 46,055 hectares, about 51% of the irrigated area recorded in the shapefiles. The location of the UMVs in space and their associated project construction dates are shown in Figure 1.

788 of our sample UMVs are in Dagana, corresponding to 459 unique projects and covering a total of 33,112 hectares. Similarly, our analysis sample in Podor contains 576 UMVs distributed across 401 projects over 12,943 hectares. In general, projects in Podor are much smaller than in Dagana. The average project area in Podor is 32.3 hectares while it is 72.2 in Dagana. This is partly due to large agribusinesses being mostly concentrated in Dagana while there are very few in Podor.

In order to better understand whether the analysis sample is representative of the broader UMVs in the Podor and Dagana delegations, Table A2 compares the means of observable characteristics between UMVs that are included and excluded from the analysis. Table A2 reveals that the UMVs in our sample are somewhat more likely to be private and are relatively newer compared to excluded UMVs. The latter is partly mechanical, due to our removal of projects completed before 1988. In terms of number of members and share of female members, both samples appear similar.

Using data from 2021, the UMVs in our sample have a similar average size compared to excluded UMVs, though their initial size was slightly larger. The analysis sample also has less self-reported abandoned area and consequently slightly greater currently exploited area compared to excluded UMVs.

Importantly, while Table A2 shows that UMVs in our analysis sample differ along a number of observable dimensions from those that are excluded, we also show that our analysis sample is quite similar to the set of all irrigated land in the SRV in terms of outcomes. Specifically, we construct the 1985-2019 NDVI series for UMVs in our analysis and for a set of hand-drawn irrigation project polygons compiled by Zwart (2017). Appendix Figure A1 shows that, while there is a slight divergence in the cumulative density functions of NDVI by sample, they are extremely close over much of the range of observed NDVI values, suggesting that our sample is broadly representative in terms of cultivation levels.

3.2 Satellite Imagery

We construct cultivation measures using satellite imagery spanning the 1985-2019 period. In particular, we use Tier 1 Thematic Mapper, Enhanced Thematic Mapper+, and Operational Land Imager imagery from Landsats 5, 7, and 8, respectively. These sensors record multispectral imagery data for 30m$\times$30m pixels roughly every two weeks. Landsat 5 was operational from March 1984 to May 2012, Landsat 7 from January 1999 to April 2022, and Landsat 8 launched in April 2013 remains operational. For the event study analysis, we restrict the temporal coverage of the study to 1985-2019, as the imputation event study estimator we use requires a set of projects that have not yet been completed for identification and all observed projects are completed by 2020. There is relatively worse temporal coverage prior to the year 1999, as only Landsat 5 was operational and produced relatively few science-quality images (see Appendix Figure A2 for monthly imagery availability over the study area). We apply a cloud mask based on the Landsat CFmask and apply the relevant scaling factors to radiance values before any additional processing.

Our analysis is based on a measure derived from multispectral imagery, the Normalized Difference Vegetation Index (NDVI), which has been widely used to study agricultural production due to its strong correlation with biomass. NDVI exploits the fact that chlorophyll reflects electromagnetic radiation with wavelength greater than approximately 0.69 μm, but absorbs visible light (Myneni et al., 1995). This causes plants to have high reflectance values in near-infrared bands, but low values in the adjacent red band.[6]

NDVI is calculated as

[6]The near infrared and red bands correspond to 0.77-0.9 μm and 0.63-0.69 μm respectively for the Landsat 5 Thematic Mapper.

$$NDVI = \frac{Near\ Infrared - Red}{Near\ Infrared + Red}$$

with greater values indicative of greater biomass. While the potential range is from -1 to 1, values less than 0 are rarely observed in our sample.

We calculate NDVI in every available image of the study region between 1985 and 2019 and take the pixel-level maximum, either for each of the three agricultural seasons in a given year (rainy, cold off-season, hot off-season) or combining the two off-seasons (cold off-season and hot off-season). To measure agricultural cultivation, we consider December-March as the cold off-season, April-July as the hot off-season, and August-November as the rainy season.

To get a better sense of how NDVI captures changes in agricultural production, we collect the red, green, and blue (RGB) bands of images over an area with varying NDVI for the 2020 hot off-season, and create a composite image with RGB values of a given pixel associated with the date of maximum NDVI for that pixel. In a sense, this composite shows the area as it would look if every pixel attained its maximum NDVI at the same time. Figure 3 shows that the UMV outlined in bright green, which has a highest average NDVI of the UMVs that appear in the image at 0.7, looks dark green in the RGB image, consistent with successful cultivation. In contrast, the UMV outlined in pink, which has the lowest NDVI in the image at 0.21, looks brown, consistent with null or unsuccessful cultivation. This 'smell test' suggests our use of maximum NDVI works as intended at capturing cultivation patterns from satellite images.

We next deal with the definition of a cutoff for what we classify as a cultivated plot. We propose a measure for whether a pixel is cultivated based on a simple threshold rule. If a pixel achieves a max NDVI in a given season greater than 0.3, we classify that pixel as cultivated. Our use of the max NDVI value is adequate given that development of biomass over a season may differ between locations due to management decisions, weather, and crop choice. Furthermore, given the sparsity of imagery in the early years of our sample, an ideal classification method would work well with imagery captured at any point during the season. Using a relatively low threshold that is greater than the maximum level generated by background vegetation allows for discrimination of cultivation prior to crop maturity.

The threshold of 0.3 is selected based on the fact that there are very few (less than 5%) pixels in uncompleted projects that exceed this value during the off-season. Appendix Figure A3 shows that, for the majority of years we study, the 95th percentile of NDVI values in projects which have not been completed is less than 0.3, and it is less than 0.4 in all but four years, suggesting that NDVI rarely reaches above these levels in the absence of irrigation. We also conduct a series of robustness checks using alternative thresholds and raw NDVI demonstrating that results are robust to threshold choice. UMV-level cultivation rates are calculated by taking the simple average of this binary pixel-level use status.

3.3 Phone Survey

To elucidate the reasons for land non-use, we conducted a phone survey with farmers in the delegation of Podor in September 2023. We started our sampling frame with the satellite-derived land use presented in Section 3.2. First, we pre-selected 30 UMVs in the delegation of Podor that exhibited multiple years of non-use, based on the estimated land use methodology presented in Section 3.2.[7] Next, we worked with SAED to obtain contact information for representatives of 20 of those UMVs.

For each of the 20 projects selected, we visited a representative of the UMV, and interviewed them using a general in-person survey.[8] At the end of these general in-person surveys, we asked each respondent to give us the contact information of 10-15 farmers in their UMV *who have land they do not use, either continuously or intermittently.* In the end, we received the contacts of 206 producers and successfully conducted interviews with 162 of them, obtaining information on 268 different plots.

The phone survey asks about the history of land use in plots operated by the farmer during the three off-seasons of 2021-2023. The survey asks about reasons for lack of use, separately for land that is never used in the three off-seasons (continuous non-use) and for land used some of the time (intermittent non-use). Potential responses include water availability, financial constraints, labor availability, and land quality, among others. We also asked respondents whether the main issue constraining use is a problem at the level of the farmer, the UMV, or the project, and whether the issue could be potentially solved by the farmer, within the UMV, or would require external intervention beyond the producer organization.

We end the survey with a series of preference elicitation questions about constraint-alleviating services that the respondent would be interested in receiving. In line with our constraints of focus, we consider the following six service options: (i) put you in touch with a financial institution, (ii) connect you to someone who is interested in renting out land, (iii) connect you to someone who is interested in

[7]In making this pre-selection, we restricted the sample to UMVs which are within 5 kilometers from the main highway road in order to minimize travel time to the chosen sites.

[8]This in-person survey is different in content from the phone survey and was mainly fielded to verify the validity of the satellite image land use methodology discussed in Section 3.2.

renting in land, (iv) provide you with information on how to improve land quality, (v) put you in touch with someone that can provide reliable labor for your farm, or (vi) coordinate a meeting with your neighbors to facilitate improvements to the drainage and irrigation system. We present farmers with a random selection of 3 of the 15 possible pairwise combinations of these improvement options and ask them to choose which one of the two they would prefer. These light-touch options are all costless to the farmer and do not imply any subsidy. The intent was to glean information about which type of problems are perceived by the farmers as solvable.

4 Econometric Specification

We organize our data as a panel of UMVs, each of which has a potentially different irrigation treatment date (i.e. staggered roll-out). Given this data structure, we implement an event study specification that ensures our estimates have a plausibly causal interpretation. Specifically, we implement the imputation estimator developed by Borusyak, Jaravel and Spiess (2024). To do this, we first posit a model of potential outcomes in the absence of irrigation:

$$y_{u,t}(0) = \alpha_u + \delta_{t,s(u)} + \varepsilon_{u,t} \tag{1}$$

where u indexes UMVs, $s(u)$ maps UMVs to administrative sectors, and t indexes years. This model can be augmented to allow for estimation of pre-trends. Letting $E_{p(u)}$ denote the year the project associated with UMV u is constructed and $pre = \{-9, -8, ..., -1\} \setminus \{-3\}$, this specification can be written as:

$$y_{u,t}(0) = \alpha_u + \delta_{t,s(u)} + \sum_{k \in pre} \beta^k 1\{t - E_{p(u)} = k\} + \beta^{-10} 1\{t - E_{p(u)} \leq -10\} + \varepsilon_{u,t}. \tag{2}$$

The indicator for three years prior to project construction is omitted because it is the first that all observations share.[9] Both equations 1 and 2 are estimated using only observations of UMVs in years prior to project construction. Once estimated, the coefficients and associated standard errors (clustered at the project level) of equation 2 allow for an evaluation of pre-trends to partially assess the validity of the

[9]Under the proposed model of untreated potential outcomes, the choice of which event time indicator(s) to omit is asymptotically irrelevant, as all estimated pre-treatment coefficients will converge to 0. Omitting an event time indicator earlier than -1, the event time which is typically omitted, allows for more straightforward visual inspection of pre-trends in the periods immediately prior to treatment, when anticipation effects are most probable.

parallel trends assumption underlying causal identification. The model of untreated potential outcomes implies that $\beta^k = 0$ for all values of $k \in (\{-10\} \cup pre)$, which can be formally evaluated via an F-test.

Next, we use the estimated parameters of equation 1 to impute untreated potential outcomes for post-treatment observations, $\hat{y}_{u,t}(0) = \hat{\alpha}_u + \hat{\delta}_{t,s(u)}$. Post-treatment observation-specific treatment effects can then be simply estimated as $\hat{\tau}_{u,t} = y_{u,t}(1) - \hat{y}_{u,t}(0)$. While these UMV specific estimates are potentially interesting in their own right, performing statistical inference on them is impossible without some restriction of treatment effect heterogeneity and the estimator underlying their computation does not necessarily exhibit desirable properties, such as consistency.[10] The estimates are unbiased however, and given a set of regularity conditions described by Borusyak, Jaravel and Spiess (2024) averages of the observation-specific treatment effects can be consistently estimated. For this reason, our estimands of interest are averages of observation-specific treatment effects, which can be obtained by assigning weights $w_{u,t}$ to each post-treatment observation and taking the weighted average

$$\hat{\tau}(\{w_{u,t}\}) = \frac{\sum_{u,t} w_{u,t}\hat{\tau}_{u,t}}{\sum_{u,t} w_{u,t}} \tag{3}$$

where $w_{u,t} = 0$ for pre-treatment observations. To estimate typical event study effects, we set $w_{u,t} = 1\{t - E_{p(u)} = k\}$ for various values of $k \geq 0$. Borusyak, Jaravel and Spiess (2024) also provide a method for conservative estimation of clustered standard errors of $\hat{\tau}(\{w_{u,t}\})$. We extend their statistical inference procedure to allow for clustering on irrigation projects, rather than UMVs, given that this is the level at which treatment is assigned.

Aside from avoiding potential biases due to negative weighting that can arise when using the more traditional two-way fixed effects estimator in empirical settings such as ours with staggered treatment timing and (potentially) heterogeneous treatment effects, the imputation estimator confers advantages over other recently-proposed estimators. In particular, our panel is not complete, with observations much more likely to be missing earlier in the panel, rather than later. These missing data pose no identification issue for the imputation estimator, so long as we are able to estimate $\hat{\alpha}_u$ and $\hat{\delta}_{t,s(u)}$. This is because once $\hat{\alpha}_u$ and $\hat{\delta}_{t,s(u)}$ are estimated, we can calculate $\hat{\tau}_{u,t}$ for any UMV-year for which we observe the treated outcome. Under the proposed model of untreated potential outcomes, additional pre-treatment observations only improve the precision with which $\hat{\alpha}_u$ and $\hat{\delta}_{t,s(u)}$ are estimated, and thus the precision of $\hat{\tau}(\{w_{u,t}\})$.

[10] This is because in short panel settings, the unit fixed effects α_u are not consistently estimated.

5 Impact of Irrigation Project Completion on Land Use

5.1 Aggregate Effect

We begin by presenting preliminary evidence of the impact of irrigation project completion in the raw data. Figure 4 displays monthly averages of NDVI in 2015, a year for which we have an abundance of images that allow this type of month-level analysis, for UMVs in projects completed in 2014 or earlier compared to those completed in 2016 or later. In the figure, the orange shaded box corresponds to the hot off-season, the purple shaded box to the rainy season, and the white area to the cold off-season. The dashed orange line is average NDVI in completed projects, while the solid green line displays average NDVI among projects that have yet to be completed.

Three facts emerge from this raw data figure. First, it is striking that measured NDVI is uniformly higher in completed projects throughout the year, including in the rainy season. Second, the figure also shows that the gap in NDVI is largest in the hot off-season, when NDVI is almost double in complete vs incomplete projects. Lastly, the pattern over the course of the year suggests that, for pre-completion projects, agricultural production is concentrated in the rainy season, whereas for post completion observations NDVI is at high levels in both the hot off-season and the rainy season, consistent with irrigation allowing for two agricultural cycles in a year. The pattern holds up when we analyze the NDVI separately by delegation in Appendix Figures A4 and A5, though the difference is slightly weaker in Podor.

Our main results for the impact of irrigation project completion on land cultivation rates are presented in Figure 5. The figure displays coefficient estimates and 95% confidence intervals of the event study estimation procedure described in section 4 with the NDVI threshold-based measure of UMV cultivation rate as the outcome. Recall that the outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI over the season (either April-July for the rainy season, or August-March for the off-season) is greater than 0.30, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level to obtain the share of project land in use. The X-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level.

First, we highlight that the figure provides compelling evidence that there are no pre-trends in cultivation rates based on the pattern of coefficients that are tightly centered around zero up to 10 years before the project begins. Second, cultivation rates start picking up substantially beginning in the first year after completion. This is especially true for the off-season cultivation rates. Cultivation rates are remarkably stable at around 20 percentage points higher in the treated UMVs for the first twenty years. This represents a six-fold increase from 4 to 24 percentage points in the non-rainy season and a tripling from 10 to 30 percentage points cultivated in the rainy season.

The estimates are even larger at around 40 percentage points from years 20 to 25 in the off-season. These are very large effects, considering that average cultivation rates prior to project completion were only 10 percent in the rainy season and 4 percent in the off-season across the sample, implying an increase of about an order of magnitude in land cultivation rates 25 years post project completion in the off-season. For the rainy season, the impacts still represent a substantial doubling of cultivated area at the end of the analysis period. These results dispel the notion that irrigation projects could have only a temporary effect that then disappears. Instead, impacts on land use seem large and sustained.[11]

These patterns broadly hold up when we analyze separately by delegation in Appendix Figures A7 and A8. The only differences between delegations we observe are at the tail ends of the event study. Beyond year 20, confidence intervals in Podor increase substantially - due to few observations which can serve as controls for UMVs with such long post completion data - while in Dagana the increasing trend in cultivation rates is very clear post year 20 during the off-season. We also perform robustness checks with alternative thresholds of 0.4 and 0.5 as well as raw NDVI. The results, shown in Appendix Figures A9-A11, are qualitatively similar, though effects are mechanically smaller in magnitude for greater thresholds. Our estimates for raw NDVI are about 60% larger than those in BenYishay et al. (2024), which are directly comparable to ours. One difference in contexts is that irrigation in Mali is used to extend and supplement agriculture during the rainy season, whereas in Senegal it allows for cultivation in the off-season.

We can also consider average impacts over time for all completed projects. To do so, we set $w_{u,t} = 1\{t - E_{p(u)} > 0\} \times 1\{t == s\}$ for values of $s \in \{1988, \ldots, 2018\}$, generating estimates of the average effect of irrigation in year s.

[11] For this analysis, we considered as a single season the hot and the cold off-seasons. We do this because farmers engage in at most one cultivation cycle in the off-season, either cold or hot but not both, due to temporal overlap between harvesting and land preparation activities. In Appendix Figure A6, we show results for both off-seasons to provide evidence that effects are also evident in both off-seasons when considered separately.

Figure 6 shows the estimated effect of irrigation on off-season cultivation by year. Unlike the event study results, which show fairly stable effects across years since calendar project completion, we find substantial heterogeneity in the effects of irrigation across years. Irrigation had little impact on off-season cultivation rates prior to the year 2000, with estimated effects in all years for which we have data less than 10 percentage points.

Following this period of limited success, there is a discontinuous jump of over 20 percentage points in the effect of irrigation in year 2000. Treatment effects tend to continue growing in subsequent years and are generally statistically significant, though standard errors are very large for the last two years, due to the limited number of remaining control UMVs. The effect of irrigation on off-season cultivation peaks in 2015, at just under 50 percentage points.

One potential spurious cause for the discontinuous jump in 2000 is the launch of Landsat 7 in 1999, which dramatically increased the availability of imagery. Contrary to this hypothesis, Appendix Figure A2 shows that, while there was a dramatic reduction in the number of months with missing imagery after the Landsat 7 began producing science-quality images in May of 1999, most years prior to 2000 have imagery covering the majority of either the cold or hot off-season. This allays any concerns that the jump in 2000 is due to a change in the availability of imagery.

5.2 Heterogeneity in Patterns of Use across UMVs

Given that the impact of irrigation rose substantially in 2000, a natural question is whether this is due in part or in full to a cohort effect. We explore this question in Figure 7, separately estimating treatment effects by year for cohorts of projects completed between 1988-1999, 2000-2009, and 2010-2018. While the new cohort of projects completed between 2000 and 2009 does perform better than older projects at the turn of the century, the oldest projects still experience a sizeable jump in cultivation in 2000 and consistent increase thereafter. These results suggest that, while there is some cohort effect, much of the heterogeneity we observe by year holds across all projects.

Another characteristic of projects that predicts differential impacts of irrigation is whether they were publicly or privately developed. Estimating treatment effects by year for UMVs in public and private projects separately shows that public projects generally perform better, with the starkest difference in the period from 2000 to 2007 (Appendix Figure A12). After 2008, there is substantial convergence, though public projects still outperform private projects in the majority of years. As noted earlier, estimates become substantially noisier for later years given the shortage of UMVs which remain untreated and can serve as clean controls.

This sudden improvement in the cultivation rate followed the economic reforms of the agricultural sector that took place in the early 2000 (section 2.3). A combination of easier access to land that encouraged land development by private actors where irrigation was accessible to them, better access to credit, and input subsidies all favored the agricultural development in the SRV. In addition, in the years after 2000, there was an emphasis on large annual rehabilitation works by SAED. The market reforms of the agricultural sector and an emphasis on rehabilitation appear in the data to have had large impacts on SRV land utilization rates.

5.3 UMV-Level Land Use Classification

Both the event study and the average cultivation rates paint a positive picture of the impact of irrigation project completion on land cultivation patterns in the Senegal river valley. However, this large and fairly stable average impact masks over significant heterogeneity across UMVs with different characteristics and over calendar time, as demonstrated by the preceding results. We thus proceed to characterize this heterogeneity more generally by classifying individual UMV-level time series into four use classes associated to different degrees of average cultivation rates post project completion.

Figure 8 shows one randomly selected example of each class that is associated with a project completed between 2000 and 2005. The top left panel is an example of a "High Use" project UMV, defined as one in which average cultivation rates exceed 67 percent in the post-completion period. In this particular example, cultivation rates shoot up dramatically immediately after the project is finalized - to around 87 percent - and in some years reaching close to 100 percent cultivation rates.

We also observe what we call "Medium Use" cases, in which cultivation rates are between 33 and 67 percent on average post project completion. In the top right panel of the Figure, we observe an example that exhibits a gradual increase in the cultivation rate, reaching 50 percent by year 5, but then stabilizing for a few years. It is only by year 10 that this project reaches above 75 percent cultivation rate for the first time, thus dragging down the average.

The bottom left panel shows an example of a "Low Use" case, defined as those UMVs with an average cultivation rate lower than 33 percent. In this example, the low average post completion cultivation rate is due to a very slow ramp-up and a low ceiling. More than 10 years post completion, cultivation rates are below 30 percent and they almost never pass the 50 percent mark in over 20 years after irrigation is introduced.

Finally, the bottom right panel displays an example of "Intermittent Use". Here, we observe the very high cultivation rates post completion, but also years with low or even zero cultivation rates in an alternating pattern. To generalize this concept quantitatively, for each UMV we calculate the average share of pixels that change cultivation status (from zero to one and vice versa) between years $t-1$ and t in post-project completion years. We use this metric to classify projects as having intermittent land use if cultivation rates vary by more than 10 percent annually on average, excluding years where either observation is missing.

In order to show the overall prevalence of these different use classes, Figure 9 plots UMVs according to their average share of land used and the share of pixels which change cultivation status in the post project completion off-seasons. Here we maintain the definition of "Intermittent Use" with at least 10% of pixels changing off-season cultivation status in the average year, and classify the other UMVs into "Low Use" as those UMVs with an average share in use below 33 percent, "Medium Use" with between 33 and 67 percent, and "High Use" with above 67 percent. In the Figure, each circle size is proportional to the UMV's land area. We observe that 18 percent of land can be classified as Low Use, 37 percent as Intermittent Use, 29 percent as Medium Use, and only 16 percent as High Use. Histograms with the distribution of all four categories of land use by number of units and land area are presented in Appendix Figure A13. Measured in terms of either units or land area, the intermittent use class is the most prevalent land use category. High use UMVs represent a smaller share of land than of number of units, showing that successful high-use UMVs tend to be relatively small in total size.

Because cultivation rates have changed so much over time, it is possible that some UMVs would have a different use class if analyzed using only more recent data. Appendix Figure A14 shows the distribution of use classes when only using data for 2010-2019. While there is an increase in the share of high use UMVs, this is almost entirely driven by reductions in low and medium use. The share of UMVs and land that is intermittently used has remained approximately constant, even when only using more recent data.

In Appendix Figure A15, we show the spatial dispersion of all four use class categories. The figure shows that all four different use classes are scattered throughout the SRV and not concentrated in particular parts of the region, which would raise concerns about unobserved determinants of land use, such as weather patterns.

We can also estimate the main event study figure for the four different land use classes, shown in Appendix Figure A16. Not surprisingly, off-season impacts of project completion are largest for the high use category, and lowest for the low use category. Intermittent and medium use class types track each other closely in terms of impacts on cultivation rates over time. In the rainy season, the effects are less distinguishable between all categories with the exception of the low use class for which we observe low usage rates even during the rainy season.

5.4 Intermittent Usage is not Due to Land Rotation

The finding that 37 percent of all land presents intermittent usage patterns is novel and deserves further exploration. One concern is that observed intermittent usage is in fact due to land fallowing that allows the land to recover before being used again. However, efficient fallowing in a UMV should take the form of rotation of fields in and out of cultivation while keeping overall cultivation shares steady for output and consumption smoothing purposes.

In Appendix Figure A17, we show the share of UMV-level land use changes that are either not rotated at all or totally rotated. Intuitively, total rotation refers to a change in land area into or out of cultivation that is compensated within the same UMV by a change in cultivation in the opposite direction, such that overall area under cultivation does not change. Following the same logic, 'no rotation' refers to a change in area cultivation status for some portion of land in a UMV that is not compensated at all by an opposing change in other parts of the UMV.[12]

In the case of successful UMVs with high cultivation rates, we observe that there are high levels of total rotation. In more than 30 percent of cases, cultivation status changes are fully compensated elsewhere in the UMV. This is very different in low use UMVs, where total rotation is observed in less than 10 percent of observations.

Similarly, we observe that 'no rotation' — where any change in cultivation status corresponds to the total change in cultivation at the UMV level — is more common in the low use class at close to 60 percent of cases. For the successful high use class, no rotation is less prevalent, only observed in close to 40 percent of observations.

We also observe that intermittent use projects behave much more like low use projects than high use projects. This strongly suggests the intermittent use class is not engaged in some type of efficient rotation system, but rather that variability may be caused by time varying constraints to land use.

[12]Formally, rotation is defined as the UMV total change in cultivation minus the UMV net change in cultivation normalized by the UMV total change in cultivation. A value of 1 (total rotation) indicates that an equal amount of land became cultivated and stopped being cultivated between adjacent years. A value of 0 indicates that all change in cultivation was either land becoming newly cultivated or land no longer being cultivated. No Rotation is defined as Rotation< 0.05 and Total rotation is defined as Rotation> 0.95. Net change in cultivation is defined as the difference between area newly cultivated and area that stopped being cultivated, relative to the previous year.

Intermediate levels of rotation are shown in Appendix Figure A18. For all levels of rotation, the intermittent use class looks like a mixture of the low use and medium use classes in terms of rotation patterns.

5.5 Predictors of Use Class

We documented that UMVs broadly fall into four classes of land use patterns: low, intermittent, medium, and high use. A natural question that arises is whether any characteristics of projects or producer organizations are associated with particular use classes. To explore this question, we focus on the most recent 10 years of data which display high levels of usage and allows a comparison across characteristics that is not confounded with early period observations characterized by low usage rates. Table 1 shows the distribution of UMVs across the four use classes by three sets of characteristics: cohorts, private vs publicly managed, and whether the UMV is an agribusiness.

Panel A shows that UMVs from the oldest project cohort are most likely to be low use, whereas the other cohorts are more likely to be in the other use classes. Notably, all three cohorts are mainly distributed between the intermittent and high use classes, suggesting that general improvement in the effects of irrigation over time has led to generally higher levels of cultivation, either constantly or sporadically.

Subsection 5.2 showed that publicly developed projects tend to experience greater benefits from irrigation in terms of off-season cultivation rates than private projects. This finding is echoed in Panel B, which shows that UMVs in private projects have higher shares among the low use type (3 times larger share than public projects) and have 1/3 lower prevalence among the high use class. So even when we focus on the last ten years of data, public projects continue to outperform private ones in terms of land usage.

Panel C shows that agribusinesses are clearly mostly classified as high use (3 out of 4 cases), whereas non-agribusinesses are only classified as high use in about 1 out of 3 cases. While these results are purely descriptive and should not be interpreted as causal, they do reveal a pattern which broadly aligns with our heterogeneity analysis of subsection 5.2.

5.6 Contributions to Uncultivated Land

While the land use rate has rapidly increased since 2000, unused land remains a major issue, with 25% of total land in completed projects being unused in 2019 (Appendix Figure A19). We have documented that UMVs can be classified based on their cultivation rate and share of land that changes cultivation status in a typical year, and that there are a number of characteristics which predict these use patterns. One question that remains is how much each use class contributes to the total uncultivated area which could feasibly be used for production. We can directly measure this number for UMVs in our analysis sample by calculating the share of pixels with a maximum off-season NDVI less than 0.3, multiplying by the area of the UMV, and aggregating by use class.

Figure 10 shows the total off-season uncultivated area of UMVs within completed projects by use class and year. Notably, in almost every year UMVs in the intermittent use class contribute the greatest amount of uncultivated land. Medium use UMVs also contribute about the same amount of uncultivated land each year as low use UMVs. Between 2000 and 2019, about 11,234 ha per year, or one third of the total land in completed projects, was left uncultivated in intermittent and medium use UMVs alone. These findings suggest that substantial increases in cultivation and output could be achieved by targeting interventions to increase land utilization in areas with some evidence of use.

6 Causes for non-use of land

In what follows, we present results from our dedicated farmer survey that allows exploring the bottlenecks holding back over 50 percent of farmland in the SRV in the low and intermittent land use categories. Recall that the phone survey sampled farmers that were known in their community for having some land either intermittently or continuously unused. We asked farmers about the history of land use in their plots during the three off-seasons of 2021-2023.

We partition farmers' plots into three different categories: (i) plots with continuous non-use between 2021 and 2023, which are plots that were not used at all in the three off-seasons, (ii) plots with intermittent non-use between 2021 and 2023, which are plots used in some but not all of the three off-seasons and (iii) plots with continuous use between 2021 and 2023, which are plots used in each of the three off-seasons. In total, we observe 76 plots continuously unused, 126 plots intermittently unused, and 66 plots continuously used between 2021 and 2023.

Panel A of Table 2 presents aggregate statistics for the constraints to land use analysis. Among our sample of farmers with some unused land, we find 45.6 percent of land intermittently unused (59.2 hectares), 31.5 percent of land continuously unused (40.9 hectares), and 22.9 percent of land continuously used (29.7 hectares). Plots in each category are on average about the same size — half a hectare.

In Panel B of Table 2, we present the main constraints to land use by type of constraint separately for continuously unused and intermittently unused land. This breakdown is useful as continuously unused land (column 1) may well be completely unfit for agriculture. Intermittently used land, on the other hand, is certain to be viable for agriculture and inquiring about reasons for its non-use is exactly what we want to shed light on. The percentages shown in each column represent the area-weighted share of plots that cited the constraint in the left column as their main constraint to land cultivation. These percentages thus represent area shares and they add up to 100 in each column.

Note, first, that in spite of all these plots being in projects with completed irrigation infrastructure, inadequate access to water is still cited as the main constraint to land cultivation. This is true for 61.6 percent of continuously unused land and 22.5 percent of intermittently unused land.[13]

Farmer interviews revealed that irrigation canals are not always able to provide water at sufficient levels for cultivation. This can be due to lack of canal maintenance or to land that is far away from the canal. An example of this is the case of a women-only UMV whose water canals have been ruptured for 5 years. These women have not used their land for cultivation ever since.

In the case of intermittent non-use due to lack of water, one farmer reported that his UMV originally had exclusive use to a canal getting its water directly from the river. However, a recent nearby public infrastructure program developed 200 hectares for irrigation, but did not enlarge the canal nor brought in another water source. So now the same canal supplies more projects, with water getting periodically scarce for them, leading to intermittent use of the land.

Financial constraints were found to be quantitatively more important drivers of intermittent non-use of plots (47.8 percent of area) compared to continuous non-use (25.9 percent of area). In one interview, farmers in a large UMV said they were severely financially constrained, to the point where they couldn't afford to get a pump, despite close proximity to a major canal. For another UMV, land was never used due to lack of a water pump capable of bringing sufficient water to the fields. They mentioned that diesel pumps would be adequate but they didn't have the means to afford them.

In theory, functioning credit markets could alleviate some of these financial con-

[13] Stars in column (2) of Table 2 represent the significance level associated with the coefficient ρ in the regression of the form: $Y_i = \alpha + \rho X_i + \epsilon_i$, where Y_i is the variable on the first column of the table, X_i is a dummy variable that takes the value of 1 if the plot is intermittently not used and a value of 0 if it is continuously not used, with standard errors clustered at the farmer level. The sample covers only plots of respondents interviewed in the phone survey. Each observation in the analysis corresponds to a plot in the survey.

straints to deal with issues related to pumps. In practice, however, many farmers reported having trouble accessing financing. Many farmers reported not knowing how to approach getting financing on their own. Others knew how but had trouble finding credit. For example, farmers from one UMV reported having difficulties finding financing for the last few years for a large public project, to the point where they resorted to a wealthy private processor that gave them credit at a heavy premium (in the form of rice bags at harvest). Finally, some farmers mention having no desire to borrow for religious reasons or for fear of not being able to pay back. Farmers from one UMV reported having only one pump for their entire land, although the pump does not have the capacity to supply all of the plots. They do not take bank loans to deal with this issue because they are afraid of not being able to repay them.

The 'Total other constraints' category was significantly higher for intermittently non-used land. These included subcategories 'Projects that are being renovated', as well as problems with 'access to labor'.

Overall, interviewing farmers with some unused land revealed several new insights. First, fallowing was never cited as a reason for non-use, consistent with result of the rotation analysis presented above. Second, and more surprisingly to us, water access continues to be a major constraint to land use, especially for continuously unused land. Third, financial constraints were also important causes of non-use, especially among intermittently unused plots.

After asking farmers about the main cause that explained lack of land cultivation in a plot, we asked whether the issue was one that affected the farm/household only, or whether the problem was common to the whole UMV, or was an issue external to the UMV. For example, if lack of water was identified as being the culprit for non usage of a plot, the farmer was asked to explain if lack of water was due to himself not having brought available water from a nearby canal to his plot (a farm level problem), or whether the UMV was not sending water to the canal serving the farmer (a UMV level problem), or whether it was due to some reason external to the UMV. As an example of the latter, farmers cited the main canal serving the whole UMV not being supplied with enough water. Further, we also asked if the "solution" to the identified constraint was something addressable by the farmer himself, the UMV managers, or would have to be addressed by actors external to the UMV to be solved (for example by SAED).

This analysis is presented in Table 3 for continuously unused land. Recall that for these lands, water access was identified as the most important explanation for non-use. What this table reveals is that, for the most part, the water access problem causing non-use is not seen as something the farmer can address on his own. In almost all cases, the solution is seen as needing to be addressed by actors external to the UMV (Column 6, first and second rows).

In Table 4, we present the analogous analysis for intermittently unused plots. Recall that for these types of plots, financial constraints were quantitatively the most frequent reason for non-use. In this case, both the level of the problem and the solution tend to be at the level of the UMV itself.

The second most important cause for non-use among intermittently unused plots is water access. Table 4 (top panel) shows that for plots intermittently unused due to water constraints, the level of the problem and possible solutions are again not considered something that the farmer himself can address. Rather, both the problem and the solution are more likely to be at the UMV level, or external to the UMV. Since this refers mostly to canals not being well constructed from the start or needing major maintenance interventions, we identify water access issues as something SAED could potentially improve to increase land usage, addressable via improved canal maintenance and infrastructure projects.

6.1 Revealed Preference for Different Constraint-Alleviating Services

We now present results from our survey module asking questions about potential constraint-alleviating services that the respondent could be interested in receiving at zero cost. We asked farmers to express preference between two of the following options selected at random: (i) put you in touch with a financial institution (revealing an interest in actions to alleviate credit constraints), (ii) connect you to someone who is interested in renting out land (revealing an interest in expanding farm size) (iii) connect you to someone who is interested in renting in land (revealing an interest in reducing farm size) (iv) provide you with information on how to improve land quality (revealing an interest in solving land quality issues) (v) put you in touch with someone that can provide reliable labor for your farm (revealing an interest in addressing labor constraints), or (vi) coordinate a meeting with your neighbors to facilitate improvements to the drainage and irrigation system (revealing an interest in addressing water and drainage issues). Our purpose with this relative ranking analysis was to provide options that were costless to the farmer, yet could tell us about the kinds of problems that seemed solvable to them.

To reduce overburdening with an excessive number of similar questions, we presented farmers with a random selection of 3 of the 15 possible pairwise combinations of these improvement options and asked them to choose their preferred options among the two alternatives presented in each pair. We then aggregated these preferences into a single ranking. The aggregation was done using the Schulze method, which has desirable properties for pairwise comparisons and takes into account the full ranking of preferences rather than just the first choice(s). The method identifies the strongest option that is preferred to other options when presented together in pairs. The strongest option is the one most preferred by a majority of respondents in pairwise comparisons. Once the first option is selected, the process is repeated to choose the second option. This second option is the one that wins the most in a series of paired match-ups with other options not selected as the first one. The process is then repeated until all options have been ranked.

In Table 5, we present average relative attractiveness of these different options in a decreasing order. The highest-ranked option by farmers was 'information on how to improve soil quality'. This is in contrast to what we found in Table 2, where land suitability and soil drainage issues were deemed quantitatively unimportant for non-use. We interpret this as something which is of interest to farmers because it could improve outcomes in their continuously used plots (comprising 22.9 percent of their land).

The second most attractive option to farmers was 'improvements to the irrigation and drainage system'. Strong preference for this option is consistent with the results on non-use among farmers as solution to important causes of non-use in both continuously and intermittently unused land.

7 Conclusion

In this paper, we use satellite imagery to evaluate the impact of irrigation infrastructure development on land use in a context where there are no data available on cultivated land and agricultural output over time. We study the case of the Senegal River Valley, where a very large irrigation development program started in the 1980s with construction of two large dams on the Senegal River and has since gradually expanded land under irrigation using canal infrastructure to reach a total of 128,000 ha.

To fill the land use data gap, we turn to a collection of over 3,000 Landsat images collected between 1985 and 2019. We convert the information into a pixel level indicator of whether land is cultivated or not. We then combine this information with administrative data on irrigation infrastructure projects completed between 1988 and 2019. In doing so, we produce a UMV level panel of cultivation rates by season spanning over three decades, along with date of project completion and the nature of the project. We use an event-study specification of agricultural land use to identify the effect of irrigation infrastructure.

Our results provide strong evidence that cultivation rates increase markedly beginning in the first year after project completion, especially during the dry season. On average, we estimate that 20 years after project completion, cultivated land increases by a factor of six during the off-season and by a factor of three during the rainy season.

The intensification of land use was low in the years prior to 2000, but increased dramatically after introduction of policies aimed at boosting investment in agriculture in the early 2000 followed by a sequence of infrastructure rehabilitation projects. Land non-use rates in the off-season fell from an average of 70 percent in 2000-08 to 25 percent at the end of the observation period in 2019.

Analyzing heterogeneity across production units, we find that, despite overall progress in cultivation rates, more than half of production units make either low or intermittent use of the land. High use of the land is seen more frequently in the most recent cohort of projects (2010-19), in projects developed by public agencies rather than privately, and in agribusiness projects. We observe, however, some convergence over time within the three categories of projects.

In spite of the large improvement in cultivation rates since 2000, the non-cultivation problem within completed projects remains a recurrent issue affecting 25 percent of the land as of 2019.

Using a survey of production units, we elicit perceived reasons for under-utilization of irrigation potential. We find that water access problems are a large constraint on cultivation due to deficient infrastructure construction or maintenance issues. Farmers thus consider that improving the irrigation and drainage systems are the most important interventions needed to reduce cultivation failures, and that these interventions require initiatives taken beyond the farm level. These water access constraints are most clearly associated to lands that are continuously unused. In second place, farmers in the SRV seem to be constrained by financial or credit constraints, and these are more prevalent among farmers that only used their land intermittently.

More broadly, this research provides methods to exploit publicly available satellite data to conduct ex-post impact evaluation of important policies for developing country agriculture. We hope these tools will be broadly used by other researchers. One of the most novel findings to come out of these new data sources and methods is the importance of intermittent land usage, a quantitatively important phenomenon that hitherto has not been widely appreciated — most likely because it requires high frequency data on usage patterns over the same plot, something low frequency farmer surveys have trouble identifying. Intermittent land usage is a pheonomenon amenable to policy intervention and a potentially low cost way of boosting agricultural output that may become increasingly relevant in the context of a changing climate.

References

Asher, Sam, Alison Campion, Douglas Gollin, and Paul Novosad. 2022. *The Long-Run Development Impacts of Agricultural Productivity Gains: Evidence from Irrigation Canals in India.* Centre for Economic Policy Research London, UK.

BenYishay, Ariel, Rachel Sayers, Kunwar Singh, Seth Goodman, Madeleine Walker, Souleymane Traore, Mascha Rauschenbach, and Martin Noltze. 2024. "Irrigation Strengthens Climate Resilience: Long-Term Evidence from Mali Using Satellites and Surveys." *PNAS Nexus*, 3(2): 1–9.

Blakeslee, David, Ram Fishman, and Veena Srinivasan. 2020. "Way Down in the Hole: Adaptation to Long-Term Water Loss in Rural India." *American Economic Review*, 110(1): 200–224.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event Study Designs: Robust and Efficient Estimation." *Review of Economic Studies.*

Bourgoin, Jeremy, and Djibril Diop. 2023. "Atlas d'un Territoire en Transition - Regards sur le Delta du Fleuve Sénégal." CIRAD-ISRA.

Bruckmann, Laurent. 2018. "Crue et Développement Rural dans la Vallée du Sénégal: Entre Marginalisation et Résilience." *Belgeo. Revue Belge de Géographie*, 2.

Burke, Marshall, Anne Driscoll, David B Lobell, and Stefano Ermon. 2021. "Using Satellite Imagery to Understand and Promote Sustainable Development." *Science*, 371(6535).

Duflo, Esther, and Rohini Pande. 2007. "Dams." *Quarterly Journal of Economics*, 122(2): 601–646.

Faye, M., M.B. Fofana, and Jean-François Bélières. 1995. "Présentation de la Banque de Données de la SAED pour le Suivi des Aménagements Hydro-Agricoles et des Organisations Paysannes." In *Nianga, Laboratoire de l'Agriculture Irriguée en Moyenne Vallée du Sénégal. Colloques et Séminaires*, , ed. Pascal Boivin, Ibrahim Dia, André Lericollais, Jean Christophe Poussin, Christian Santoir and Sidi Mohamed Seck, 513–533. Paris:ORSTOM, ISRA.

Hein, Lars, and Nico De Ridder. 2006. "Desertification in the Sahel: A Reinterpretation." *Global Change Biology*, 12(5): 751–758.

Higginbottom, Thomas P, Roshan Adhikari, and Timothy Foster. 2023. "Rapid expansion of irrigated agriculture in the Senegal River Valley following the 2008 food price crisis." *Environmental Research Letters*, 18(1): 014037.

Higginbottom, Thomas P, Roshan Adhikari, Ralitza Dimova, Sarah Redicker, and Timothy Foster. 2021. "Performance of Large-Scale Irrigation Projects in Sub-Saharan Africa." *Nature Sustainability*, 4(6): 501–508.

Hornbeck, Richard, and Pinar Keskin. 2014. "The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought." *American Economic Journal: Applied Economics*, 6(1): 190–219.

IPAR. 2015. "Subventions des Intrants Agricoles au Sénégal : Controverses et Réalités."

Le Roy, Xavier. 2000. "La Difficile Mutation de l'Agriculture Irriguée dans la Vallée du Fleuve Sénégal." *Territoires en Mutation*, 165–177.

Le Roy, Xavier. 2011. "Crédit et Production Agricole dans la Vallée du Fleuve Sénégal." IRD Working Paper ird-00593726.

Lobell, David B, George Azzari, Marshall Burke, Sydney Gourlay, Zhenong Jin, Talip Kilic, and Siobhan Murray. 2020. "Eyes in the Sky, Boots on the Ground: Assessing Satellite- and Ground-Based Approaches to Crop Yield Measurement and Analysis." *American Journal of Agricultural Economics*, 102(1): 202–219.

Lobell, David B, Marshall B Burke, Claudia Tebaldi, Michael D Mastrandrea, Walter P Falcon, and Rosamond L Naylor. 2008. "Prioritizing Climate Change Adaptation Needs for Food Security in 2030." *Science*, 319(5863): 607–610.

Lobell, David B, Stefania Di Tommaso, Calum You, Ismael Yacoubou Djima, Marshall Burke, and Talip Kilic. 2019. "Sight for Sorghums: Comparisons of Satellite-and Ground-Based Sorghum Yield Estimates in Mali." *Remote Sensing*, 12(1): 100.

MAER. 2015. "Rapport de Performance 2014." Ministère de l'Agriculture et de l'Equipement Rural, Direction de l'Analyse, de la Prévision et des Statistiques Agricoles.

Ministère du Développement Rural. 1984. "Nouvelle Politique Agricole." République du Sénégal.

Murillo-Sandoval, Paulo J, Emma Gjerdseth, Camilo Correa-Ayram, David Wrathall, Jamon Van Den Hoek, Liliana M Dávalos, and Robert Kennedy. 2021. "No Peace for the Forest: Rapid, Widespread Land Changes in the Andes-Amazon Region Following the Colombian Civil War." *Global Environmental Change*, 69: 102283.

Myneni, Ranga B, Forrest G Hall, Piers J Sellers, and Alexander L Marshak. 1995. "The Interpretation of Spectral Vegetation Indexes." *IEEE Transactions on Geoscience and Remote Sensing*, 33(2): 481–486.

Rosa, Lorenzo, Davide Danilo Chiarelli, Matteo Sangiorgio, Areidy Aracely Beltran-Peña, Maria Cristina Rulli, Paolo D'Odorico, and Inez Fung. 2020. "Potential for Sustainable Irrigation Expansion in a 3°C Warmer Climate." *Proceedings of the National Academy of Sciences*, 117(47): 29526–29534.

République du Sénégal. 1995. "Programme d'Ajustement Sectoriel Agricole, Lettre de Politique de Développement Agricole."

SAED, and JICA. 2019. "Preparatory Survey on Senegal River Valley Irrigated Rice Farming Improvement Project in Republic of Senegal." Japan International Cooperation Agency and Société d'Amenagement et d'Exploitation des Terres du Delta du fleuve Sénégal et des Vallées du fleuve Senegal et de la Falémé.

Seck, Abdoulaye. 2016. "Fertilizer Subsidy and Agricultural Productivity in Senegal." International Food Policy Research Institute (IFPRI) AGRODEP Working Paper 24.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225(2): 175–199.

UN. 2024. "World Population Prospects 2024: Summary of Results." United Nations Department of Economic and Social Affairs, Population Division.

USAID. 2017. "Évaluation des Expériences et du Potentiel Hydroagricole pour la Petite Irrigation dans les Zones Nord et Sud du Sénégal." Feed the Future Senegal-Naatal Mbay, U.S. Agency for International Development.

USDA. 2000. "Senegal Grain and Feed - Rice Annual." United States Department of Agriculture (USDA) Gain Report SG0003.

Wuepper, David, Haoyu Wang, Wolfram Schlenker, Meha Jain, and Robert Finger. 2023. "Institutions and Global Crop Yields." NBER Working Paper 31426.

Xue, Jinru, Baofeng Su, et al. 2017. "Significant Remote Sensing Vegetation Indices: A Review of Developments and Applications." *Journal of Sensors*, 2017.

Zwart, Sander J. 2017. "A Detailed Map of Irrigation Infrastructure along the Senegal River by 2015." AfricaRice Center GIS Report 4.

Figures

Figure 1: Senegal River Valley Study Areas and Irrigated Land Expansion Dates

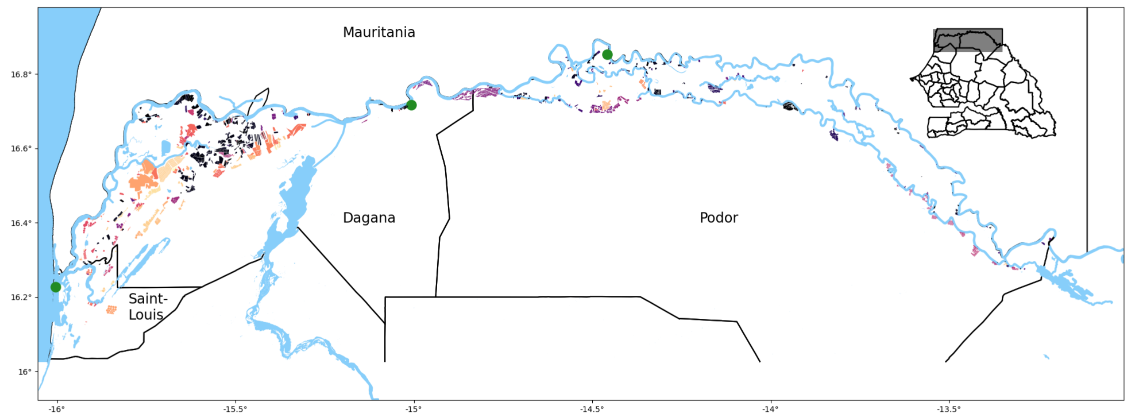

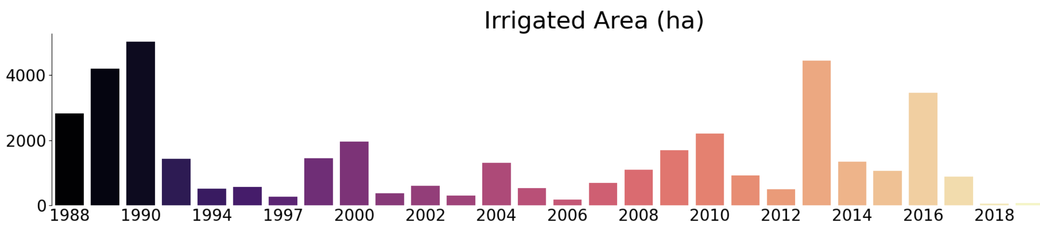

Notes: The map shows productive units (Unités de Mise en Valeur - UMVs) included in the analysis colored by their year of completion, with the histogram along the bottom panel displaying how colors map to years along with the growth in area equipped for irrigation in that year. Water bodies are shown in blue and the three green dots show the location of the capitals of the three departments (level 2 administrative units) in the SRV: Saint-Louis, Dagana and Podor. For the purposes of the analysis, UMVs contained within the department of Saint-Louis are included in the Dagana delegation. The small inset at the top right shows all the departments of Senegal, with the grey rectangle showing the study region.

Figure 2: Land Use Over Time in Dagana and Podor

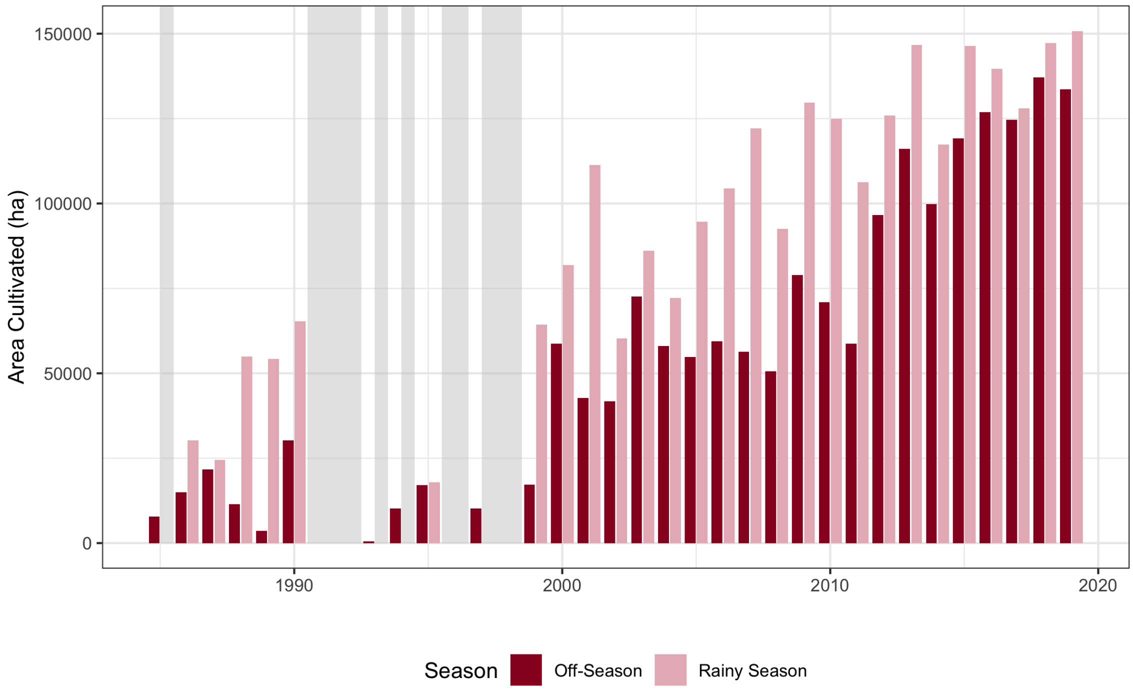

Notes: This figure shows patterns of land use over time in Dagana and Podor. The sample consists of all irrigation projects included in the raw SAED shapefiles as well as any additional area included in the shapefile of Zwart (2017) but not in the SAED shapefiles. The Y-axis denotes the area in hectares that is cultivated. Cultivation is defined at the pixel-level as having a maximum NDVI within the season greater than 0.3. There is no available data for the years 1991, 1992, 1996, and 1998. Additionally, there is no available Rainy Season data for the years 1985, 1993, 1994, and 1997. Missing data is indicated by grey areas.

Figure 3: Visualizing NDVI

Notes: This figure displays an NDVI-priority mosaic of a region in Dagana during the hot off-season 2020. Red, green, and blue values were selected from the image with the greatest NDVI value at that pixel. The UMV with the pink outline corresponds to the lowest NDVI in the figure and the bright green to the greatest. The maximum hot off-season NDVI values in 2020 for these example UMVs are given by the text in their respective color. Other UMVs included in the analysis sample are outlined in black.

Figure 4: NDVI by Month Before and After Project Completion

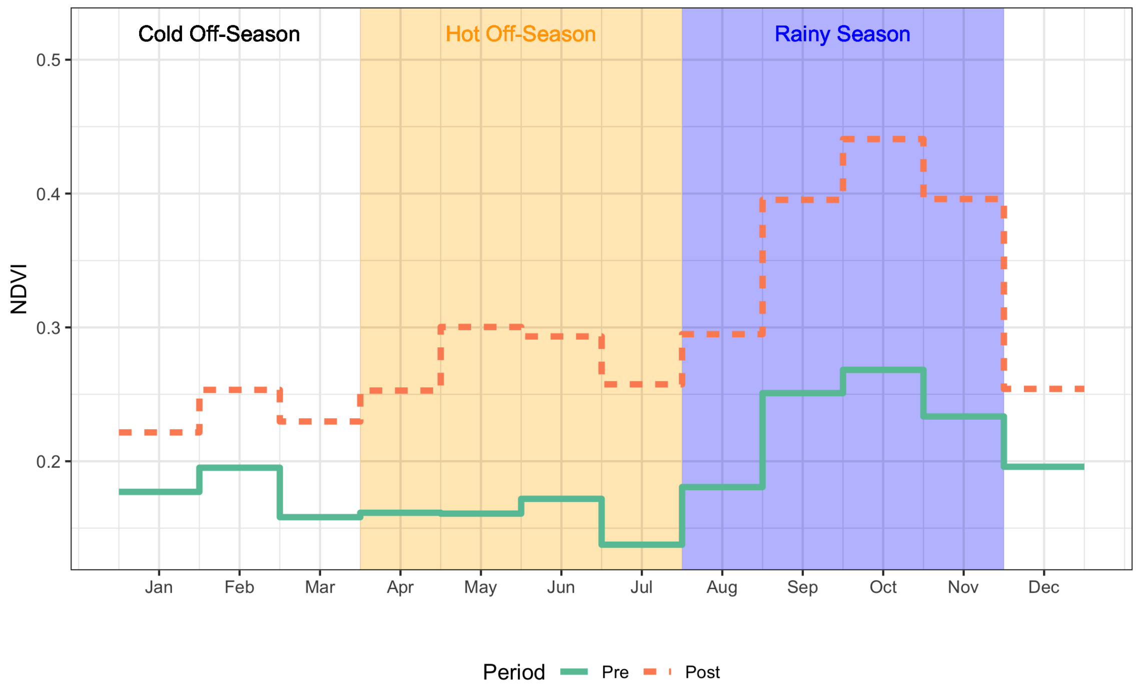

Notes: This figure displays the monthly average maximum NDVI series by construction status of the project for the year 2015. UMVs in the “Pre” group were completed in 2014 or earlier whereas UMVs in the “Post” group were completed in 2016 or later. The orange shaded box corresponds to the hot off-season. The purple shaded box corresponds to the rainy season. The white area for December-March corresponds to the cold off-season.

Figure 5: Effect of Irrigation Project Completion on Cultivation Rate

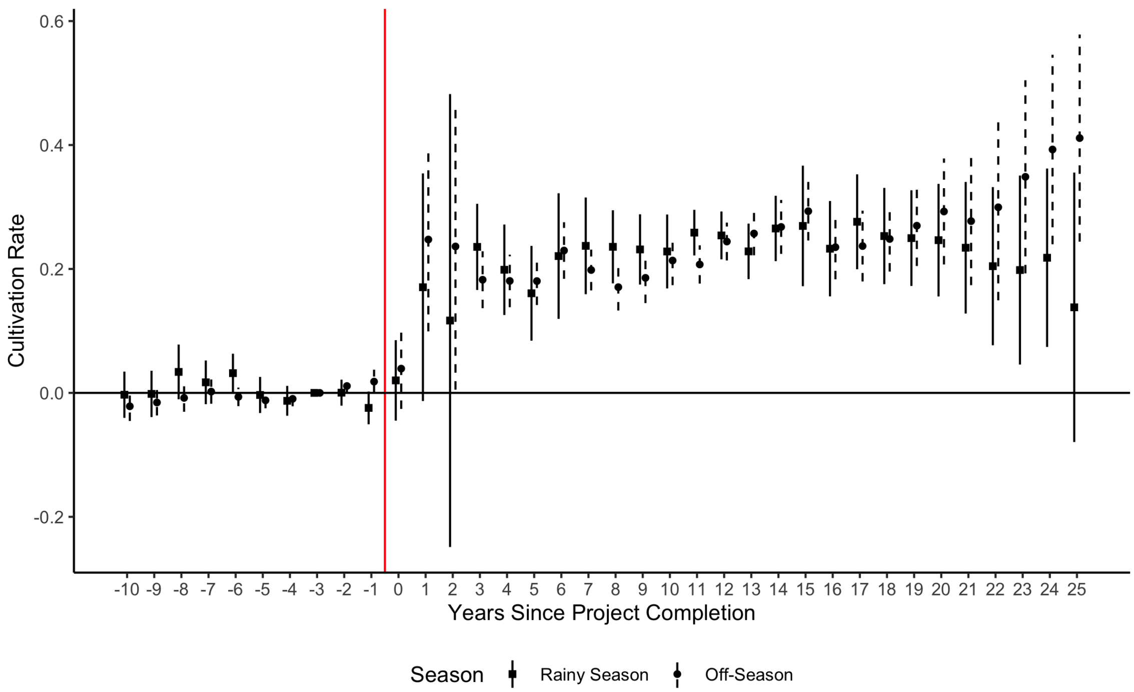

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of UMV land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure 6: Effects of Irrigation Project Completion by Year

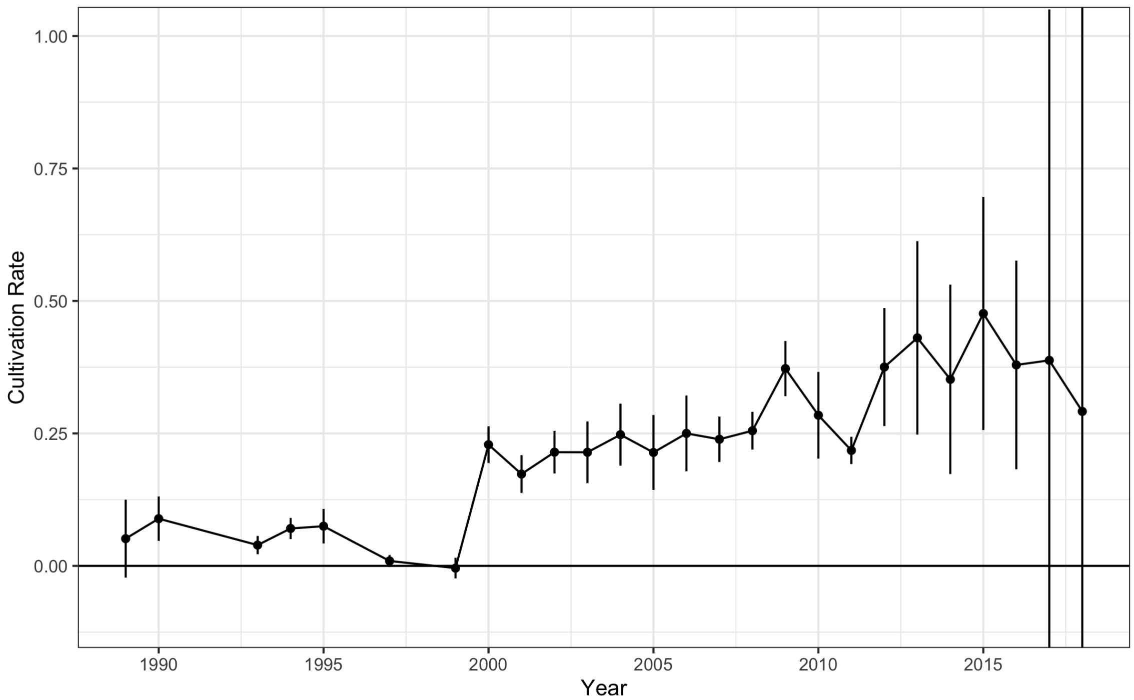

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effect of irrigation project completion on off-season cultivation rate by year. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-March to get a proxy for the share of UMV land in use during the off-season. The average cultivation rate prior to project completion is 0.0402 across the sample. For each year s, we set the weights $w_{u,t} = 1\{t - E_{p(u)} > 0\} \times 1\{t == s\}$ for values of $s \in \{1988, \ldots, 2018\}$, generating estimates of the average effect of irrigation in s. The x-axis shows the year for which the effect is estimated. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from conservative estimates of standard errors clustered at the project level. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure 7: Effects of Irrigation Project Completion by Year and Cohort

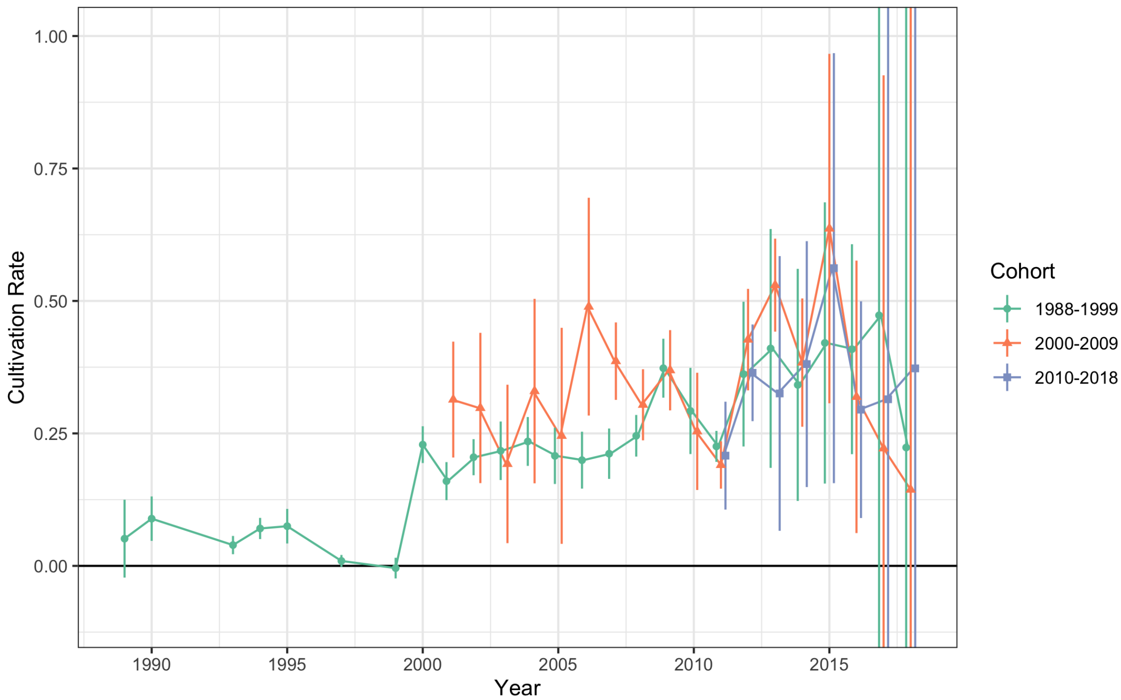

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effect of irrigation project completion on off-season cultivation rate by year and project cohort. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-March to get a proxy for the share of UMV land in use during the off-season. The average cultivation rate prior to project completion is 0.0402 across the sample. The x-axis shows the year for which the effect is estimated. The dots are coefficient estimates while the vertical lines are 95% confidence intervals constructed from conservative estimates of standard errors clustered at the project level. Colors and shapes represent project cohorts, with green circles indicating projects completed between 1988 and 1999, orange triangles between 2000 and 2009, and blue squares between 2010 and 2018. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure 8: UMV-Level Cultivation Series Examples

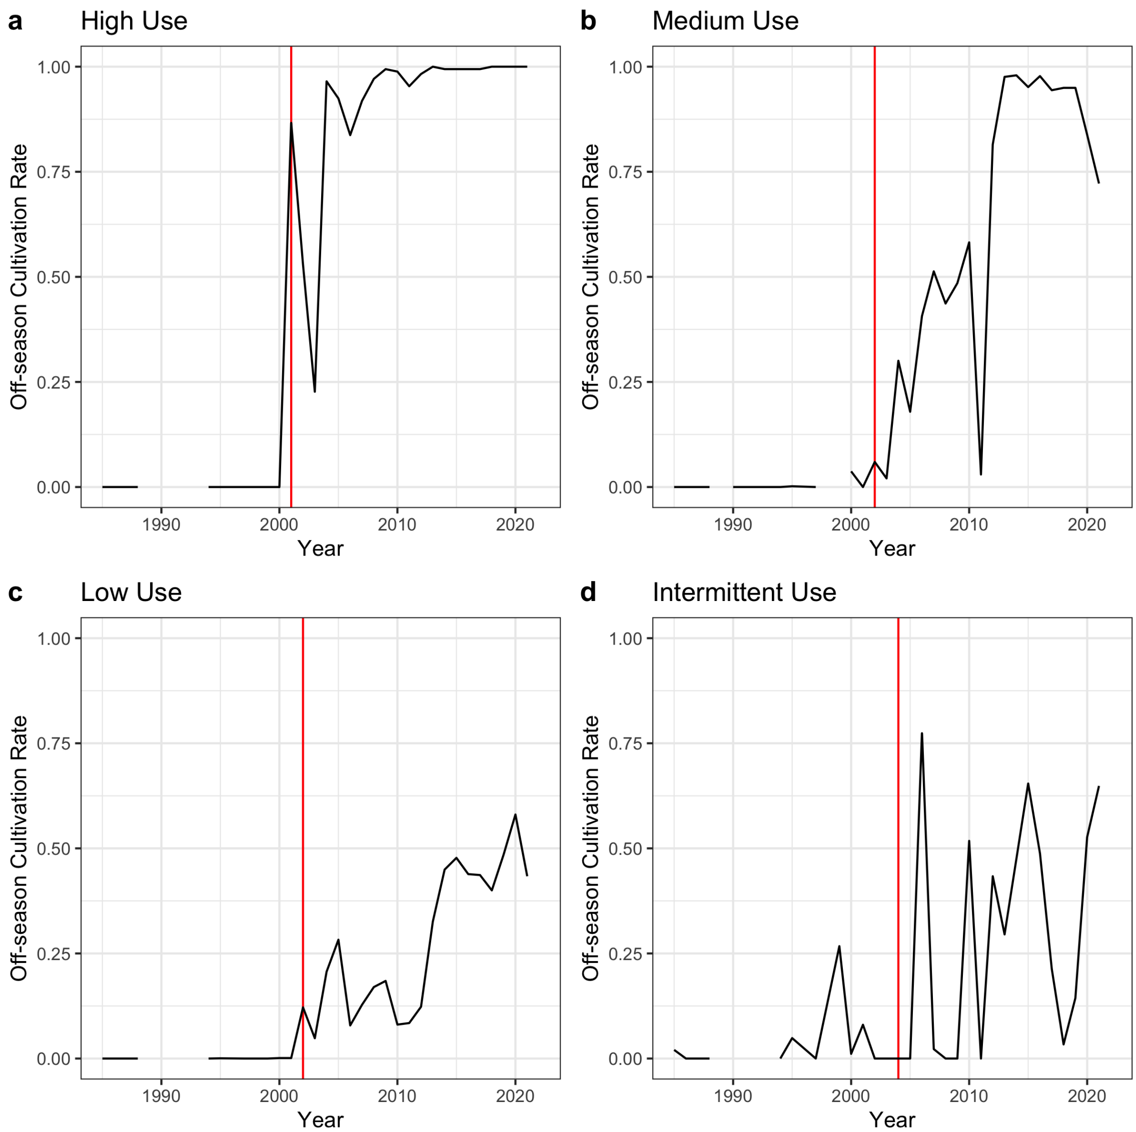

Notes: This figure displays the off-season cultivation rate time series for four UMVs. The vertical red line in each panel is the creation date of the project the UMV is associated with. Panels a-d correspond to high, medium, low, and intermittent use classes respectively. For each use class, the example was selected randomly from the set of UMVs in projects completed between 2000 and 2005.

Figure 9: Definition of Use Classes

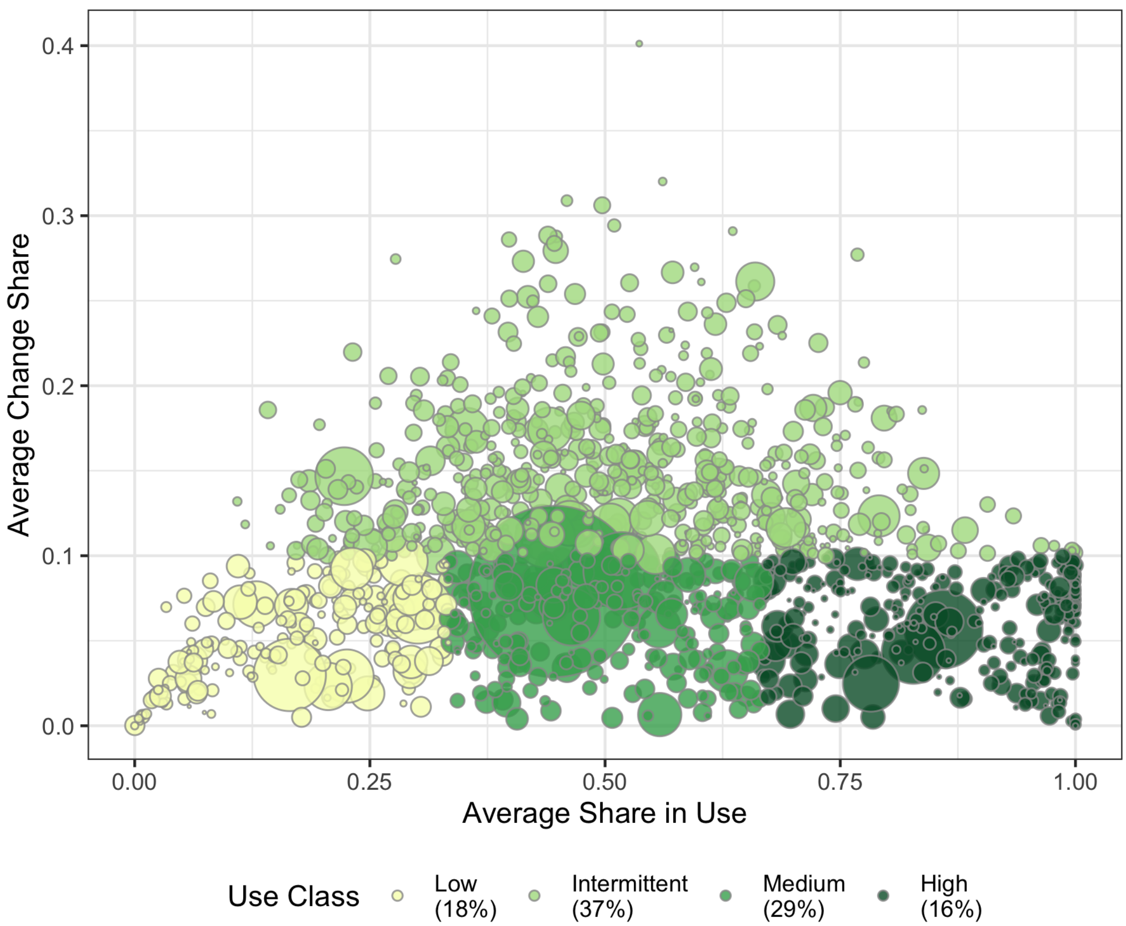

Notes: This figure displays UMVs colored by their use class. The x-axis represents the average share of land cultivated during the off-season in years after project completion. The y-axis shows the average share of land that changes cultivation status during the off-season between the previous and current year for years after project completion. "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year. Area of circles is proportional to the area of the UMV represented by the circles. The shares reported in the legend are based on area.

Figure 10: Uncultivated Area by Use Class

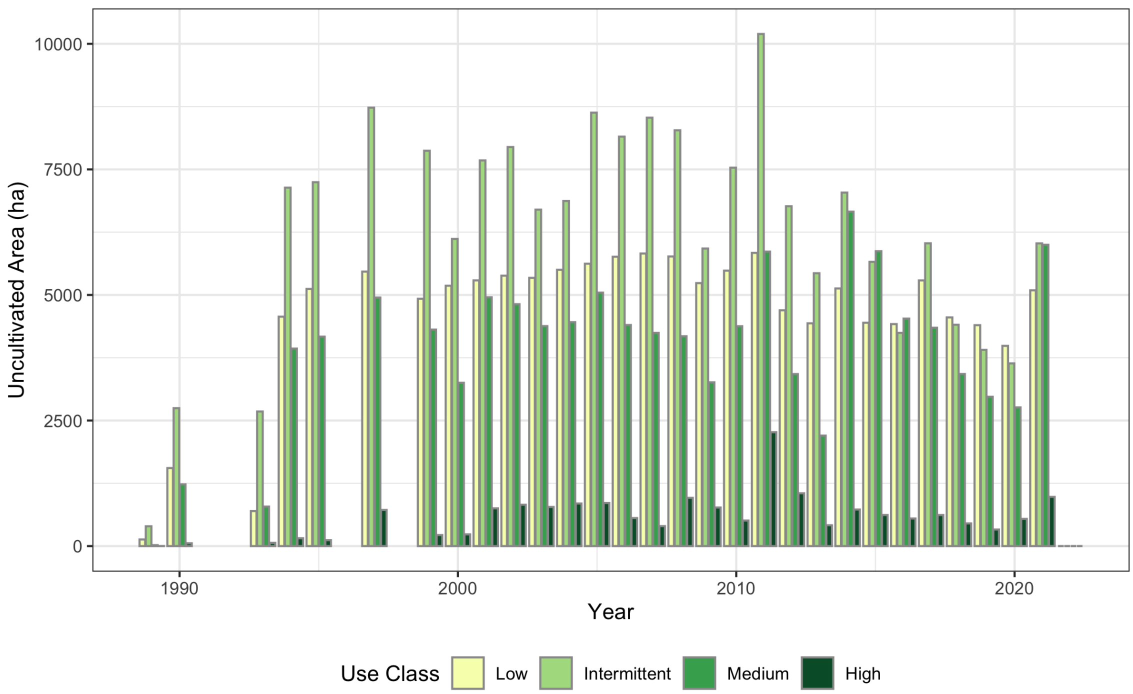

Notes: This figure shows the total uncultivated area by use class and year. For each year and UMV, we calculate the cultivation rate and subtract it from unity. This rate is then multiplied by the area of the UMV and summed across UMVs within a use class×year to get total uncultivated area for the use class. “Low Use” are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Medium Use” are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “High Use” are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Intermittent Use” are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Tables

Table 1: Shares of UMV by Use Class: 2009-2019

	Use Class				
	Low	Intermittent	Medium	High	N
Panel A: Cohort					
Completed 1985-1999	0.11	0.42	0.13	0.35	800
Completed 2000-2009	0.09	0.39	0.12	0.40	234
Completed 2010-2019	0.05	0.33	0.09	0.53	330
Panel B: Project Type					
Private	0.12	0.40	0.13	0.34	795
Public	0.04	0.38	0.10	0.48	569
Panel C: Agribusiness					
Non-Agribusiness	0.10	0.42	0.12	0.36	1203
Agribusiness	0.01	0.19	0.07	0.73	161

Notes: This table displays the share of UMVs with a given characteristic in each use class. The sum across columns may not necessarily sum to exactly 1 due to rounding. "N" represents the number of UMVs in the group described in the first column of the table. "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Table 2: Main Constraints to Land Use

	Continuous Non-Use (1)	Intermittent Non-Use (2)	Continuous Use (3)
Panel A: Aggregate numbers			
Number of plots	76	126	66
Share of plots (%)	28.36	47.01	24.63
Total area (ha)	40.90	59.21	29.65
Share of total area (%)	31.52	45.63	22.85
Average area of plot (ha)	0.54	0.47	0.45
Panel B: Average share of constraints weighted by land size (%)			
Waterways not well constructed	23.5	10.8**	–
Low water level at source point	38.1	6.3***	–
Other water constraints	0.0	5.4**	–
Total water constraints	*61.6*	*22.5****	–
No functional pump	8.6	15.9	–
No access to credit	8.3	18.5*	–
High cost of credit	0.0	0.7	–
No desire to borrow	8.3	8.4	–
Other financial constraints	0.7	4.3**	–
Total financial constraints	*25.9*	*47.8****	–
Infertile/unsuitable land	4.2	1.0	–
Soil drainage issues	1.2	4.4	–
Renovation of projects	1.0	11.6**	–
Labor constraints	0.0	3.7**	–
Input constraints	1.5	3.7	–
Other constraints	4.6	5.3	–
Total other constraints	*12.5*	*29.7***	–

Notes: This table shows the main constraints separately for the different type of non-use in the sample of the phone survey. "Continuous Non-Use" in column (1) refers to plots that were entirely not used by respondents between 2021 and 2023. "Intermittent Non-Use" in column (2) refers to plots with land used only for one or two of the three off-seasons between 2021 and 2023. 'Continuous Use' refers to plots that were entirely used for each off-season between 2021 and 2023. The shares in each column represents the share of plots that cited the constraint in the first column as their main constraint. The shares are weighted by the size of the constrained areas and they add up to 100 for each column. Stars in column (2) represent the significance level associated with the coefficient ρ in the regression of the form: $Y_i = \alpha + \rho X_i + \epsilon_i$, where Y_i is the variable on the first column of the table, X_i is a dummy variable that takes the value of 1 if the plot is intermittently not used and a value of 0 if it is continuously not used, with standard errors clustered at the farmer level. ***: p-value $<$ 0.01, **: p-value $<$ 0.05 and *: p-value $<$ 0.1. The sample covers only plots of respondents interviewed in the phone survey. Each observation in the analysis corresponds to a plot in the survey.

Table 3: Level of Problems and Solutions for Land Continuously Unused

	Level of Problem			Level of Solution			
	Farm/HH	UMV	External	Farm/HH	UMV	External	% Area
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Water constraints							
Waterways not well constructed	0.0	36.4	63.6	0.0	4.5	95.5	23.5
Low water level at source point	0.00	24.0	76.0	0.0	0.0	100.00	38.1
Other water constraints	–	–	–	–	–	–	0.0
Financial constraints							
No functional pump	0.0	25.0	75.0	0.0	12.5	87.5	8.6
No access to credit	25.0	25.0	50.00	0.0	25.0	75.0	8.3
No desire to borrow	100.0	0.0	0.0	20.0	40.0	40.0	8.30
Other financial constraints	0.0	0.0	100.0	0.0	0.0	100.0	0.7
Other constraints							
Infertile/unsuitable land	0.0	25.0	75.0	0.0	0.0	100.0	4.20
Soil drainage issues	0.0	100.0	0.0	0.0	0.0	100.0	1.2
Renovation of projects	0.0	0.0	100.0	0.0	0.0	100.0	1.0
Labor constraints	–	–	–	–	–	–	0.0
Input constraints	0.0	100.0	0.0	0.0	0.0	100.0	1.5
Other constraints	66.7	33.3	0.0	33.3	66.7	0.0	4.6

Notes: This table shows the level of the problem and the level of the potential solution for land continuously not used, which includes all plots that were entirely not used by respondents between 2021 and 2023. The first column shows the main constraint reported by the respondent for the plot-season pair. Columns (1)-(3) show the level at which the constraint binds. Columns (4)-(6) show the level at which the respondent expects the potential constraint alleviation tool to be at. Column (7) show the share of total area subject to the constraint. Each cell represents the share of total land that falls under the level of problem/solution specified in the corresponding column. Numbers in the same row of columns (1), (2) and (3) add up to 100. Numbers in the same row of columns (4), (5) and (6) add up to 100. The sample is restricted to land continuously not used from Table 2.

Table 4: Level of Problems and Solutions for Land Intermittently Unused

	Level of Problem			Level of Solution			
	Farm/HH	**UMV**	**External**	**Farm/HH**	**UMV**	**External**	**% Area**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Water constraints							
Waterways not well constructed	12.1	49.2	38.7	4.8	23.4	71.8	23.5
Low water level at source point	0.0	20.0	80.0	0.0	0.0	100.0	5.8
Other water constraints	0.0	71.4	28.6	0.0	0.0	100.0	5.4
Financial constraints							
No functional pump	0.0	78.9	21.1	0.0	68.4	31.6	15.9
No access to credit	7.7	76.9	15.4	0.0	61.5	38.5	16.5
High cost of credit	0.0	100.0	0.0	50.0	50.0	0.0	0.7
No desire to borrow	72.7	27.3	0.0	18.2	50.0	31.8	7.5
Other financial constraints	25.4	60.3	14.3	0.0	28.6	71.4	4.3
Other constraints							
Infertile/unsuitable land	0.0	0.0	100.0	0.0	0.0	100.0	0.3
Soil drainage issues	0.0	0.0	100.0	0.0	0.0	100.0	4.4
Renovation of projects	0.0	25.0	75.0	0.0	0.0	100.0	9.3
Labor constraints	100.0	0.0	0.0	40.0	60.0	0.0	2.1
Input constraints	0.0	100.0	0.0	0.0	0.0	100.0	0.4
Other constraints	11.1	88.9	0.0	11.1	77.8	11.1	4.0

Notes: This table shows the level of the problem and the level of the potential solution for land intermittently not used, which includes all plots with land used only for one or two of the three off-seasons between 2021 and 2023. The first column shows the main constraint reported by the respondent for the plot-season pair. Columns (1)-(3) show the level at which the constraint binds. Columns (4)-(6) show the level at which the respondent expects the potential constraint alleviation tool to be at. Column (7) show the share of total area subject to the constraint. Each cell represents the share of total land that falls under the level of problem/solution specified in the corresponding column. Numbers in the same row of columns (1), (2) and (3) add up to 100. Numbers in the same row of columns (4), (5) and (6) add up to 100. The sample is restricted to land intermittently not used from Table 2.

Table 5: Ranking of Improvement Options in Phone Survey

Improvement Option	**Relative Ranking**
Provide you with information on how to improve soil quality	1
Coordinate a meeting to facilitate improvements to the drainage and irrigation system	2
Put you in touch with someone for reliable labor	3
Put you in touch with a financial institution	4
Connect you to someone interested in renting out land	5
Connect you to someone interested in renting in land	5

Notes: This table shows the ranking of six improvement options based on responses to 3 randomly selected pairwise questions presented to respondents. The aggregation and ranking of responses is done using the Schulze method, which identifies the strongest option that is preferred to other options when presented together in pairs, and considers as the strongest option the one most preferred by a majority of respondents in pairwise comparisons. Once the first option is selected, the process is repeated to choose the second option among the remaining options in the same manner. The process is then repeated until all options are ranked.

Online Appendix Figures

Figure A1: External Validity of Analysis Sample

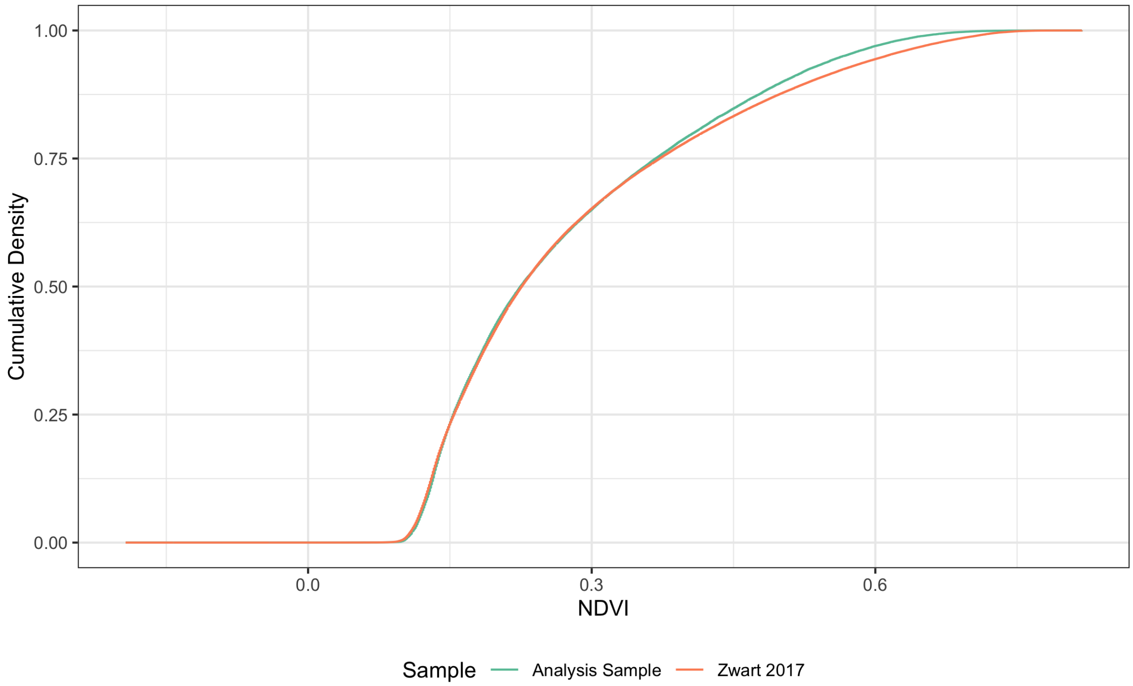

Notes: This figure shows the cumulative density of off-season NDVI observations in our analysis sample and the sample defined by Zwart (2017). In our analysis sample, each observation is a UMV-year-level mean of pixel-year-level maximum off-season NDVI. In the Zwart (2017) sample, each observation is the mean of pixel-year-level maximum off-season NDVI by polygon. The off-season runs from December to July.

Figure A2: Availability of Imagery Over Time

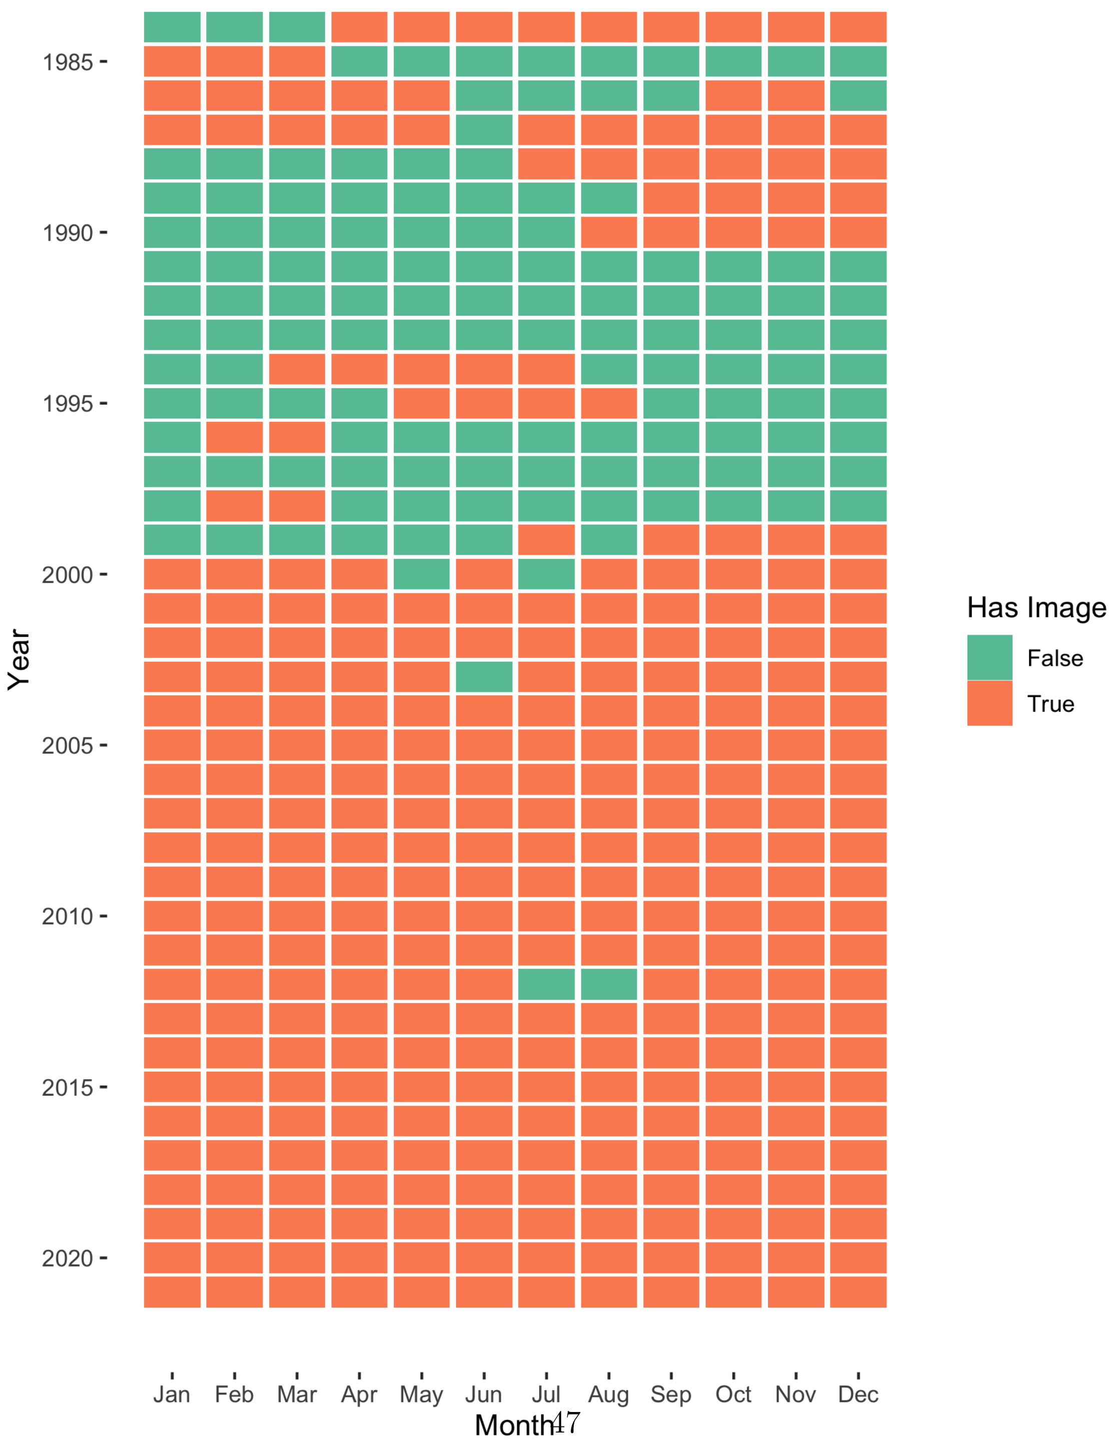

Notes: This figure displays the monthly availability of imagery of the study region. The box for a month is marked "True" if there is at least one usable Landsat image intersecting the study region.

Figure A3: 95th Percentile of NDVI in Uncompleted Projects by Year

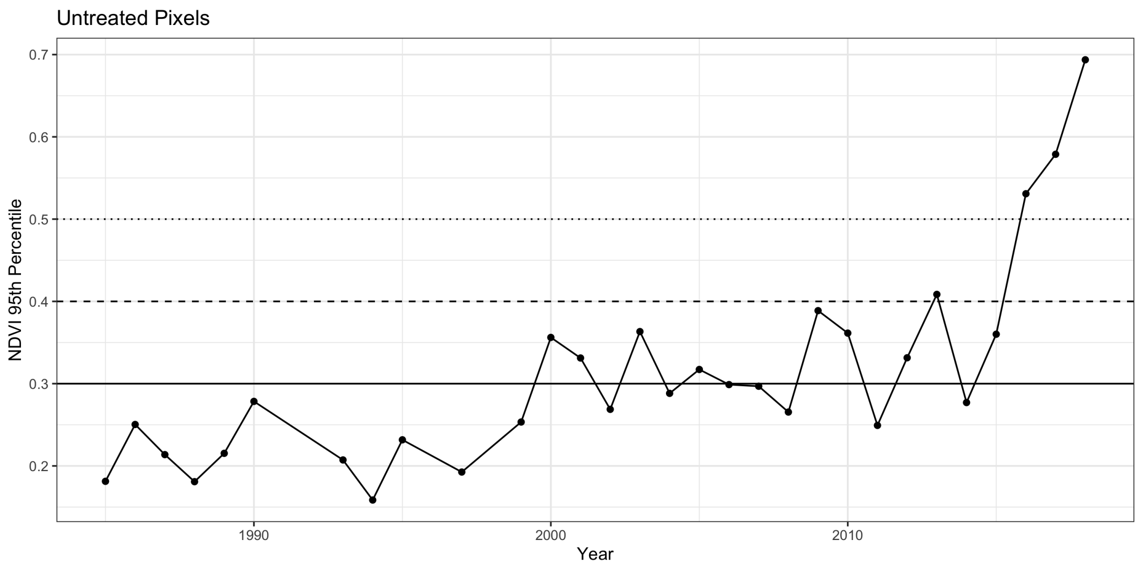

Notes: This figure shows the 95th percentile of NDVI values for pixels in projects that were not completed by the year given on the x-axis. Our main results use a threshold of 0.3, shown with a solid horizontal line. We show robustness using thresholds of 0.4 and 0.5, shown as dashed and dotted horizontal lines respectively.

Figure A4: NDVI by Month Before and After Project Completion: Dagana

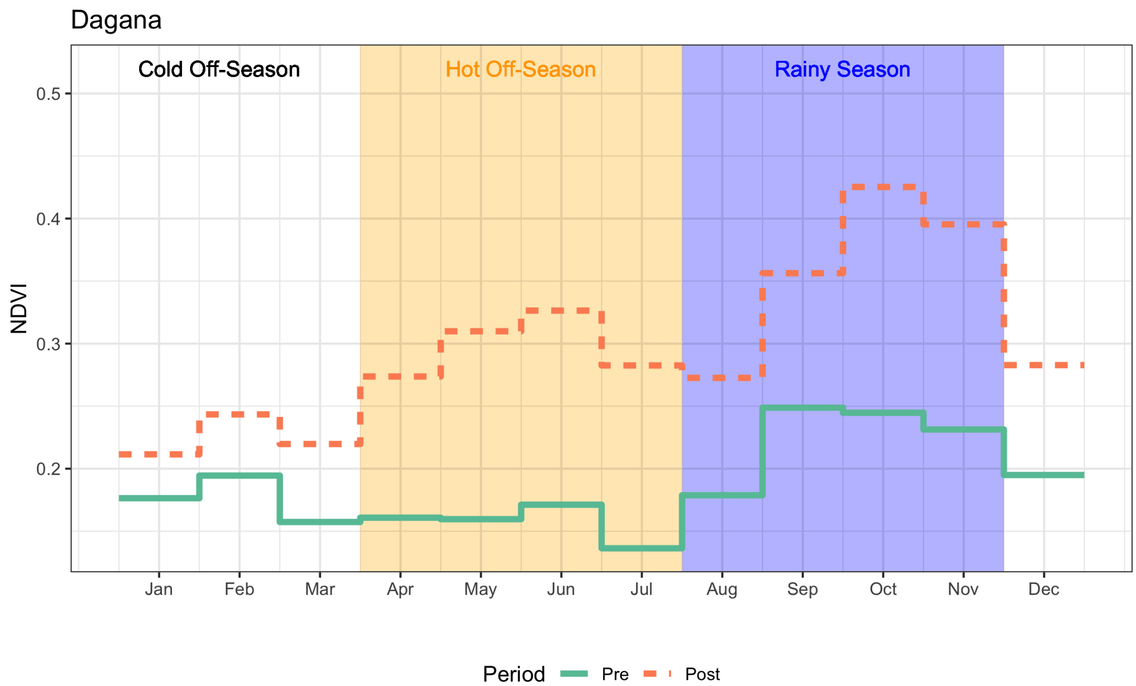

Notes: This figure displays the monthly average maximum NDVI series by construction status of the project for 2015-2017. The orange shaded box corresponds to the hot off-season. The purple shaded box corresponds to the rainy season. The white area for December-March corresponds to the cold off-season. The sample is limited to projects in Dagana.

Figure A5: NDVI by Month Before and After Project Completion: Podor

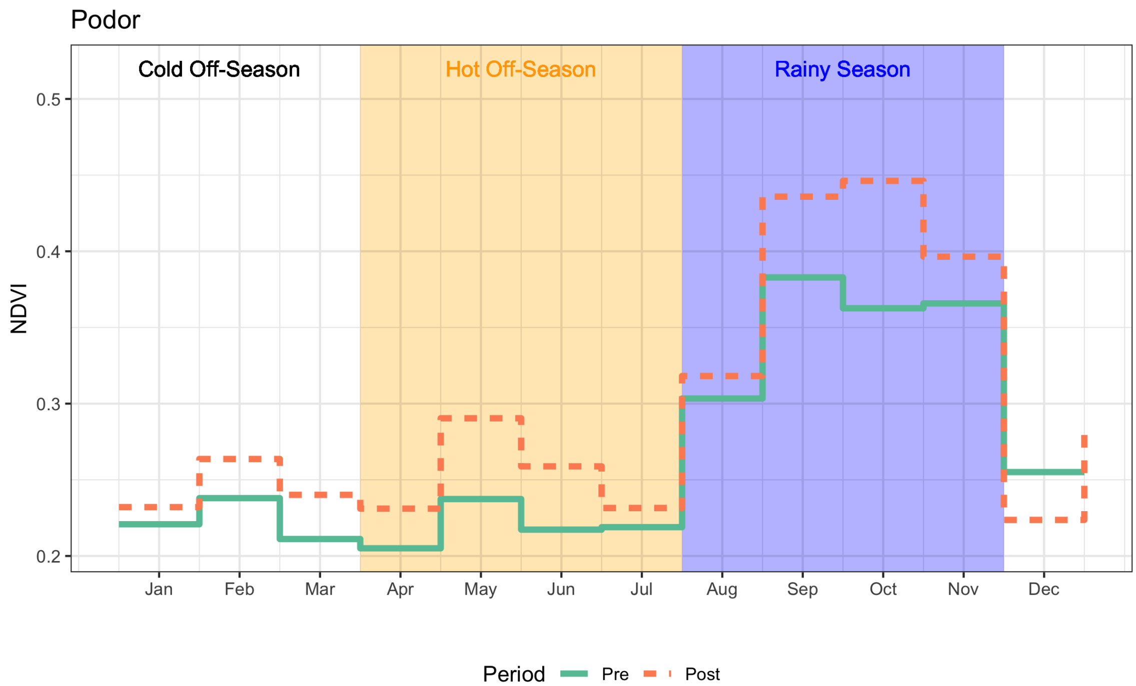

Notes: This figure displays the monthly average maximum NDVI series by construction status of the project for 2015-2017. The orange shaded box corresponds to the hot off-season. The purple shaded box corresponds to the rainy season. The white area for December-March corresponds to the cold off-season. The sample is limited to projects in Podor.

Figure A6: Effect of Irrigation Project Completion on Cultivation Rate

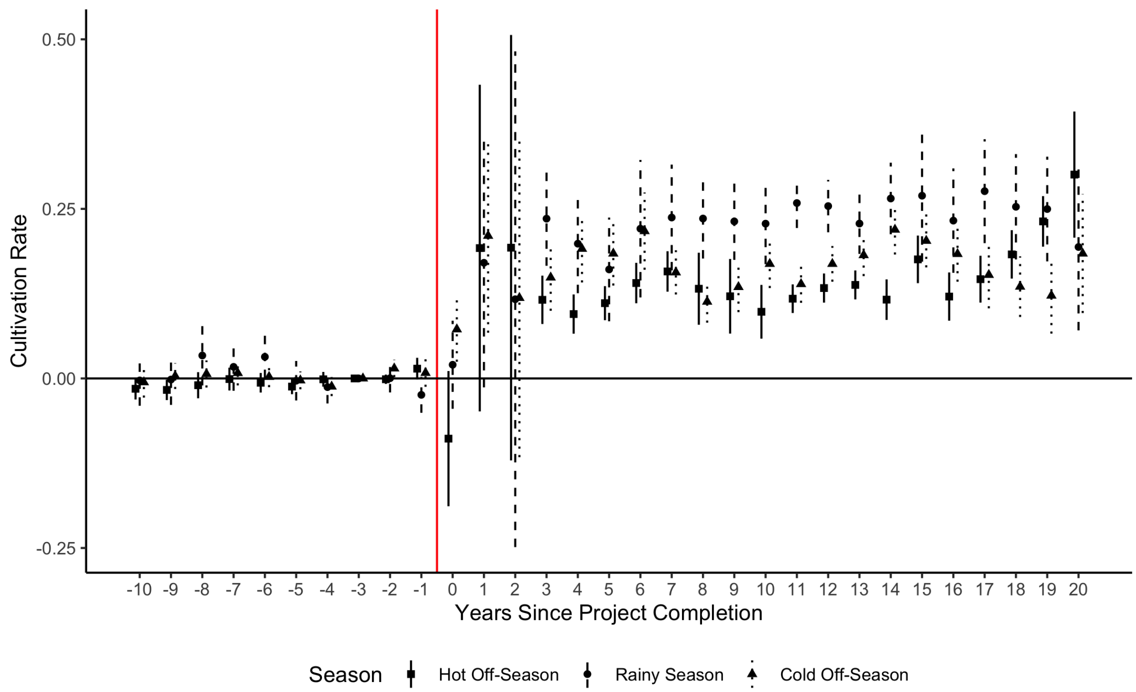

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.3 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the project level over the period April-July (solid lines with squares) corresponding to the hot off-season, August-November (dotted lines with circles) corresponding to the rainy season, and December-March (long dashed lines with triangles) corresponding to the cold off-season, to get a proxy for the share of project land in use. Average cultivation rates prior to project completion are 0.009 (April-July), 0.097 (August-November) and 0.028 (December-March) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A7: Effect of Irrigation Project Completion on Cultivation Rate: Dagana

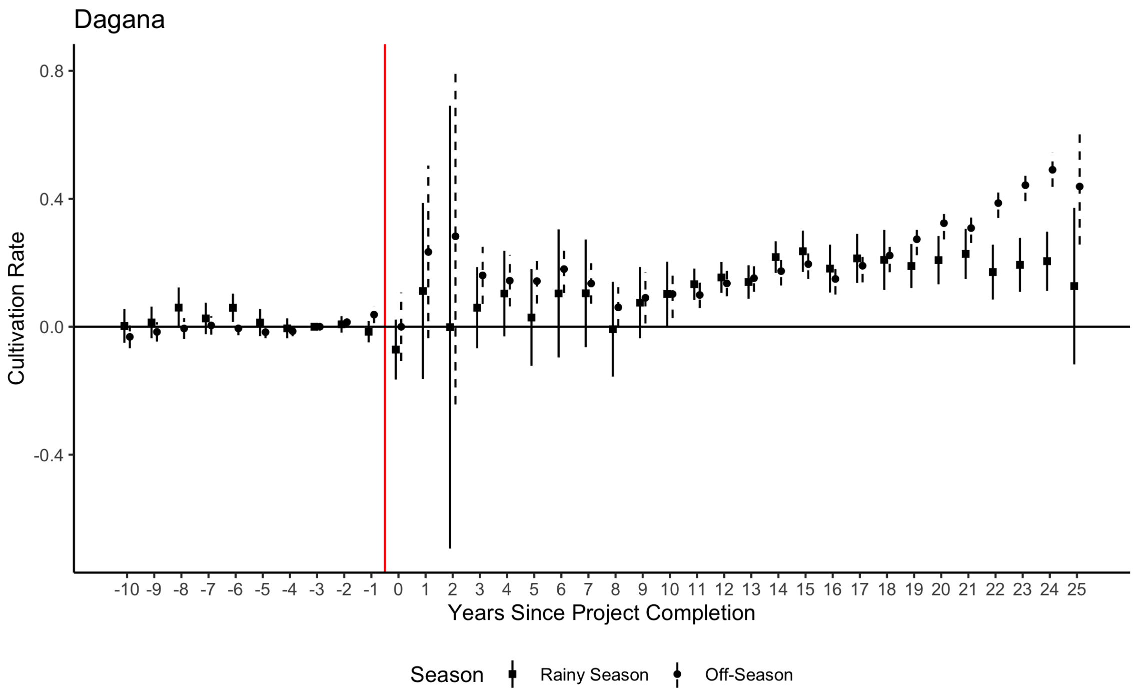

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of UMV land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 788 UMVs within Dagana in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A8: Effect of Irrigation Project Completion on Cultivation Rate: Podor

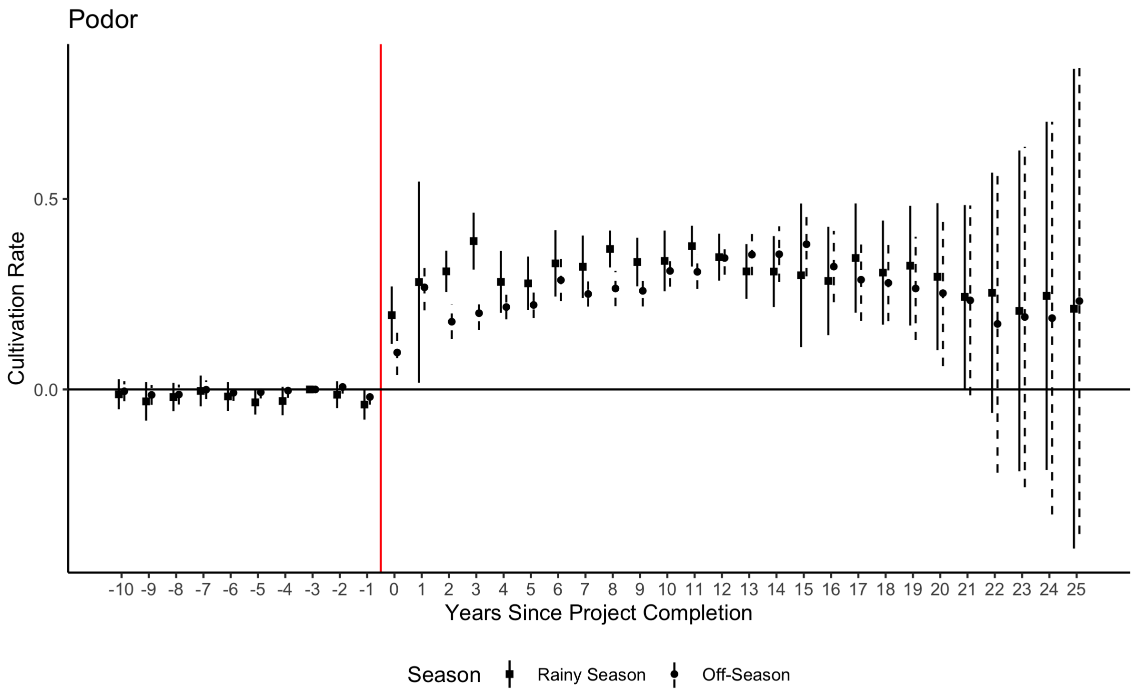

This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of UMV land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 576 UMVs within Podor in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A9: Effect of Irrigation Project Completion on Cultivation Rate (NDVI$\geq$ 0.4)

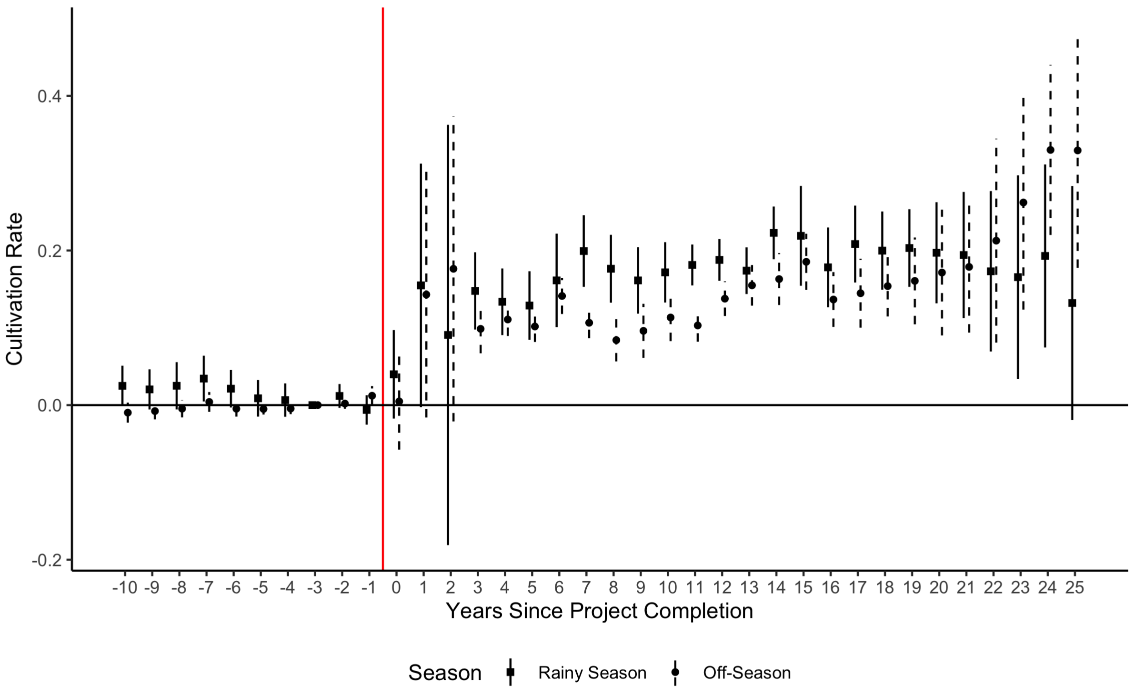

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with an alternative NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.40 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of UMV land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A10: Effect of Irrigation Project Completion on Cultivation Rate (NDVI≥ 0.5)

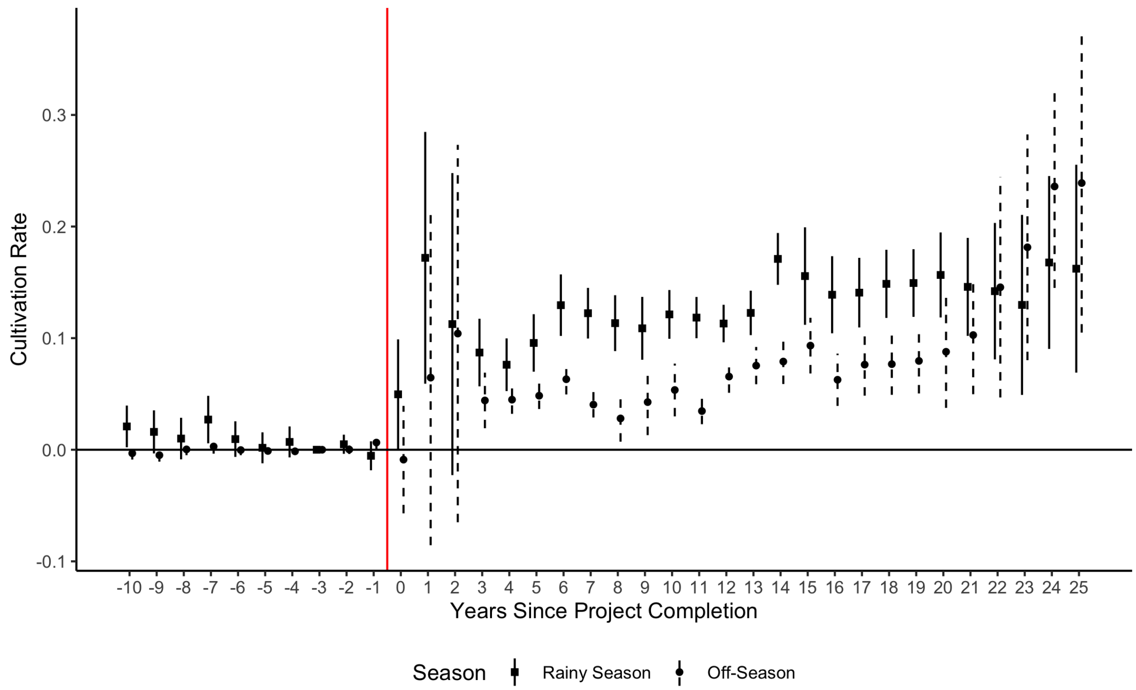

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with an alternative NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.50 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of UMV land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A11: Effect of Irrigation Project Completion on NDVI

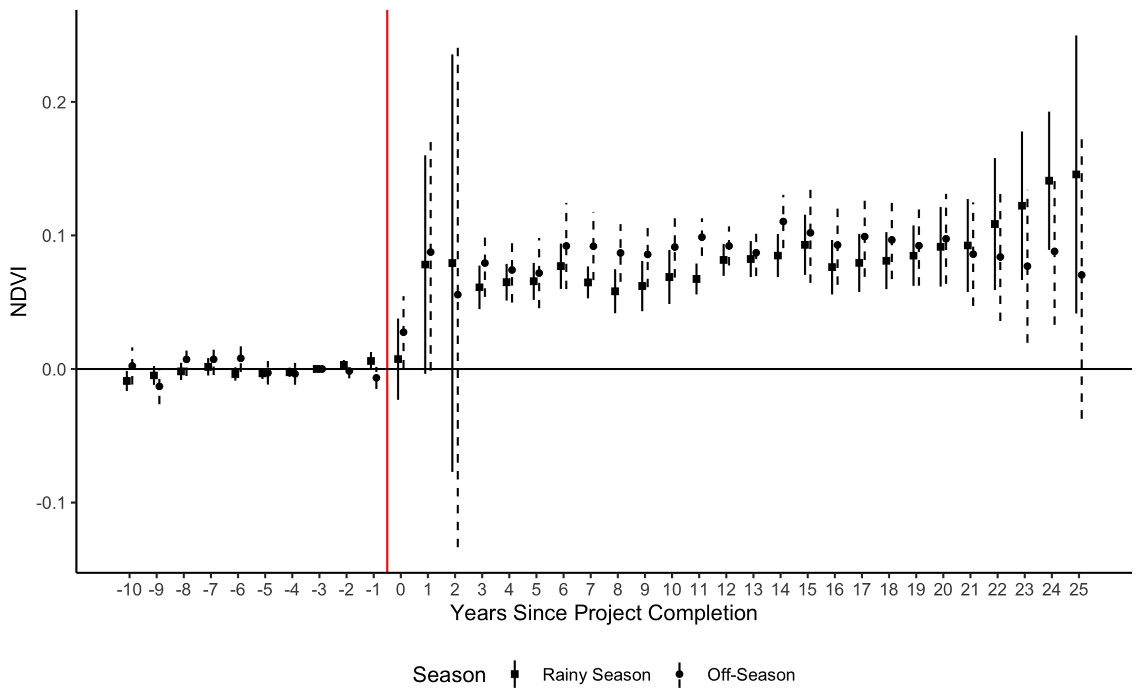

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with raw NDVI as the outcome. The outcome is constructed by averaging pixel-level NDVI at the UMV level over the period August-November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of UMV land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A12: Effects of Irrigation Project Completion by Year and Sector

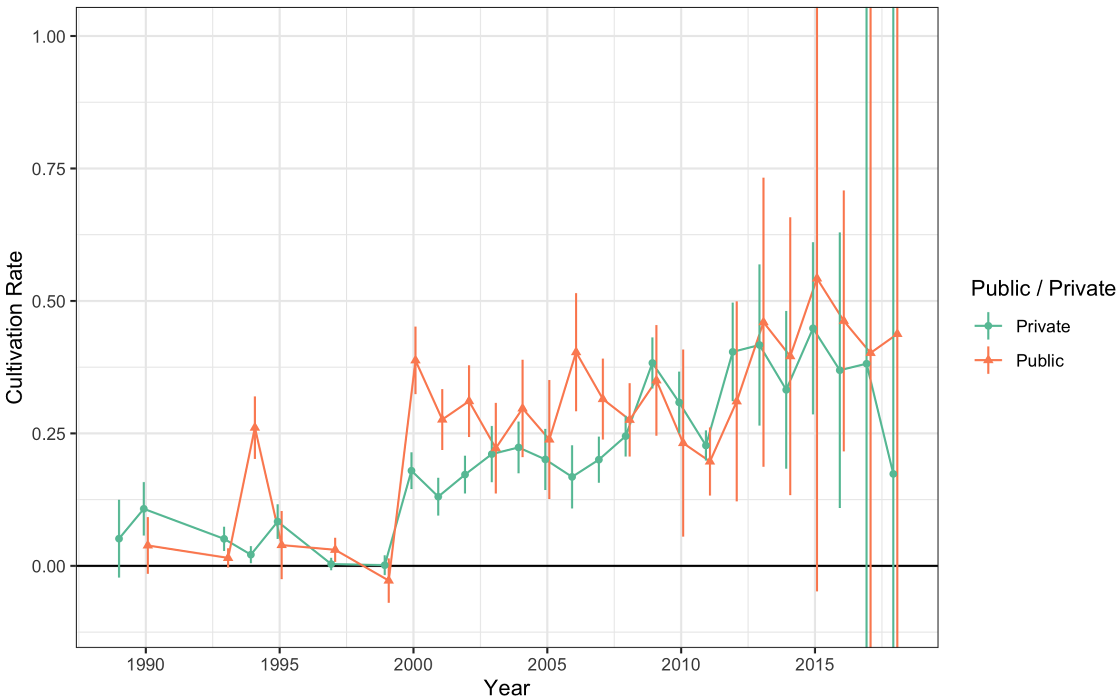

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effect of irrigation project completion on off-season cultivation rate by year and whether the project was publicly or privately developed. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the UMV level over the period August-March to get a proxy for the share of UMV land in use during the off-season. The average cultivation rate prior to project completion is 0.0402 across the sample. The x-axis shows the year for which the effect is estimated. The dots are coefficient estimates while the vertical lines are 95% confidence intervals constructed from conservative estimates of standard errors clustered at the project level. Colors and shapes represent the sector which developed the projects, with green circles indicating private development, orange triangles public development. The sample is restricted to the 1364 UMVs in projects constructed in 1988 or later which could be merged with the shapefile.

Figure A13: Distribution of Use Classes

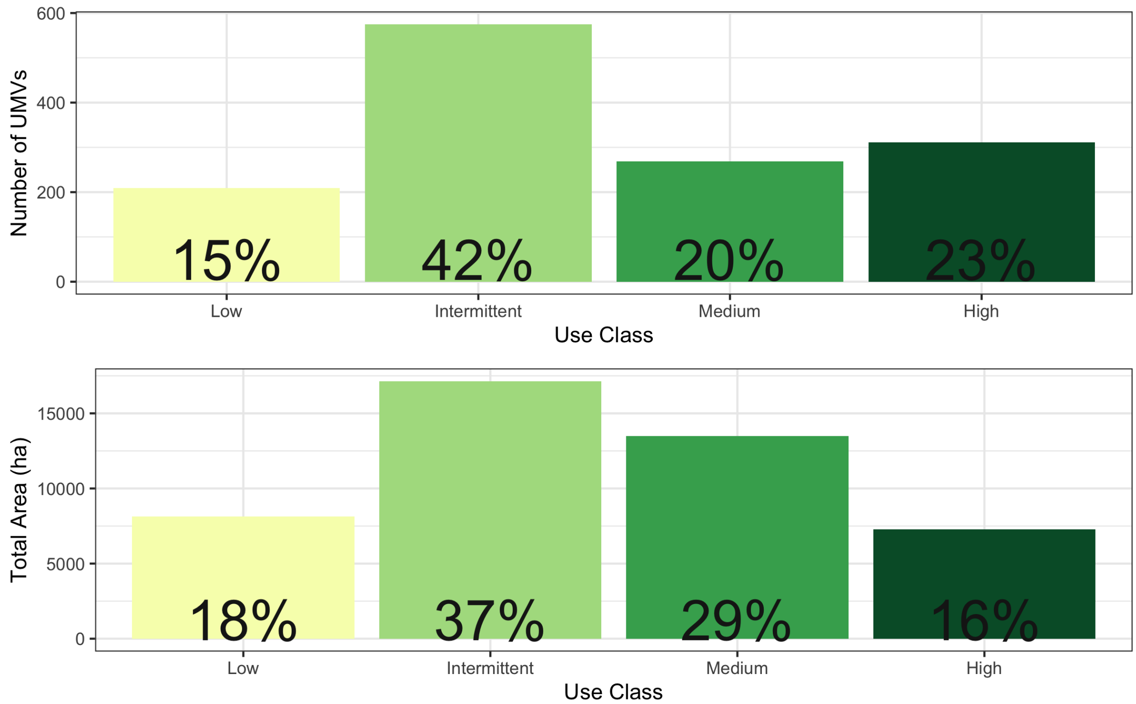

Notes: This figure shows the distribution of UMVs by use class. The top row shows the number of UMVs by use class. The bottom row shows the total area in hectares of UMVs by use class. "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Figure A14: Distribution of Use Classes: 2010-2019

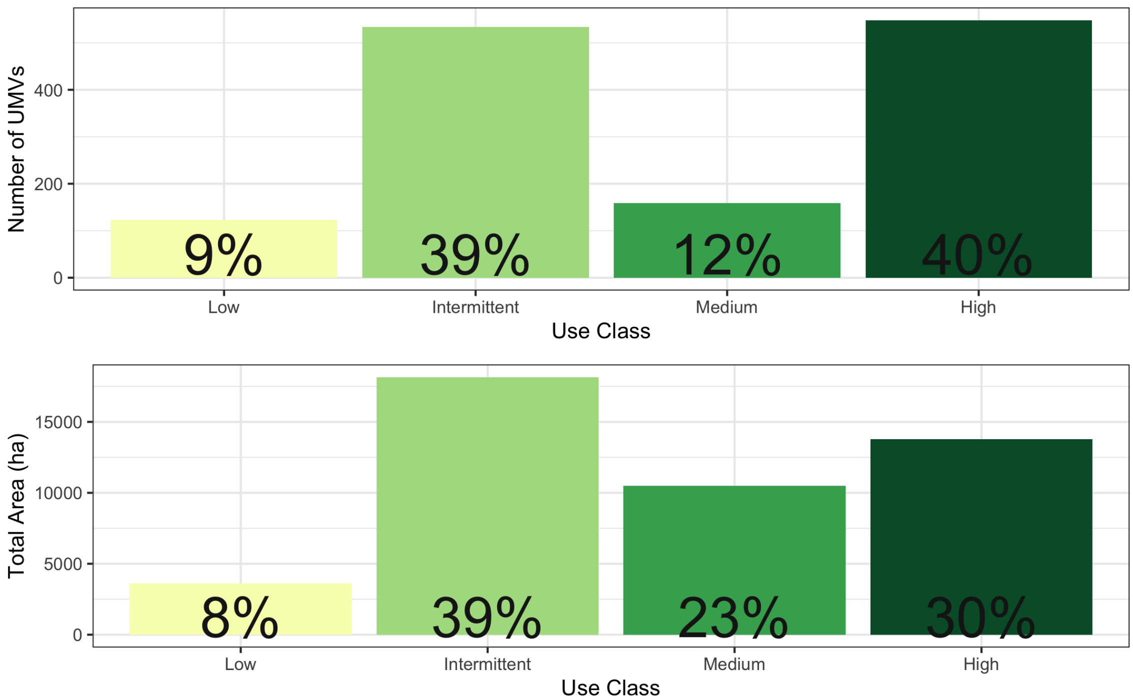

Notes: This figure shows the distribution of UMVs by use class when calculated using only data on use and changes in cultivation status from 2010-2019. The top row shows the number of UMVs by use class. The bottom row shows the total area in hectares of UMVs by use class. "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Figure A15: Map of Use Classes

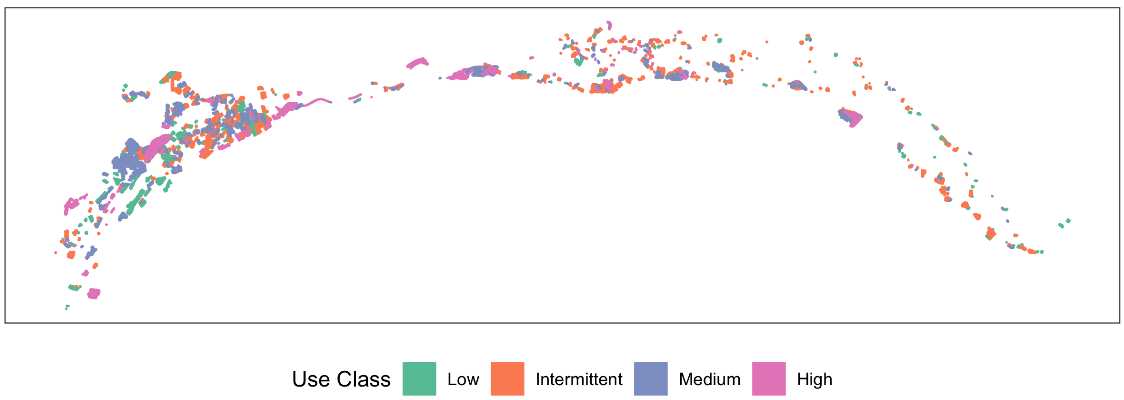

Notes: This map shows the spatial dispersion of UMVs by use class across the SRV. “Low” are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Medium” are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “High” are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Intermittent” are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Figure A16: Cultivation Rate Event Study by Use Class

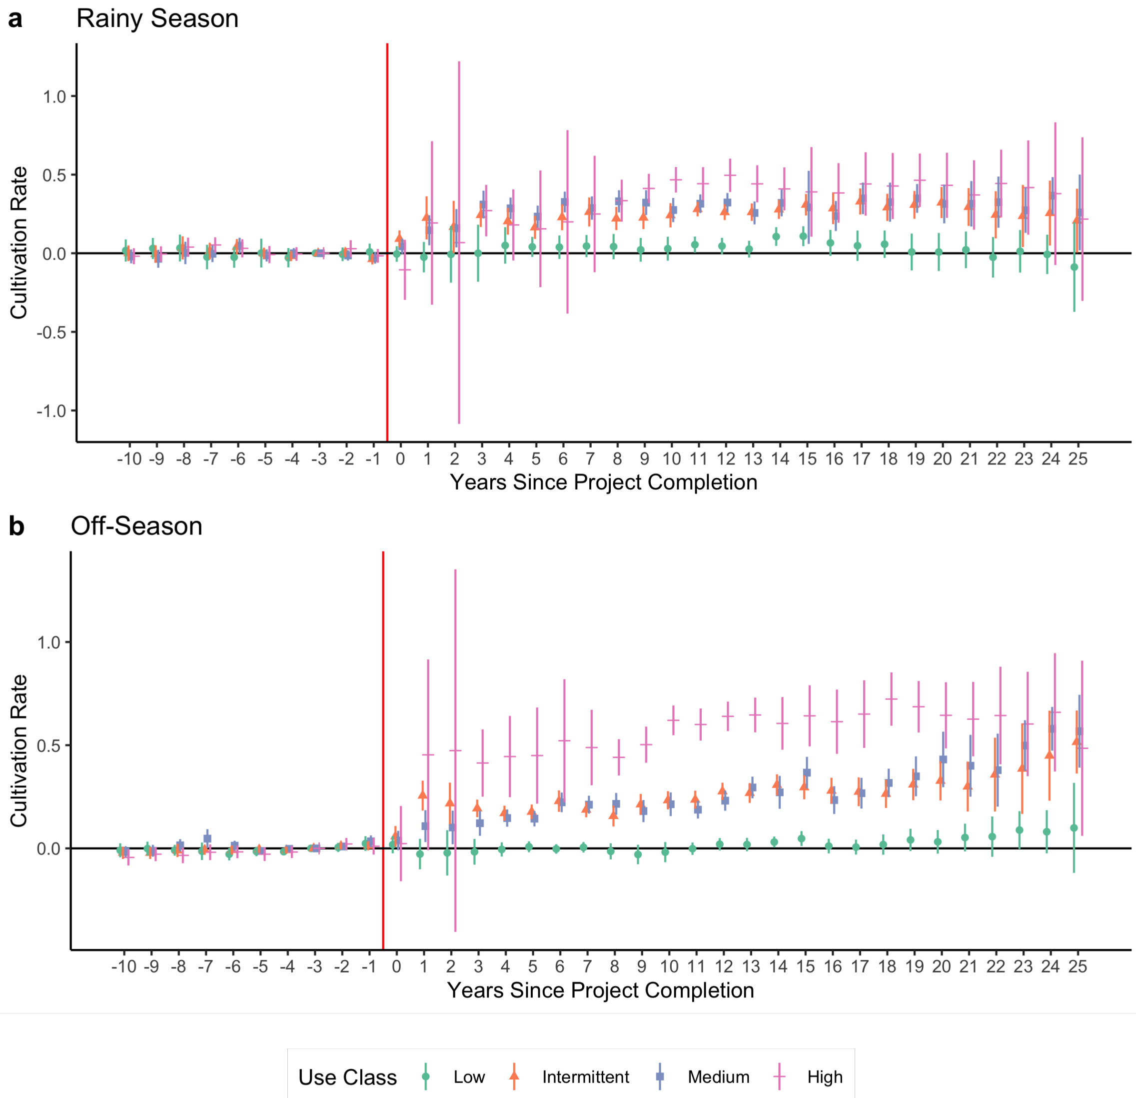

Notes: This figure shows the event-study results on the change in cultivation post project completion, by use class. Panel A shows the event-study results for the rainy season (August-November) and Panel B shows the event study results for the off-season (December-July). "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Figure A17: Rotation by Use Class

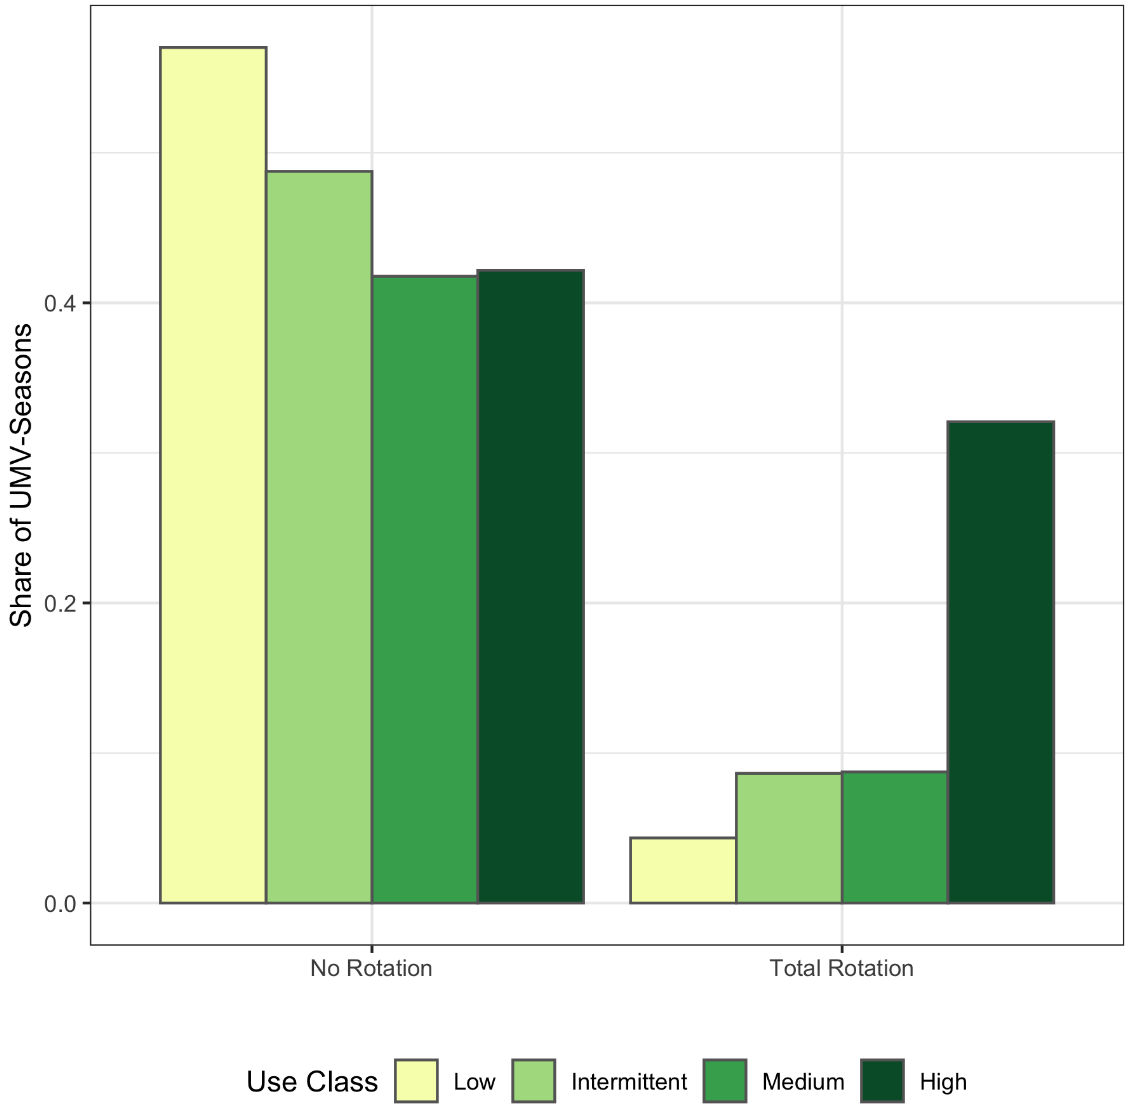

Notes: This figure shows the share of UMV-level land use changes that are either not rotated at all or totally rotated. Rotation is defined as the pixel-level total change in cultivation minus the pixel-level net change in cultivation normalized by the pixel-level total change in cultivation. A value of 1 indicates that an equal amount of land became cultivated and stopped being cultivated between adjacent years. A value of 0 indicates that all change in cultivation was either land becoming newly cultivated or land no longer being cultivated. No Rotation is defined as Rotation < 0.05 and Total rotation is defined as Rotation > 0.95. "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Figure A18: Rotation Histograms by Use Class

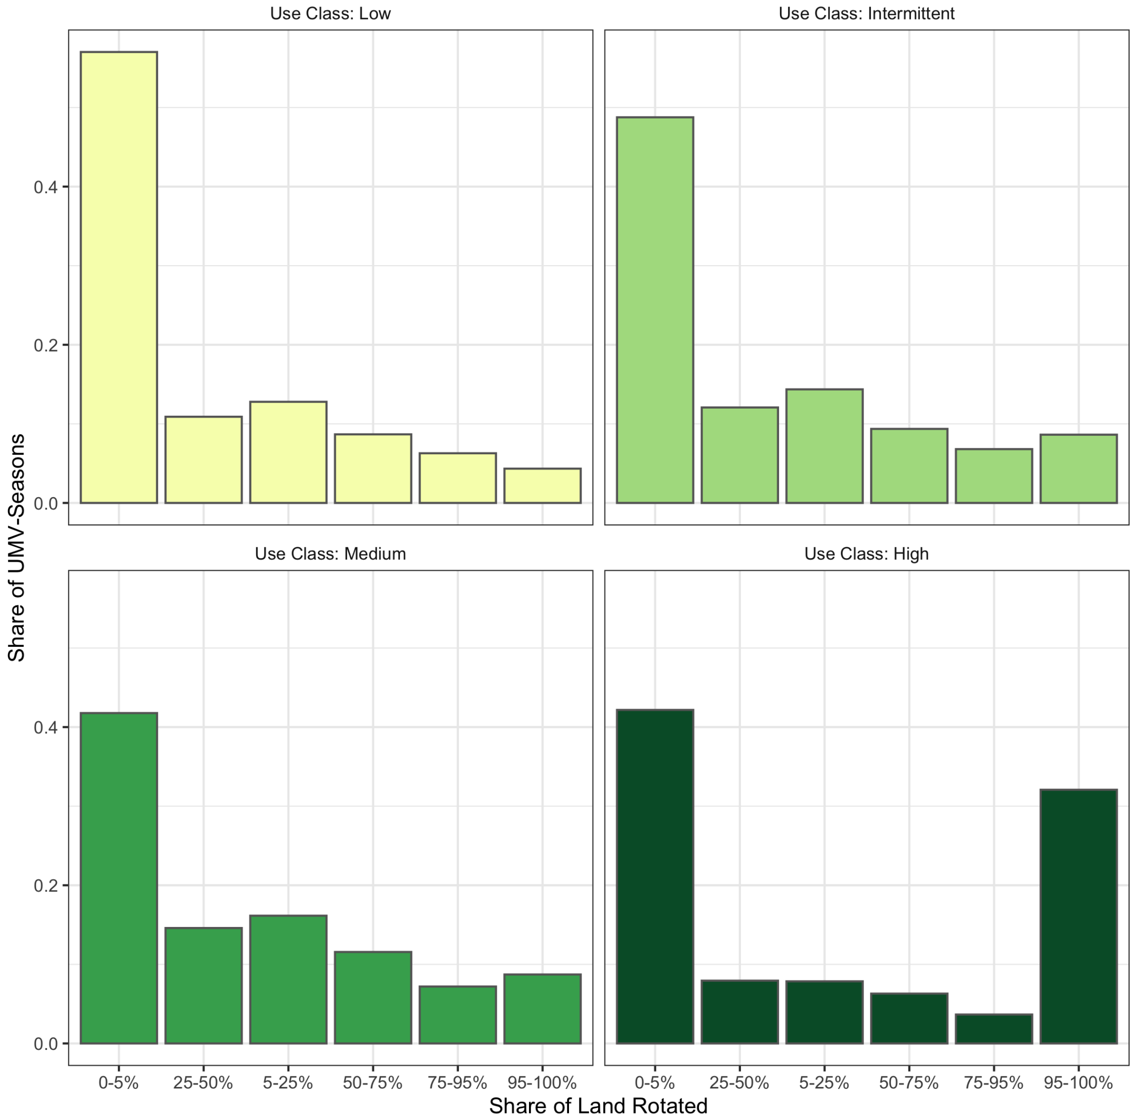

Notes: This figure shows the share of UMV-level land use changes by level of rotation and use class. Rotation is defined as the pixel-level total change in cultivation minus the pixel-level net change in cultivation normalized by the pixel-level total change in cultivation. A value of 1 indicates that an equal amount of land became cultivated and stopped being cultivated between adjacent years. A value of 0 indicates that all change in cultivation was either land becoming newly cultivated or land no longer being cultivated. "Low Use" are UMVs with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Medium Use" are UMVs with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "High Use" are UMVs with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. "Intermittent Use" are UMVs with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Figure A19: Share of Land Cultivated in Completed Projects

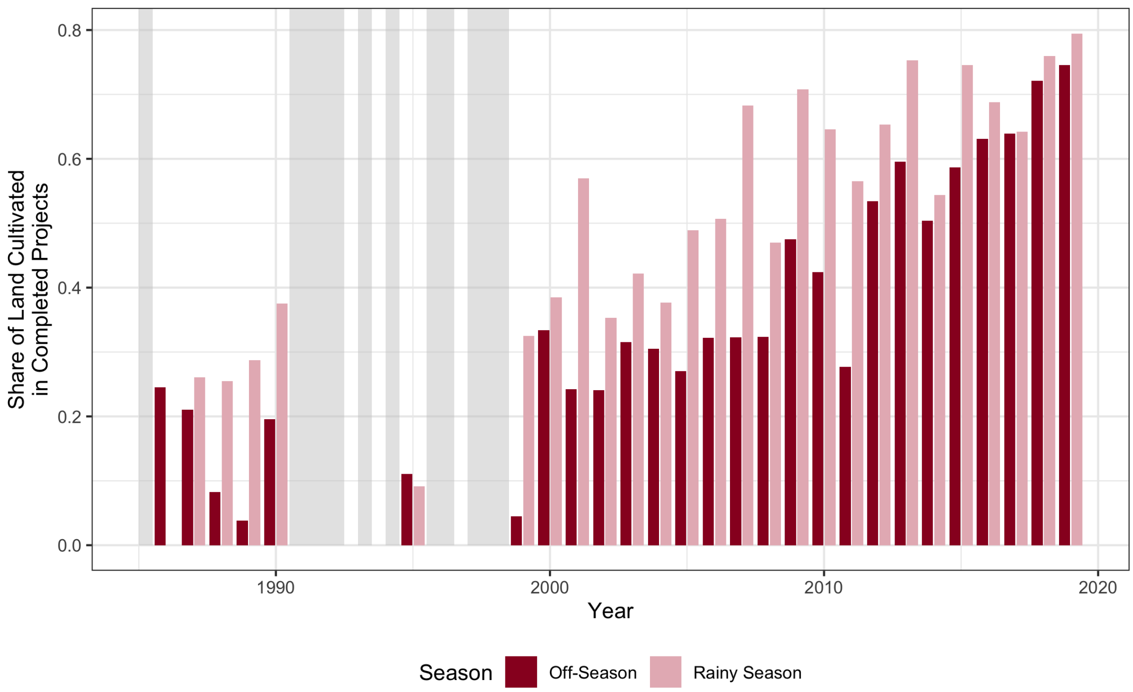

Notes: This figure shows patterns of land use over time in completed irrigation projects in the SRV. The Y-axis denotes the share of area in completed projects that is cultivated by season. Cultivation is defined at the pixel-level as having a maximum NDVI within the season greater than 0.3. The rainy season runs from August to November and the off-season runs from December to July. There is no available data for the years 1991, 1992, 1996, and 1998. Additionally, there is no available Rainy Season data for the years 1985, 1993, 1994, and 1997. Missing data is indicated by grey areas.

Online Appendix Tables

Table A1: Crops by Season and by Delegation

Crop	Crop Calendar			Area ('000s ha)		
	Sowing	Growing days	Harvest	Dagana	Podor	Total
Hot off-season						
Rice[a]	Jan - Mar	110-150	May - June	34.35	8.42	42.77
Groundnuts	Feb - Apr	90-125	June - July	6.75	4.83	11.57
Rainy season						
Rice[a]	Jun - Jul	80-90	Sept - Nov	11.24	4.78	16.02
Millet	Jun - Jul	75-90	Sept - Oct	1.52	1.75	3.27
Maize	Jul	75-90	Sept - Oct	0.94	2.90	3.84
Recession period						
Sorghum	Sept - Oct	80-90	Dec - Jan	0.83	6.50	7.33
Beans	Nov - Dec	75-100	Feb - Mar	1.20	3.40	4.60
Cold off-season						
Onion[b]	Oct - Dec	110-150	Mar - May	2.48	2.01	4.49
Tomato[b]	Oct - Dec	110-150	Mar - May	1.64	–	–
Potatoes[b]	Nov - Jan	75-100	Feb - Apr	2.4	–	–
All year						
Cassava	–	240-365	–	1.39	0.25	1.64
Sugarcane	–	365-380	–	12.00	–	12.00

Notes: This table shows the planting cycles of the main crops present in Dagana and Podor, as well as their importance based on planted area. The hot off-season runs from April to July. The rainy season runs from August to November. The recession period runs from September to February. The cold off-season runs from December to March. The area data are from the Direction de l'Analyse, de la Prevision et des Statistiques Agricoles (DAPSA) of Senegal, unless otherwise specified. (a): Data are for 2018 and are sourced from JICA reports. (b): Data are for 2020 and are sourced from the Centre de Gestion et d'Economie Rurale de la Vallée du Fleuve Senegal (CGERV)'s platform. The crop calendars are from the FAO calendar crop.

Table A2: Comparison of Sites in Analysis Sample vs. Excluded Sample

	Analysis Sample		**Missing Information**		
	Mean	**N**	**Mean**	**N**	**P-Value**
	(1)	(2)	(3)	(4)	(1)-(3)
General characteristics					
Share of private projects (%)	58.68	1377	46.87	2509	0.00
Creation year	1999.60	1377	1993.20	2309	0.00
Number of members	44.33	1059	40.84	1779	0.22
Share of women members (%)	16.39	963	14.73	1622	0.14
Land use					
Shapefile area (ha)	34.01	1377	33.42	1055	0.88
Area initially reported (ha)	24.34	1088	21.12	2323	0.07
Area extended (ha)	5.17	1088	4.40	2323	0.23
Area abandoned (ha)	1.84	1088	2.45	2323	0.07
Area currently exploited (ha)	27.68	1088	23.06	2323	0.02
Area per member (ha)	9.24	1059	3.39	928	0.00
Number of observations	–	1377	–	2568	–

Notes: This table displays the means of select variables, along with the number of non-missing observations and the p-value corresponding to differences in the two means. Columns (1)-(2) represent UMVs included in our analysis and satisfying the following two conditions: (i) the UMV was successfully georeferenced, and (ii) the UMV has a non-missing project construction date after 1988. Columns (3)-(4) correspond to UMVs excluded from our analysis because of one of the following reasons: (i) the UMV has a missing project construction date (N=110), (ii) the UMV was constructed before 1988 (N=945), (iii) the UMV was not successfully georeferenced (N=1513). Column (5) shows the p-value corresponding to the t-test for the difference in means between UMVs in column (1) and columns (3). All projects in the delegations of Podor and Dagana are included in the table.

Online Appendix Figures NOT USED IN TEXT

Figure B1: UMVs by creation year

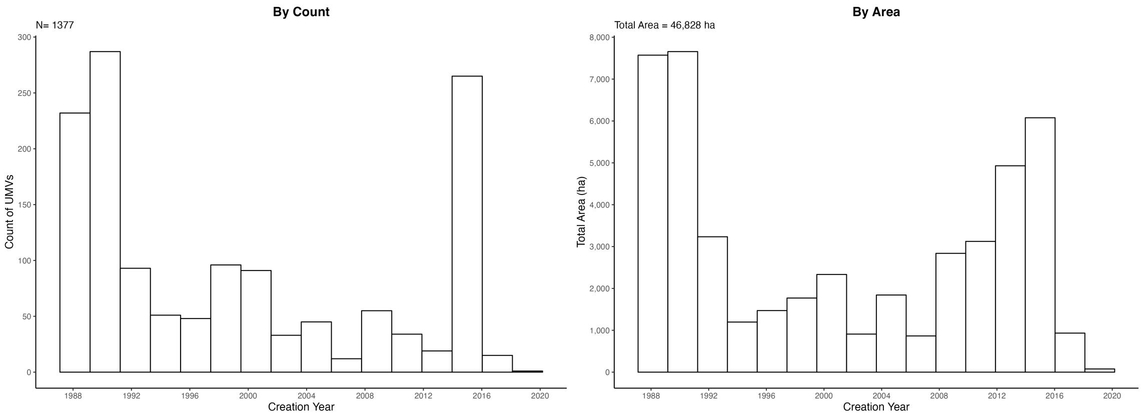

Notes: This figure displays the evolution in the number of the UMVs and their area over time. The x-axis shows the creation year, which corresponds to the first year of operation as recorded by SAED. The left-hand side figure displays the histogram of the number of new UMVs by creation year. The right-hand side figure displays the histogram of the total area of the new UMVs by creation year. The figure includes all UMVs with creation year after 1987 in the delegation of Dagana and Podor (our analysis sample). UMVs with missing creation date or with creation year before 1988 are excluded from the figure.

Figure B2: Visualizing NDVI Changes

Notes: This figure displays examples of UMVs with especially small and large increases in NDVI. The top set shows the five Dagana UMVs with the smallest difference between median April-July maximum NDVI before and after project completion. The bottom set shows the five Dagana UMVs with the largest difference between median April-July maximum NDVI before and after project completion. The UMV of interest is outlined in red and neighboring UMVs are shown in black. Within each set, the top row shows the year with the closest to median April-July maximum NDVI pre-completion and the bottom row shows the year with the closest to median April-July maximum maximum NDVI post-completion. Images are composited with priority given to pixels with greater NDVI.

Figure B3: Effect of Irrigation Project Completion on Cultivation Rate by Year of Completion

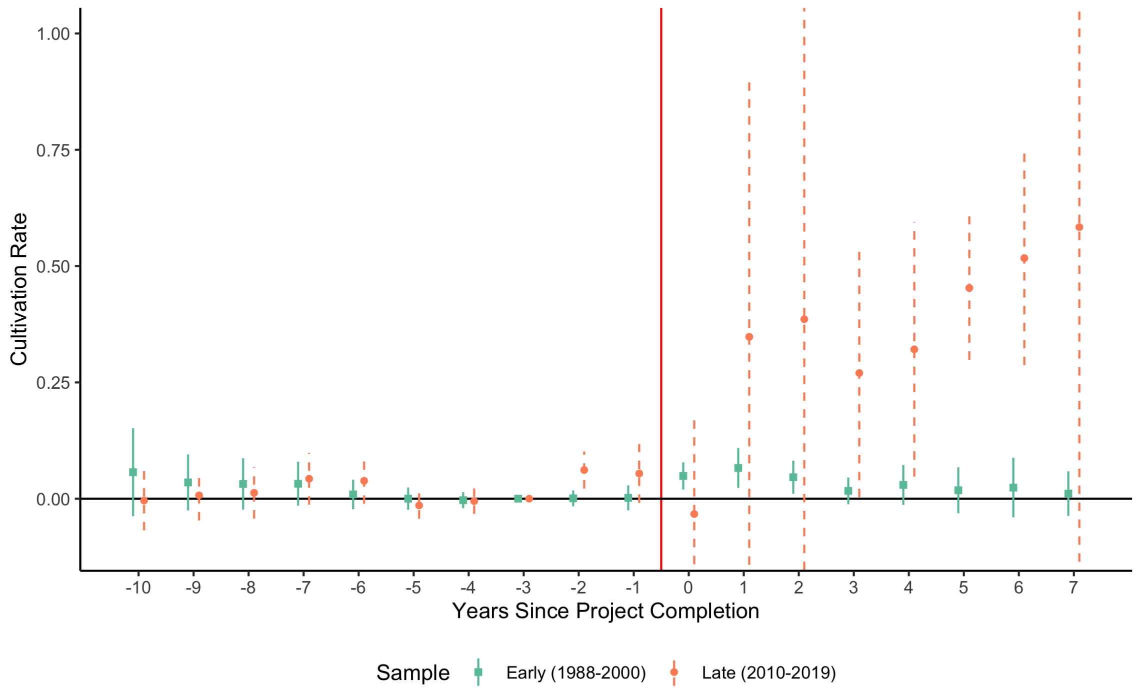

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.25 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the project level over the period December-July to construct a proxy for the share of project land in use. Estimates are shown separately for early projects (completed between 1988 and 2000) and late projects (completed between 2010 and 2019). Average cultivation rates prior to project completion are 0.029 (early projects) and 0.041 (late projects). The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients.

Figure B4: Effect of Irrigation Project Completion on Cultivation Rate by Project Type

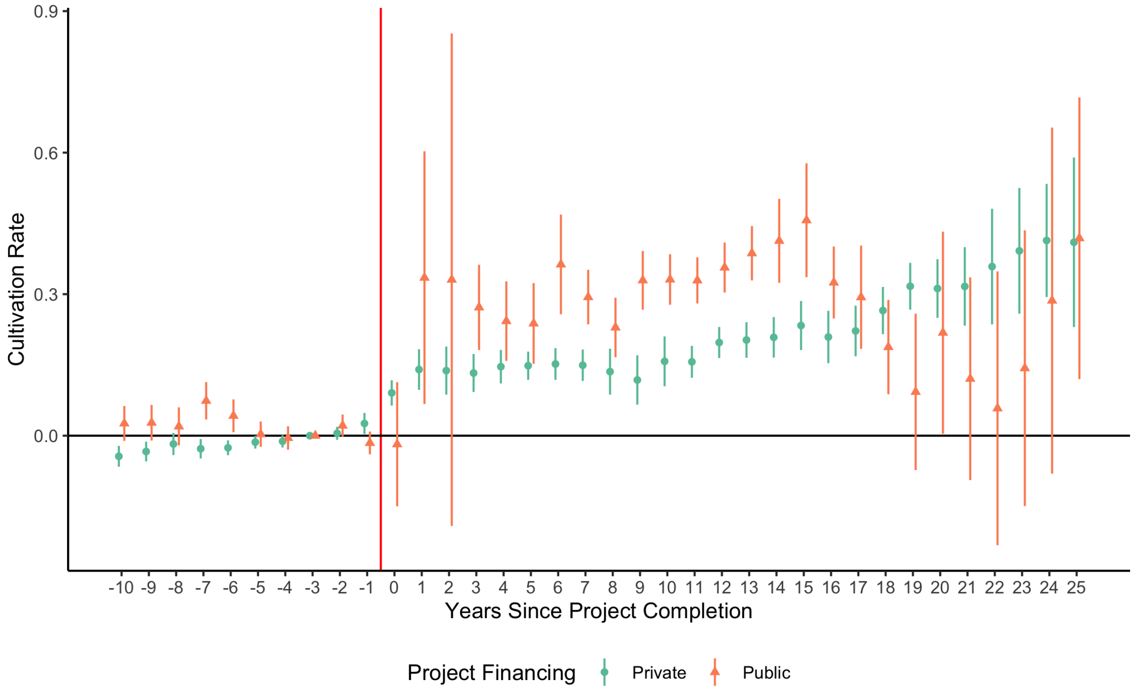

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.25 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the project level over the period December-July to construct a proxy for the share of project land in use. Estimates are shown separately for public and private projects. Estimates for public projects are in red, and those for private projects are in green. Average cultivation rates prior to project completion are 0.037 (public projects) and 0.045 (private projects). The x-axis shows the number of years since the project creation date, which is the first year of operation as recorded by SAED. The red line separates the treatment leads from the treatment lags. The dots are coefficient estimates while the vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients.