# Factor Market Failures and the Adoption of Irrigation in Rwanda*

Maria Jones Florence Kondylis John Loeser
Jeremy Magruder

January 14, 2022

**Abstract**

Factor market failures can limit adoption of profitable technologies. We leverage a plot-level spatial regression discontinuity design in the context of irrigation use by farmers provided free access to water. Using irrigation boosts profits by 43-62%. Yet, farmers only irrigate 30% of plots because of labor costs. We demonstrate inefficient irrigation use, by showing farmers irrigating one plot reduce their irrigation use on other plots. This inefficiency is largest for smaller households and wealthier households, suggesting labor market frictions constrain use of irrigation.

*Jones: Development Impact Evaluation, World Bank, mjones5@worldbank.org; Kondylis: Development Impact Evaluation, World Bank, fkondylis@worldbank.org; Loeser: Development Impact Evaluation, World Bank, jloeser@worldbank.org; Magruder: UC Berkeley, NBER, jmagruder@berkeley.edu. We thank the editor, Esther Duflo, for her invaluable guidance as well as three anonymous referees for their useful feedback. We also thank Abhijit Banerjee, Chris Barrett, Paul Christian, Alain de Janvry, Simeon Djankov, Andrew Foster, Doug Gollin, Saahil Karpe, Arianna Legovini, Elisabeth Sadoulet, John Strauss, Duncan Thomas, Chris Udry, and seminar audiences at Cornell University, Georgetown University, Georgia State University, Harvard University/MIT, Michigan State University, North Carolina State University, Northwestern University, Oxford University, University of California, Berkeley, University of Pittsburgh, University of Southern California, and the World Bank. We thank the European Union, the Global Agriculture and Food Security Program (GAFSP), the World Bank Rwanda Country Management Unit, the World Bank i2i fund, 3ie, and IGC for generous research funding. Emanuele Brancati, Anna Kasimatis, Roshni Khincha, Christophe Ndahimana, and Shardul Oza provided excellent research assistance. Finally, we thank the technical staff at MINAGRI, the staff of the LWH project implementation unit, and the World Bank management and operational teams in Rwanda for being outstanding research partners. We are particularly indebted to Esdras Byiringiro, Jolly Dusabe, Hon. Dr. Gerardine Mukeshimana, and Innocent Musabyimana for sharing their deep knowledge of Rwandan agriculture with us. Magruder acknowledges support from NIFA. This study was approved by UC Berkeley IRB (2016-06-8861), the Rwanda National Ethics Committee, and IPA IRB (10951), and is registered in the AEA RCT Registry (AEARCTR-0001323). The views expressed in this manuscript do not reflect the views of the World Bank. All errors are our own.

# I Introduction

Limited adoption of productive technologies is a prominent explanation of low agricultural productivity in sub-Saharan Africa (World Bank, 2007). Existing productive technologies may be underutilized due to inefficiencies in the markets faced by farming households (Udry, 1997). A recent literature has provided robust evidence that these market failures distort technology adoption, most commonly through experimental manipulation of markets for risk, credit, and information (De Janvry et al., 2017).

Evidence is thinner on the role of constraints to adoption generated by failures in factor markets for land and labor. Land and labor markets are characterized by substantial frictions in developing countries (Fafchamps, 1993; Udry, 1997; LaFave and Thomas, 2016), even where these markets are particularly active (Kaur, 2014; Breza et al., 2018). Economic theory suggests land and labor market failures reduce agricultural productivity by generating inefficient allocations of labor and land across farms (Fei and Ranis, 1961; Benjamin, 1992). More recent empirical work has found that these inefficiencies are quantitatively important (Udry, 1997; Adamopoulos and Restuccia, 2014; Adamopoulos et al., 2017; Foster and Rosenzweig, 2017; Adamopoulos and Restuccia, 2018).

In this paper, we evaluate a potentially transformative technology: irrigation. Irrigation boosts agricultural productivity in several ways: it adds additional agricultural seasons, enables cultivation of water-intensive crops, and reduces production uncertainty. In the Rwandan context of this paper, irrigation use primarily enables high-value horticulture cultivation in the dry season. Absent irrigation, farmers either cultivate lower-value banana or leave their fields fallow in dry seasons.

Surface water irrigation systems like the ones we study require substantial public investment; returns to that investment can only be realized if farmers choose to use irrigation. While irrigation water is free to farmers in these schemes, horticulture is associated with increased use of complementary inputs, including labor, fertilizer, and seeds. Market failures, including in factor markets, therefore have the potential to generate inefficient irrigation adoption as they induce a wedge between shadow prices and market prices of these inputs. We demonstrate that despite apparent high returns, adoption of irrigation is partial, inefficient, and distorted by incomplete land and labor markets. This distortion has direct policy relevance: constraints posed by incomplete factor markets may limit the return on investment for productive infrastructure.

We proceed in 3 steps. First, we establish that irrigation is a productive technology, but adoption is partial. Second, we demonstrate that this partial adoption is inefficient. Third, we show that labor market failures generate constraints to adoption of irrigation.

We begin by estimating the returns to newly-constructed hillside irrigation schemes in Rwanda, making use of a plot-level spatial discontinuity design. We sample plots within 50 meters of scheme boundaries, determined by gravity fed canals which originate from a distant water source and must maintain a consistent gradient along the hillside. We survey 931 cultivators on 1,688 plots for 4 years. We then compare plots just inside the command area, which have access to water for irrigation, to plots just outside the command area, which do not. Treatment on the treated estimates reveal that irrigation enables the transition to dry season cultivation of horticulture. While we find no effects on rainy season yields, labor, or inputs, dry season estimates correspond to 43% - 62% growth in annual cash profits. To our knowledge, this is the first study to use a natural experiment to estimate the returns to irrigation in sub-Saharan Africa; our estimate is almost identical to an estimate from Duflo and Pande (2007) in India.[1] Despite the large effects we estimate, adoption is low: only 30% of plots are irrigated 4 years after canals became operational. At this level of adoption, the sustainability of hillside irrigation systems is in doubt: even the large gains in cash profits to adopters are unable to generate enough surplus to pay for routine maintenance costs.

We investigate the effect of irrigation on inputs to shed light on what might determine farmers' decisions to adopt irrigation. In this context, the dominant input associated with irrigation is households' own labor. The shadow wage that prices household labor is notoriously difficult to value, but if this labor were valued at the market wage, estimated effects on household labor would be 6 times as large as estimated effects on expenditures on hired labor and other inputs, and estimated effects on profits would fall from 43% - 62% to -9% - 2%. Valuing household labor at the market wage may not be appropriate: rural market wages are likely to be inefficiently

[1]Existing work that estimates the returns to irrigation using natural experiments is predominantly from groundwater irrigation in South Asia, leveraging variation in slope characteristics of river basins (Duflo and Pande, 2007), aquifer characteristics (Sekhri, 2014), or well-failures (Jacoby, 2017) for identification. Estimates of the return to irrigation in Africa include Dillon (2011), who estimates the returns to irrigation using propensity score matching in Mali. More broadly, Dillon and Fishman (2019) review the literature on the impacts of surface water irrigation infrastructure.

high in developing countries (Kaur, 2014; Breza et al., 2018), and labor market failures in rural areas may generate heterogeneity in the shadow wage (Singh et al., 1986; Benjamin, 1992; LaFave and Thomas, 2016). Heterogeneity in the shadow wage would then cause inefficient adoption of irrigation across households.[2] Alternatively, these results could also be consistent with unconstrained profit maximization if farmers have heterogeneous returns to or costs of adopting irrigation (Suri, 2011) and optimize at market wages.

We build on seminal agricultural household models (Singh et al., 1986; Benjamin, 1992) to derive a test for inefficient adoption of irrigation in the spirit of Udry (1997). We model households' production and consumption decisions in the presence of uncertainty, plot-level heterogeneity, and failures in insurance, credit, and labor markets. Consistent with our reduced form results, we model access to irrigation as a labor- and input-complementing increase in plot-level productivity. Our test is based on the efficient separability of production decisions across plots. With complete markets, farmers maximize profits on each plot and access to irrigation on one plot does not affect production decisions on other plots. In contrast, when there are failures in land and other markets, access to irrigation on one plot causes farmers to reduce labor and input use on other plots. This test is joint for the null of frictionless land markets: if land markets are frictionless, then markets should reallocate land to farmers who can cultivate most profitably.

We implement our test for inefficient adoption, exploiting the plot-level discontinuity in access to irrigation. We test whether farmers who have a plot just inside the command area reduce their input use on their other plots compared to farmers who have a plot just outside the command area. We find large substitution effects, strongly rejecting complete markets: an additional irrigated plot caused by access to irrigation is associated with a 64 - 70 percentage point decrease in the probability of irrigating a second command area plot. We find similarly large effects for adoption of horticulture, household labor, and inputs. These results confirm a simple descriptive analysis, which shows that few households are able to irrigate more than one command area plot. Applying these results, a simple back-of-the-envelope calculation implies that, absent this substitution, adoption of irrigation would be at least 34% higher. Moreover, the presence of this substitution implies current adoption of irrigation is

[2] This heterogeneity could only exist if there were frictions in at least one other market in addition to labor markets.

inefficient: different households make different adoption decisions on technologically identical plots because of their access to irrigation on their other plots.[3]

The previous test shows that inefficient adoption of irrigation is caused by failures of land markets, and at least one other market; however, it does not establish which other market fails. We produce two tests that suggest that labor market constraints, as opposed to financial constraints, bind in our context.

First, we extend the model and propose a test for whether labor market frictions contribute to inefficient adoption in this context. We consider the effects of household size and wealth on input substitution across plots in the presence of insurance, credit, and labor market failures. We demonstrate that only labor market failures can explain irrigation access on one plot leading to greater input substitution across plots for richer households, and decreased input substitution across plots for larger households. We then estimate differential substitution with respect to household size and wealth to test for labor market failures. We find exactly this pattern: households with two additional members substitute 37% - 78% less than average size households, while one standard deviation wealthier households substitute 45% - 90% more than average wealth households. As these patterns of differential substitution can only be explained by labor market failures, and not credit or insurance market failures, these results imply that labor market failures cause substitution and contribute to inefficient adoption of irrigation.

We then complement this result with experimental evidence. We conduct three randomized controlled trials with the farmers who have access to irrigation. Two of these trials focus on characteristics peculiar to irrigation systems: irrigation scheme land taxes and failures of operations and maintenance; we find neither plausibly affects farmers' adoption decisions in our context. In the third experiment, we distribute minikits which contain all necessary inputs for horticulture cultivation to randomly selected farmers. Previous work has shown providing free minikits targets credit, risk, and information constraints: it reduces costs of growing horticulture under irrigation, basis risk, and costs of experimentation, respectively (Emerick et al., 2016; Jones et

[3]With sufficient time, these sites could reach an equilibrium in which this misallocation would have slowly been corrected by markets (Gollin and Udry, 2019). However, we note that our results are 4 years after initial access to irrigation, and we do not observe dynamics after 2 years. This is sufficient for our results to have meaningful implications for the long run sustainability of these schemes. Our results also complement evidence from the United States which suggests that initial allocations can persist for many decades even with seemingly well functioning land markets (Bleakley and Ferrie, 2014; Smith, 2019).

al., 2021). We find no effects of receiving minikits on adoption of horticulture in our context, in contrast to existing work, and conclude that financial and informational constraints are unlikely to be a primary explanation for low and inefficient adoption of irrigation.

This paper demonstrates that there are large revenue returns to hillside irrigation in Rwanda, but that frictions in land and labor markets generate inefficient adoption of the technology. Our context allows us to demonstrate that separation failures induce differential adoption of irrigation on technologically identical plots. In doing so, we expand on tests from classical agricultural household models that investigate factor market frictions (Singh et al., 1986; Benjamin, 1992; LaFave and Thomas, 2016; Dillon and Barrett, 2017; Dillon et al., 2019) to provide ground-level evidence on the mechanisms behind misallocation (Adamopoulos and Restuccia, 2014; Adamopoulos et al., 2017; Foster and Rosenzweig, 2017; Adamopoulos and Restuccia, 2018). Adoption of irrigation in Rwanda is constrained by factor market frictions, even when water is provided for free; this conclusion suggests the urgent need for future research in the role of incomplete factor markets in technology adoption.

This paper is organized as follows. Section II describes the study context and our sources of data. Section III presents our estimates of the impacts of irrigation in Rwanda. Section IV presents our model of adoption of irrigation in the presence of market failures. We implement tests of constraints to adoption and labor market failures suggested by the model in Section V, and experimental tests in Section VI. Section VII concludes.

# II Data and context

## A Irrigation in Rwanda

We study 3 hillside irrigation schemes located in Karongi and Nyanza districts of Rwanda, constructed in 2014 under the auspices of a flagship government project; a timeline of construction and our surveys is presented in Figure 1. Rainfed irrigation in and around these sites is seasonal, with three potential seasons per year. During the main rainy season ("Rainy 1"; September - January), rainfall is sufficient for production in most years. In the second rainy season ("Rainy 2"; February - May), rainfall is sufficient in an average year but insufficient in dry years. In the dry season

("Dry"; June - August), rainfall is insufficient for agricultural production for seasonal crops. Absent irrigation, agricultural production in these sites consists of a mix of staples (primarily maize and beans) which are cultivated seasonally and primarily consumed by the cultivator, as well as bananas which are sold commercially;[4] most farmers adopt either a rotation of staples, fallowing land in the dry season, or cultivate bananas.

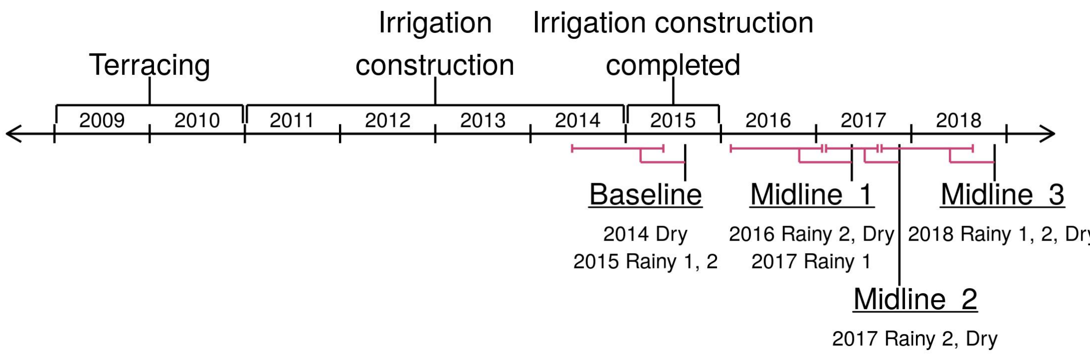


Figure 1: Timeline

*Notes:* Black lines indicate when (or the period during which) events took place, while pink lines are used to indicate survey recall periods.

Irrigation in these schemes is expected to increase yields by reducing risk in the second rainy season and enabling cultivation in the short dry season. As the dry season is relatively short, cultivating the primary staple crops is not possible for households that cultivate during the two rainy seasons, even with irrigation. Instead, irrigation enables the cultivation of short-cycle horticulture during the dry season. Horticulture production (most commonly eggplant, cabbage, carrots, tomatoes, and onions) can be sold at local markets where it is both consumed locally and traded for consumption in Kigali. Since horticultural production is relatively uncommon during the dry season in Rwanda due to limited availability of irrigation, finding buyers for these crops is relatively easy during this time. At baseline 3.2% of plots outside of the command area are planted with at least some horticulture, primarily during the rainy seasons.

[4]Staple rotations also include smaller amounts of sorghum and tubers, while there is also some cultivation of the perennial cassava, along with other minor crops. In our data, maize, beans, or bananas are the main crop for 85% of observations excluding horticulture.

The three schemes we study were constructed by the government from 2009 - 2014, with water beginning to flow to some parts of the schemes in 2014 Dry and becoming fully operational by 2015 Rainy 1 (August 2014 - January 2015). In each site, land was terraced in preparation for the irrigation works (as hillside irrigation would be infeasible on non-terraced land) and the LWH approach also includes soil fertility interventions across all parts of the site. Construction and rehabilitation of terraces in these sites began in 2009 - 2010.[5] The schemes are all gravity fed, and use surface water as the source. From these water sources, a main canal was constructed along a contour of the hillside; engineering specifications require the canal to be sufficiently steep so as to allow water to flow, but sufficiently gradual to control the speed of the flow, preventing manipulation of the path of the canal.

In all sites, sufficient water is available to enable irrigation year-round. Underground secondary pipes run down the terraces from the canal every 200 meters, with valves on the main canal controlling the flow of water into a set of secondary pipes. Farmers draw water from tertiary valves on these pipes located on every third terrace, from which flexible hoses and dug furrows enable irrigation on all plots below the canal. The "command area" for these schemes, the land that receives access to irrigation, is made up of all the plots which are below the canal and located within 100 meters of one of these tertiary valves. We define clusters of plots that share the same secondary pipe as Water User Groups. Water User Groups can be grouped into Zones, where each Zone defines the set of secondary pipes controlled by a single valve on the canal. To the extent that there is heterogeneity in plot-level water pressure, the plots nearest to the canal face the lowest pressure, as the pressure available at an individual valve is determined by the volume of water in the pipe above that valve.

Farmers who cultivate land in the command area receive access to the irrigation infrastructure for free. The primary cost to farmers of irrigating a plot in this context is the labor associated with the actual irrigation, including maintaining the dug furrows and using the hoses to apply water from the valves to their plots. At the time of the study, only 4% of farmers paid land taxes to cover operations and maintenance costs, and there were no significant challenges in operations and maintenance during

[5] Construction of these terraces would have disturbed perennial cultivation, including banana root structures. Banana pseudostems or "trees" develop from root structures and bear a single crop of fruit over about 9-12 months. As new pseudostems can be transplanted, banana cultivation is itself an annual decision and banana cultivation decisions observed in our follow up surveys are all made several cropping cycles and adoption decisions after terracing.

the sample period, as we document in Section VI. Hence, since all plots in the command area have access to irrigation, analysis will focus on reported use ("adoption") of irrigation.

We exploit a spatial discontinuity in irrigation coverage to estimate the impacts of irrigation. Because the main canals must conform to prescribed slopes relative to a distant and originally inaccessible water source, the geologic accident of altitude relative to this source determines which plots will and will not receive access to irrigation water. Hence, before construction, plots just above the canal should be similar to plots just below the canal and, importantly, should be managed by similar farmers. Following construction, however, the plots just below the canal fall inside the command area and have access to irrigation, while the plots just above the canal fall outside the command area and do not have access to irrigation. Appendix Figure A1 presents a photograph of one of these hillsides. In that picture, the contour with the canal is easily identified by the sharp change in crops cultivated above versus below the canal.

## B Data

### 1 Spatial sampling

To take advantage of the spatial discontinuity in access generated by the command area boundary, we randomly sampled plots in close proximity to this discontinuity. In practice, we dropped a uniform grid of points across the site at 2-meter resolution, and then randomly sampled points from the grid, oversampling 65% of sampled points within 50m of the command area boundary to ensure density at the discontinuity.[6] In all three irrigation sites, we additionally sampled some points further from the canal inside the command area. We use these points primarily to examine experimental treatments described in Section VI. After each point was sampled, we excluded all points within 10m of that point (to avoid selecting multiple points too close together). Enumerators visited each of these sampled points, and retained the 80.1% of points which fell on cultivable land (that is, land which was not covered in bodies of water, forests, swamps, or other dense foliage). With the help of a key informant (often the village leader), they then recorded the name of the cultivator of the plot, their

[6] This procedure yields a sample of plots that is more representative of land than of households. Analysis in Section III is qualitatively similar when weighting households inversely proportionally to their number of plots (not reported).

contact information, as well as a detailed description of the plot. These distinct listed cultivator households form our main household survey sample. For each cultivator, one of their identified plots is randomly selected, which we refer to as the *sample plot*.

We use three subsamples for our main analysis in this paper (Jones et al., 2022). First, our analysis of impacts of access to irrigation in Section III leverages the spatial discontinuity in access to irrigation, and we make two restrictions to construct the subsample for this analysis. Only two of the three sites have a viable boundary of cultivable land both just inside and just outside the command area, and we therefore restrict to sample plots in these two sites. Second, we restrict to sample plots that are located within 50m of the command area boundary. Throughout the paper, we refer to this sample of plots (and associated households) as the *discontinuity sample*. Second, for our analysis of impacts of sample plot access to irrigation on production decisions on other plots in Section V, we use the discontinuity sample, but we must also restrict to households with data on another plot (which we refer to as the "LOP", and discuss in Section 2). Throughout the paper, we continue to refer to this sample of plots (and associated households) as the discontinuity sample; when it is useful to disambiguate, we refer to this sample as the "LOP sample". Third, for our analysis of experimental impacts in Section VI, we restrict to sample plots (and associated households) located in the command area in all three sites.

## 2 Survey

Our baseline survey covered 1,695 spatially sampled cultivators in August - October 2015 (Jones et al., 2022). The survey includes detailed agricultural production data (season-by-season) for seasons 2014 Dry through 2015 Rainy 2 (June 2014 - May 2015). The dates of this survey and follow up surveys, along with the agricultural seasons they cover, are presented in Figure 1. This is not a "true" baseline as some farmers had already gained access to irrigation in 2014 Dry; we therefore complement our survey data with measures of vegetation (specifically, NDVI) from satellite images from 1999 through 2018, covering 10 years before and 4 years after the construction of the irrigation schemes. In addition, relatively small parts of the site had access to irrigation as of 2014 Dry; in Appendix Figure A2 we demonstrate 2014 Dry adoption of irrigation is less than 25% of adoption in subsequent dry seasons, and in Section 1 we show balance across the command area boundary in household and plot characteristics both in our survey data and in pre-construction satellite images. A panel of plot-level production and input data are maintained for two plots, which were mapped using GPS devices for precise location and area measurement. The two plots on which panel data is collected represent the primary data for analysis; they include the sample plot (described above) and the farmer's largest other plot excluding the sample plot (defined at baseline; we refer to this as the "largest other plot" or "LOP"). To minimize attrition, we maintain these panels independently for each plot, so missing production data on one plot does not affect collection of production data on the other plot. We also collect data on household characteristics, labor force allocation, and a short consumption and food security module. In our analysis, we will focus on the sample plots to learn about the effects of the irrigation itself, and the largest other plot to learn about how access to irrigation on the sample plot impacts households' productive decisions on their other plots.

Three follow up household surveys were conducted in May - June 2017, November - December 2017, and November 2018 - February 2019 (Jones et al., 2022). In each survey, we ask for up to a year of recall data on agricultural production; based on the timing of our surveys we therefore have production for all agricultural seasons from June 2014 through August 2018, with the exception of 2015 Dry (June - August 2015) and 2016 Rainy 1 (September 2015 - February 2016).

The household sample for the follow up surveys consists of all the baseline respondents, while the plot sample for the follow up surveys consists of the sample plots and largest other plots. To maintain a panel of plots, each follow up survey round includes a "tracking survey". This survey is triggered whenever a household's sample plot or largest other plot was sold or rented out to another household, or a household stopped renting in that plot if it was not the owner ("transacted"). Specifically, we track and interview the new household responsible for cultivation decisions on that plot to record information about cultivation and production, along with household characteristics when the new household was not already in our baseline sample. Data from this tracking survey is incorporated in all our plot level analysis, limiting plot attrition.

Attrition in our survey is low, and details on attrition are presented in Appendix Table A5. Only 5.8% (6.1%) of plot-by-season observations for sample plots outside the command area in our primary analysis sample (defined in Section A) are missing during the dry season (rainy season). There are three sources of attrition: household attrition, plots transacted to other farmers that we were not successful in tracking, and plots rented out to commercial farmers who were based in the capital or internationally (from whom we were unable to collect agricultural production data). We do not find evidence of differential attrition of sample plots due to household attrition or plots transacted to other farmers that we did not track, however we do find access to irrigation causes an additional 6.3% - 9.1% of plots to be rented out to a commercial farmer. We interpret the lack of data on these plots as biasing our primary estimates of the impacts of irrigation downwards, as these plots are cultivated with productive export crops.[7] We note that we continue to collect production data on the farmer's largest other plot when sample plots are rented out, including to a commercial farmer.

## C Summary statistics on agricultural production

To motivate our analysis of the impacts of hillside irrigation, we first introduce some descriptives about irrigation and agricultural production in this context. Table 1 presents summary statistics for agricultural production from our four years of data, pooled across seasons.

**Descriptive 1.** *Irrigation in Rwanda is primarily used to cultivate horticulture in the dry season.*

Farmers in our data rarely irrigate their plots in the rainy seasons, and almost never use irrigation when cultivating staples or bananas (only 2% of plots cultivated with staples or bananas use irrigation in our data). In contrast, 92% of plots cultivated with horticulture in the dry season use irrigation.

**Descriptive 2.** *Horticultural production is more input intensive than staple cultivation, which in turn is (much) more input intensive than banana cultivation.*

The mean horticultural plot uses about 440 days/ha of household labor, 65 days/ha of hired labor, and 50,000 RwF/ha of inputs, regardless of the season in which it is planted. This contrasts with staple plots (250 days/ha of household labor, 40 days/ha of hired labor, 20,000 - 40,000 RwF/ha of inputs), and bananas (100 days/ha of household labor, 10 days/ha of hired labor, 3,000 RwF/ha of inputs).

[7] When sample plots rented out to commercial farmers, we identify the managing household as the household renting out the plot to the commercial farmer to avoid differential attrition. In addition, in results not reported, we demonstrate that impacts of irrigation on profits are unaffected by imputing profits on plots rented out to a commercial farmer as rent paid by these commercial farmers.

Table 1: Summary statistics on agricultural production

| | Staples | | | | Horticulture | | |
|---|---|---|---|---|---|---|---|
| | Staples (1) | Maize (2) | Beans (3) | Bananas (4) | All (5) | Rainy (6) | Dry (7) |
| Yield | 304 | 314 | 287 | 253 | 587 | 580 | 592 |
| Any hired labor | 0.32 | 0.32 | 0.33 | 0.17 | 0.38 | 0.41 | 0.36 |
| Hired labor (days) | 39 | 39 | 40 | 11 | 65 | 67 | 63 |
| Hired labor expenditures | 29 | 28 | 29 | 9 | 48 | 50 | 46 |
| HH labor (days) | 288 | 290 | 275 | 104 | 437 | 425 | 446 |
| Inputs | 19 | 36 | 17 | 3 | 49 | 48 | 49 |
| Profits | | | | | | | |
| Shadow wage = 0 RwF/day | 256 | 250 | 241 | 241 | 491 | 482 | 498 |
| Shadow wage = 480 RwF/day | 118 | 111 | 109 | 191 | 281 | 278 | 283 |
| Shadow wage = 800 RwF/day | 26 | 18 | 21 | 158 | 141 | 142 | 140 |
| Sales share | 0.16 | 0.19 | 0.13 | 0.50 | 0.61 | 0.59 | 0.63 |
| Irrigated | 0.02 | 0.03 | 0.02 | 0.02 | 0.62 | 0.23 | 0.92 |
| Days irrigated | 0 | 1 | 0 | 1 | 21 | 6 | 32 |
| Rainy | 0.99 | 1.00 | 1.00 | 0.49 | 0.43 | 1.00 | 0.00 |
| log GPS area | -2.53 | -2.44 | -2.53 | -2.08 | -2.79 | -2.87 | -2.73 |
| Share of obs. | 0.66 | 0.10 | 0.45 | 0.18 | 0.13 | 0.06 | 0.07 |
| # of obs. | 11,794 | 1,783 | 8,009 | 3,131 | 2,263 | 982 | 1,281 |

*Notes:* Sample averages of outcomes by crop across plot-crop-season observations are presented in this table. Yield, inputs, hired labor expenditures, and profits are reported in units of '000 RwF/ha, labor variables are reported in units of person-days/ha, and log area is in units of log hectares. All other variables are shares or indicators. For reference, the median wage in our data is 800 RwF/person-day and the exchange rate was approximately 800 RwF = 1 USD over our study period.

**Descriptive 3.** *Horticultural production generates much higher cash profits than other forms of agriculture.*

Horticultural production generates much higher cash profits (defined as yields net of expenditures on inputs and hired labor) than other forms of agricultural production in and around these sites. Plots planted to horticulture yield about 500,000 RwF/ha in cash profits, in both rainy and dry seasons. This contrasts with about 250,000 RwF/ha of cash profits producing either staples or bananas.

**Descriptive 4.** *While there are active markets in land rentals and labor, the large majority of farmers cultivate their own land using primarily their own labor.*

At baseline 89% of plots are cultivated by the owner. Across crops 86% - 90% of labor is supplied by the household, and only about 1/3 of plots use any hired labor at all. These two observations could be indicative of efficiencies in household production or of frictions in factor markets, and suggest that agricultural household models (Singh et al., 1986; Benjamin, 1992; LaFave and Thomas, 2016) form a natural framework to consider on-farm decision-making.

**Descriptive 5.** *Household labor is the primary input to production of any crop, and the economic profitability of horticulture depends critically on the shadow wage.*

Pricing on-farm household labor is notoriously difficult. If households are constrained in the quantity of labor they are able to sell on the labor market, they may work within the household at a marginal product of labor well below the market wage. Here, we see that if we value household labor allocated to horticulture at market wages, then cultivating horticulture appears less profitable than cultivating bananas. As a result, ultimately the economic profitability of horticulture relative to bananas will depend critically on the constraints on household labor supply decisions. We note that both horticulture and bananas appear more profitable than cultivating staples, which would be unprofitable if labor were valued at market wages; the ubiquity of staple cultivation in and around these sites is a first piece of evidence that farmers face a shadow wage below the market wage.[8]

# III Impacts of irrigation

## A Empirical strategy

We identify the impacts of irrigation leveraging a spatial discontinuity in access to irrigation. In Figure 2a, we visualize the link between our collection of plot maps and the discontinuity, plotting sample plots within 50 meters of one segment of the command area boundary, within which plots have access to irrigation. Plots to the left are just above the canal that runs through the middle of the figure, and as such do not have access to irrigation, while plots to the right are just below the canal, and have access to irrigation. In Figure 2b, we illustrate how we link this data to agricultural production data and establish our "first stage". We plot the average fraction of

[8]Both horticulture and bananas are also primarily commercial crops, unlike staples. Farmers may place higher value on staples if consumer prices are higher than producer prices (Key et al., 2000), or if there is price risk in production and consumption, both of which may contribute to cultivation decisions as well.

plots that are irrigated during the dry season, after the irrigation schemes became fully operational, within 10 meter bins of distance to the command area boundary – across the boundary, irrigation jumps from 7% to 24%, a 17pp increase. Under the assumption that plots just inside and just outside the command area boundary would have been similar absent irrigation, comparing outcomes on these plots provides an estimate of the causal impact of irrigation.

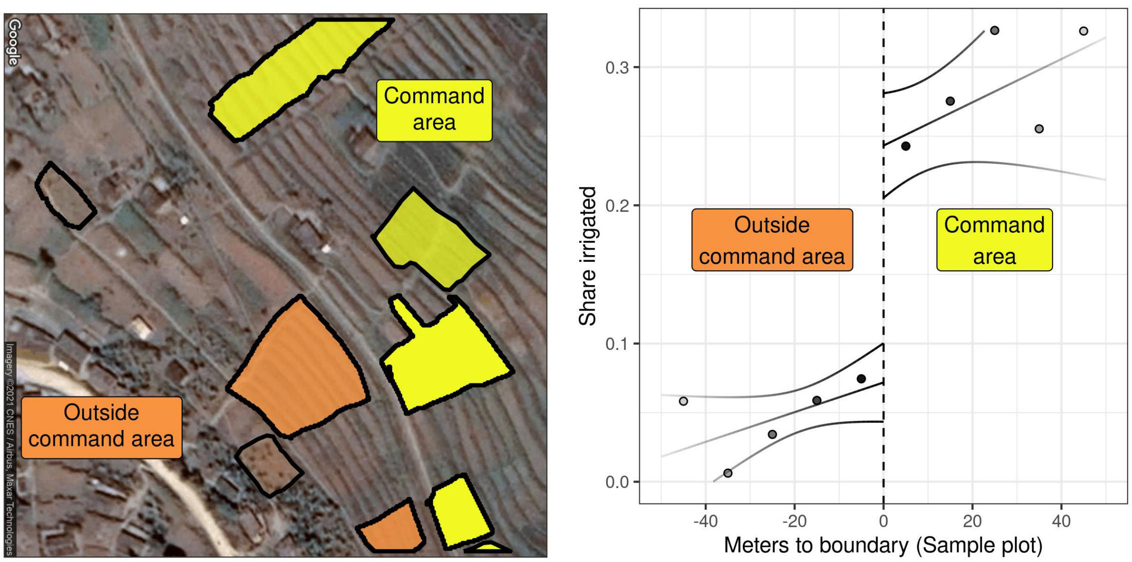


(a) Plots inside and outside command area (b) Sharp increase in irrigation at boundary

Figure 2: Estimating the impact of irrigation exploiting spatial discontinuity in access

We start our analysis with a regression discontinuity design. We restrict this and subsequent analysis to the sample plots within 50 meters of the discontinuity, consistent with our sampling strategy. We regress

$$y_{ist}^{\mathrm{SP}} = \beta_1 \mathrm{CA}_{is}^{\mathrm{SP}} + \beta_2 \mathrm{Dist}_{is}^{\mathrm{SP}} + \beta_3 \mathrm{Dist}_{is}^{\mathrm{SP}} * \mathrm{CA}_{is}^{\mathrm{SP}} + \alpha_{st} + \gamma X_{is}^{\mathrm{SP}} + \epsilon_{ist}^{\mathrm{SP}} \tag{1}$$

where $y_{ist}^k$ is outcome $y$ for plot $k$ of household $i$ located in site $s$ in season $t$, $\mathrm{CA}_{is}^k$ is an indicator for plot $k$ being in the command area, and $\alpha_{st}$ are site-by-season fixed effects meant to control for any differences across sites (including market access or prices). We use $k$ = SP to indicate the household's sample plot, as opposed to $k$ = LOP to indicate the household's largest other plot. $\mathrm{Dist}_{is}^{\mathrm{SP}}$ is the distance of the sample plot from the command area boundary (positive for plots within the command area, negative for plots outside the command area) and $X_{is}^{\mathrm{SP}}$ is the log plot area. We calculate distance using the distance of the plot boundary to the command area boundary. We estimate impacts separately during the rainy seasons and during the dry season; Appendix Figure A2 shows that adoption of irrigation remains stable across our three years of follow-up surveys within rainy or dry seasons, and therefore pool across years. Our coefficient of interest is $\beta_1$, the effect of the command area on outcome $y$.

We include controls for distance to the boundary and log plot area to address two primary potential sources of omitted variable bias. First, the canal sits at a particular contour of the hillside. Plots that are positioned relatively higher on the hillside may have different agronomic characteristics; accordingly, farmers may differentially sort into these plots. We therefore follow convention by controlling for the running variable ($\text{Dist}_{is}^{\text{SP}}$) and its interaction with treatment ($\text{Dist}_{is}^{\text{SP}} * \text{CA}_{is}^{\text{SP}}$).[9] Second, as the construction of the canal slices through plots on the hillside, this may differentially change the area of plots that are positioned above or below the canal. For example, roads are more often located higher on the hillside, leaving less room for plots to extend above the canal relative to below the canal. We anticipate this will cause plots to be relatively larger just inside the command area. As plots exhibit strong evidence of diminishing returns to scale in this context, this effect would likely bias $\beta_1$ downwards absent control.

Next, we consider additional concerns related to selection into our sample caused by access to irrigation. This may arise for two reasons. First, during the construction of the hillside irrigation schemes, forest was deliberately preserved or planted just outside of the command area in order to protect the new investment from erosion. As these forested plots are not agricultural, they are not included in our sampling strategy. This amounts to selection out of our sample of low productivity plots outside the command area, which would bias $\beta_1$ downwards. Second, marginal plots which would have been too unproductive to cultivate absent irrigation, and would

[9] A related concern is that command area plots will mechanically sit at lower elevation than non-command area plots, even after controlling for distance to the canal. Effectively, this difference generates a second dimension to the running variable (elevation), a concern common to spatial regression discontinuity designs. This precise concern motivates the spatial fixed effects analysis introduced below which allows space to have a flexible relationship with unobservables in two dimensions. In Appendix Table A3, we show that being in the command area is associated with a 21.9 meter drop in elevation when estimating Specification 1 but only a 8.9 meter drop in elevation when estimating Specification 2. These differences are objectively small; perhaps unsurprisingly, our primary estimates are also robust to including elevation as a control (not reported).

thus have been left permanently fallow, may now be sufficiently productive to be worth cultivating with access to irrigation. While our sampling strategy selected both cultivated and uncultivated plots, it did not select plots which had been left overgrown with thick bushes, as it would have been difficult to identify the household responsible for those plots. In practice, the latter is likely uncommon, as typical household landholdings are small in the hillside irrigation schemes we study (around 0.3 ha), and agricultural land is highly valued – median annual rental prices in our data are 120,000 RwF/ha, approximately 20% of annual yields.

We account for this potential source of bias using spatial fixed effects (SFE; see Goldstein and Udry (2008); Conley and Udry (2010); Magruder (2012, 2013)), which use a spatial demeaning procedure to eliminate spatially correlated unobservables, such as unobserved heterogeneity in productivity caused by soil characteristics. This spatial demeaning ensures that comparisons are made only over proximate plots. For example, if some areas of low productivity are left forested outside of the command area, but not inside, then plots inside the command area will be systematically (unobservably) less productive than plots outside the command area. However, because SFE estimators only compare neighboring plots, the low productivity plots inside the command area that are near forested low productivity areas will not have nearby comparison plots outside the command area, and therefore will not contribute to the estimation of the effect of the command area.

In practice, we define a set $\mathcal{N}_{ist}^{\mathrm{SP}}$ to be the group of five closest sample plots to the sample plot of household $i$ that are observed in season $t$, including the household's sample plot itself. Then, for any variable $z_{ist}$, define $\overline{z_{ist}} = (1/|\mathcal{N}_{ist}^{\mathrm{SP}}|)\sum_{i' \in \mathcal{N}_{ist}^{\mathrm{SP}}} z_{i'st}$. The SFE specification then estimates

$$y_{ist}^{\mathrm{SP}} - \overline{y_{ist}^{\mathrm{SP}}} = \beta_1(\mathrm{CA}_{1is}^{\mathrm{SP}} - \overline{\mathrm{CA}_{is}^{\mathrm{SP}}}) + (V_{is}^{\mathrm{SP}} - \overline{V_{is}^{\mathrm{SP}}})'\gamma + (\epsilon_{ist}^{\mathrm{SP}} - \overline{\epsilon_{ist}^{\mathrm{SP}}}) \tag{2}$$

where $V_{is}^{\mathrm{SP}}$ includes all controls from Specification 1, except the subsumed site-by-season fixed effects.

Our sampling strategy yields the following plot proximity: restricting to the sample plots in our main sample for regression discontinuity analysis, 52% of plots have their nearest 3 plots (self inclusive) within 50 meters, and 89% have 3 plots within 100m; 61% of plots have all their nearest 5 plots (self inclusive) within 100m, while 85% have all 5 plots within 150m. As reference, Conley and Udry (2010) use 500m as the bandwidth for their estimator, while Goldstein and Udry (2008) use 250m as the bandwidth; we therefore anticipate that underlying land characteristics are likely to be quite similar between each plot and its comparison plots.

When estimating Specification 1, we cluster standard errors at the level of the nearest Water User Group, the group of plots that can source water from the same secondary pipe. The spatial fixed effects in Specification 2 generate correlation between the errors of close observations. In these specifications, we present Conley (1999) standard errors, where we allow households $j$ and $j'$'s errors to be correlated if there exists a household $i$ such that $j \in \mathcal{N}_{ist}^{\text{SP}}$ or $i \in \mathcal{N}_{jst}^{\text{SP}}$, and $j' \in \mathcal{N}_{ist}^{\text{SP}}$ or $i \in \mathcal{N}_{j'st}^{\text{SP}}$.

## 1 Balance

We now use Specifications 1 and 2 to test the exogeneity of access to irrigation on the sample plot. Specifically, we examine whether the plots in our sample and the households who cultivate them are comparable at baseline. For each of these specifications, we show balance both with key controls omitted (Columns 3 and 4), and our preferred specifications which we use in our analysis with key controls included (Columns 5 and 6).

First, we show in Table 2 that our sample plots are balanced in terms of pre-construction NDVI, baseline ownership, and rentals. 88% of sample plot owners on both sides of the canal owned the land over 5 years, or prior to the start of the irrigation construction. There is, however, some imbalance on plot size; as discussed in Section A, log area (measured in hectares) is larger inside the command area than outside the command area. This imbalance is weaker in the SFE specification than in the RDD specification, such that the omnibus test fails to reject the null of balance for the SFE specification (although we reject for the RDD specification). However, we note that this imbalance would bias us against finding the effects we see in Section B on horticulture, input use, labor use, and yields, as all of these variables are larger for smaller plots both inside and outside the command area.

Following the results on plot characteristics, Table 3 examines the characteristics of households whose sample plots are just inside or just outside the command area. First, note that we find significant imbalance on half of our variables when we do not restrict our sample to the discontinuity sample (Column 1, Table 3), and the omnibus test rejects the null of balance. However, balance improves substantially in our two preferred specifications (Columns 5 and 6, Table 3) which restrict to the

Table 2: Balance: Sample plot characteristics

| | Full sample | Discontinuity sample | | | | |
|---|---|---|---|---|---|---|
| | SP CA Coef. (SE) [p-value] | Dep. var. mean (Dep. var. SD) # of obs. | SP CA Coef. (SE) [p-value] | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Dep. var. (SP, Landsat) | | | | | | |
| 100 * NDVI (pre-2009) | 1.625 | 24.783 | 0.345 | 0.294 | 0.318 | 0.290 |
| | (0.210) | (2.220) | (0.254) | (0.189) | (0.251) | (0.186) |
| | [0.000] | 931 | [0.175] | [0.121] | [0.206] | [0.119] |
| 100 * NDVI (pre-2009, Dry) | 1.553 | 24.296 | 0.022 | 0.313 | 0.017 | 0.308 |
| | (0.226) | (2.449) | (0.278) | (0.206) | (0.276) | (0.206) |
| | [0.000] | 931 | [0.936] | [0.129] | [0.950] | [0.135] |
| 100 * NDVI (pre-2009, Rainy) | 1.664 | 24.992 | 0.483 | 0.296 | 0.449 | 0.293 |
| | (0.214) | (2.325) | (0.263) | (0.198) | (0.260) | (0.194) |
| | [0.000] | 931 | [0.066] | [0.135] | [0.084] | [0.131] |
| Dep. var. (SP, Baseline) | | | | | | |
| log GPS area | 0.045 | -2.529 | 0.454 | 0.214 | | |
| | (0.077) | (1.184) | (0.123) | (0.130) | | |
| | [0.554] | 931 | [0.000] | [0.099] | | |
| Own plot | -0.012 | 0.893 | 0.010 | 0.001 | 0.005 | -0.002 |
| | (0.020) | (0.310) | (0.032) | (0.038) | (0.032) | (0.038) |
| | [0.535] | 931 | [0.747] | [0.983] | [0.883] | [0.956] |
| Owned plot >5 years | 0.045 | 0.879 | 0.018 | 0.012 | 0.006 | 0.008 |
| | (0.019) | (0.327) | (0.037) | (0.036) | (0.036) | (0.035) |
| | [0.020] | 663 | [0.631] | [0.734] | [0.863] | [0.813] |
| Rented out, farmer | 0.027 | 0.033 | 0.001 | 0.008 | -0.006 | 0.006 |
| | (0.012) | (0.178) | (0.023) | (0.027) | (0.023) | (0.027) |
| | [0.022] | 931 | [0.983] | [0.764] | [0.813] | [0.836] |
| Slope | -0.006 | 0.290 | -0.006 | 0.005 | -0.010 | 0.005 |
| | (0.013) | (0.183) | (0.020) | (0.025) | (0.019) | (0.025) |
| | [0.655] | 931 | [0.752] | [0.831] | [0.610] | [0.833] |
| Omnibus F-stat [p] | 10.2 | | 3.1 | 0.9 | 1.1 | 0.7 |
| | [0.000] | | [0.003] | [0.490] | [0.361] | [0.692] |
| Controls | | | | | | |
| Site FE | | | X | | X | |
| SP distance to boundary | | | X | X | X | X |
| SP log GPS area | | | | | X | X |
| Spatial FE | | | | X | | X |

*Notes:* Column 2 presents the mean of the dependent variable and the standard deviation of the dependent variable in parentheses, for sample plots in the discontinuity sample that are outside the command area, and the total number of observations. Columns 1 and 3 through 6 present regression coefficients on a command area indicator, with cluster-robust standard errors (described in Section A) in parentheses, and p-values in brackets. Column 1 uses the full sample, while Columns 2 through 6 use the discontinuity sample. Columns 5 and 6 use Specifications 1 and 2, respectively.

discontinuity sample; households with sample plots just inside the command area appear similar to households with sample plots just outside the command area. For Specification 1 in Column 5, we find a significant difference in whether the household head has completed primary schooling or not and an almost significant difference in the number of household members (15-64). We also estimate Specifications 1 and 2 with controls included for household head completed primary schooling and number of household members (15-64), and our results are unaffected by the inclusion of these controls (not reported). We also note that 1 out of 10 variables significant at the 10% level is what one would expect due to chance.[10]

Lastly, in Section 1, we consider the characteristics of households' largest other plots; we show that these appear similarly balanced.

## B Estimating the effects of irrigation

We visually present three key findings on the plot-level impact of access to irrigation, or "sample plot shock", on profitability. In each of these figures, we leverage our spatial discontinuity by plotting average outcomes against distance to the command area boundary in meters, with a positive sign on distance indicating that the plot is in the command area. First, we show in Figure 2b that free access to irrigation in the command area significantly increases use of irrigation in the dry season; yet, use of irrigation is far from universal in the command area. This irrigation is productive: we show in Figure 3a that access to irrigation causes large increases in yields. In Section C, we documented that irrigation is used to cultivate dry season horticulture, and its profitability depends on the shadow wage used to price household labor. In Figure 3b, we show that impacts on profits are remarkably similar to those on yields when household labor is not priced, which follows as labor is the primary input to horticulture production. However, Figure 3b reveals that access to irrigation has no impacts on profits when household labor is priced at market wages. This suggests that the labor intensity of horticulture limits adoption of irrigation, and therefore that variation in the shadow wage may shape adoption decisions.

[10] That the coefficient on household head primary completion in our test for balance is positive might be particularly concerning if education strongly predicted decisions to adopt irrigation, as this would be suggestive of endogenous selection rather than sampling noise. In contexts where information is a first order constraint to adoption, more educated farmers have been shown to be significantly more likely to adopt new technologies (Foster and Rosenzweig, 1996). Throughout this paper, we argue that information is unlikely to be a first order constraint to adoption of irrigation in this context; consistent with this, in results not reported we find that household head primary completion is weakly negatively correlated with adoption of irrigation. To the extent that more educated household heads have a higher shadow wage, this is also consistent with our argument that labor market failures drive heterogeneous decisions to adopt irrigation in this context.

Table 3: Balance: Household characteristics

| | Full sample | Discontinuity sample | | | | |
|---|---|---|---|---|---|---|
| | SP CA Coef. (SE) [p-value] | Dep. var. mean (Dep. var. SD) # of obs. | SP CA Coef. (SE) [p-value] | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Dep. var. (HH, Baseline) | | | | | | |
| HHH female | 0.041 | 0.224 | 0.048 | 0.052 | 0.046 | 0.049 |
| | (0.025) | (0.417) | (0.046) | (0.050) | (0.046) | (0.050) |
| | [0.094] | 931 | [0.298] | [0.297] | [0.316] | [0.325] |
| HHH age | 0.5 | 47.3 | 2.0 | 1.0 | 1.2 | 0.5 |
| | (0.8) | (14.5) | (1.4) | (1.8) | (1.4) | (1.9) |
| | [0.497] | 929 | [0.155] | [0.593] | [0.382] | [0.785] |
| HHH completed primary | 0.069 | 0.284 | 0.113 | 0.094 | 0.103 | 0.090 |
| | (0.025) | (0.452) | (0.047) | (0.061) | (0.048) | (0.062) |
| | [0.005] | 928 | [0.016] | [0.125] | [0.030] | [0.146] |
| HHH worked off farm | 0.023 | 0.412 | -0.038 | -0.013 | -0.022 | -0.004 |
| | (0.027) | (0.493) | (0.051) | (0.063) | (0.050) | (0.063) |
| | [0.392] | 931 | [0.458] | [0.838] | [0.662] | [0.952] |
| # of plots | 0.61 | 5.16 | 0.17 | 0.38 | 0.35 | 0.44 |
| | (0.18) | (3.37) | (0.35) | (0.46) | (0.35) | (0.45) |
| | [0.001] | 931 | [0.625] | [0.403] | [0.315] | [0.334] |
| # of HH members | 0.17 | 4.89 | 0.01 | 0.01 | 0.01 | 0.01 |
| | (0.11) | (2.17) | (0.22) | (0.25) | (0.22) | (0.25) |
| | [0.104] | 931 | [0.969] | [0.963] | [0.975] | [0.965] |
| # who worked off farm | 0.10 | 0.77 | 0.02 | 0.03 | 0.02 | 0.04 |
| | (0.05) | (0.85) | (0.08) | (0.10) | (0.08) | (0.10) |
| | [0.039] | 931 | [0.800] | [0.794] | [0.794] | [0.712] |
| # of HH members (15-64) | 0.20 | 2.58 | 0.27 | 0.21 | 0.25 | 0.20 |
| | (0.08) | (1.45) | (0.15) | (0.16) | (0.16) | (0.16) |
| | [0.007] | 931 | [0.079] | [0.187] | [0.114] | [0.222] |
| Housing expenditures | -2.3 | 50.7 | -7.1 | -20.8 | -7.8 | -22.7 |
| | (6.9) | (129.5) | (15.0) | (18.5) | (14.8) | (18.6) |
| | [0.739] | 924 | [0.635] | [0.263] | [0.600] | [0.222] |
| Asset index | 0.11 | -0.05 | 0.13 | 0.06 | 0.11 | 0.04 |
| | (0.05) | (1.00) | (0.12) | (0.14) | (0.12) | (0.14) |
| | [0.034] | 929 | [0.285] | [0.662] | [0.370] | [0.754] |
| Omnibus F-stat [p] | 3.8 | | 1.8 | 1.0 | 1.6 | 1.1 |
| | [0.000] | | [0.058] | [0.408] | [0.123] | [0.379] |
| Controls | | | | | | |
| Site FE | | | X | | X | |
| SP distance to boundary | | | X | X | X | X |
| SP log GPS area | | | | | X | X |
| Spatial FE | | | | X | | X |

*Notes:* Column 2 presents the mean of the dependent variable and the standard deviation of the dependent variable in parentheses, for sample plots in the discontinuity sample that are outside the command area, and the total number of observations. Columns 1 and 3 through 6 present regression coefficients on a command area indicator, with cluster-robust standard errors (described in Section A) in parentheses, and p-values in brackets. Column 1 uses the full sample, while Columns 2 through 6 use the discontinuity sample. Columns 5 and 6 use Specifications 1 and 2, respectively.

We formally assess these impacts using regression evidence on the impacts of access to irrigation on crop choices, on input use, and on production, presented in Table 4.[11]

First, command area plots are 16pp - 18pp more likely to be irrigated during the dry season than plots outside the command area, and almost all of this increase is explained by the transition to cultivation of high value horticulture during this dry season. In contrast, adoption of irrigation during the rainy season is much lower, with increases of just 4pp - 6pp, consistent with the descriptives from Section C. In Appendix Figure A2, we find that these levels of adoption appear stable over time, suggesting that this may be steady state adoption. This transition to dry season horticulture substitutes for cultivation of bananas, a less productive but less input intensive commercial crop; we estimate a decrease of 14pp in the command area, and as a consequence we observe no impacts on overall cultivation in the dry season. As the rotations of staples and horticulture (or simply horticulture) that replace bananas may only involve two plantings and harvests, we therefore see a modest decrease in cultivation during the rainy seasons of 5pp - 9pp on a baseline of 84%. In Appendix Table A1, we confirm the sign of these results leveraging measures of vegetation from satellite imagery – we find that command area plots have less cultivation of bananas than plots outside the command area.

Second, we find large increases in dry season input use, which are dominated by increases in household labor. These results are consistent with the transition from low-intensity bananas, into input and labor-intensive horticulture. To interpret these results, we conduct a treatment on the treated analysis under the assumption that the command area increases input use only through its effect on irrigation. Doing so, we find that adoption of irrigation increases household labor use, input expenditures, and hired labor expenditures by 440 - 450 person-days/ha, 26,000 - 37,000 RwF/ha, and 14,000 - 20,000 RwF/ha, respectively; these numbers are similar to differences

[11] In Appendix Figure A3, we present graphical evidence of the regression discontinuity for all outcomes in the dry season, complementing evidence on key outcomes in Figures 2b and 3. We note that our results in Table 4 are consistent with the visual intuition in Figures 2b, 3, and A3. We also note that the RDD graphs in Figures 2b, 3, and A3 typically demonstrate a relatively strong spatial slope within the command area, where plots further away from the canal are more likely to irrigate, cultivate horticulture, and have higher yields. There are several plausible explanations for this trend: further away from the canal, more of the land is marshland, where growing horticulture is more traditional, plots may be selected differently (for example, smaller); and the system also has greater water pressure. The first two of these explanations emphasize the value of the regression discontinuity estimator while the last means that our local estimates at the canal may be conservative.

in input intensity of dry season horticulture and bananas reported in Table 1. The impacts on household labor are particularly large — valued at a typical wage of 800 RwF/person-day, this labor would be priced at 350,000 - 360,000 RwF/ha, an order of magnitude larger than the effects on input expenditures or hired labor expenditures. Applying this labor to 0.3 ha (median household landholdings) of command area land would require roughly 4 person-months of labor during the 3 month dry season. In contrast to these dry season results, we find no effects on input use during the rainy seasons.

Third, confirming our analysis of Figure 3, we find large increases in yields as a result of access to irrigation, while the profitability of the transition to dry season horticulture enabled by irrigation depends crucially on the shadow wage at which household labor is valued. To begin, we conduct a treatment on the treated analysis suggesting adoption of irrigation increases yields by 270,000 - 400,000 RwF/ha, 43% - 63% of annual agricultural production. As horticulture is primarily commercial, each 1 RwF/ha increase in yields is associated with a 0.77 - 0.89 RwF/ha increase in sales. Once again, these results on outputs are consistent with differences between bananas and horticulture production reported in Table 1. Estimated yield impacts are much larger than estimates of impacts on input and hired labor expenditures; these results suggest irrigation increases annual yields net of expenditures by 240,000 - 350,000 RwF/ha, a 43% - 62% increase. The relatively modest adoption rates alongside these impacts on yields net of expenditures seem initially somewhat at odds; however, these impacts implicitly place no value on the large increases in household labor. If we instead value household labor at 800 RwF/person-day, the median wage we observe, these impacts vanish completely.[12]

As an additional placebo test that these results are explained by access to irrigation, in Appendix Table A4 we compare plots inside and outside the command area in 2014 Dry, before irrigation was fully operational. Consistent with Appendix Figure A2, 2014 Dry differences in adoption of irrigation are small (3pp) compared to those we document in follow up surveys (16pp - 18pp). In addition, and consistent with our discussion in Section A, we find reduced banana cultivation (8pp - 11pp) with smaller differences than in follow up surveys (14pp). Consistent with these small increases

[12] In Appendix Table A11, we estimate impacts of access to irrigation on household welfare. Although these estimates are imprecise, all point estimates are positive and some are statistically significant. These results are consistent with positive impacts of access to irrigation on profits, although smaller impacts than implied by estimates that do not value household labor.

in horticulture and modestly larger decreases in low input intensity bananas, we do not find any significant differences in input use, yields, sales, or measures of profits, suggesting that the impacts we document in follow up surveys are explained by access to irrigation. Complementary to this analysis, we find that impacts in 2015 Rainy 1 and 2 are similar to rainy season impacts in subsequent seasons, as irrigation became fully operational at this time.

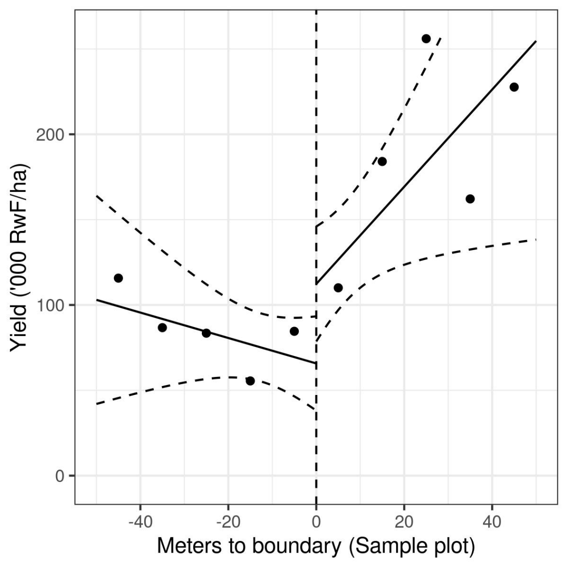


(a) Access to irrigation increases yields

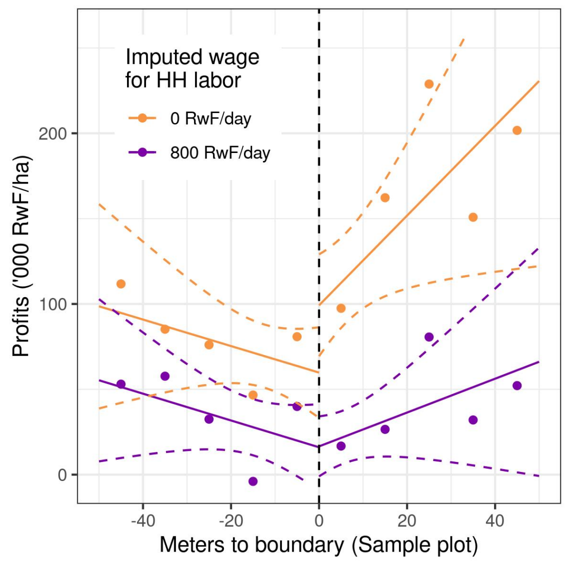


(b) Access to irrigation increases profits when household labor is not valued, but does not when it is valued at market wages

Figure 3: Irrigation increases yields, but impacts on profits depend on household shadow wage

Finally, we anticipate three categories of spillovers in our context: across household spillovers, within plot (across season) spillovers, and within household (across plot) spillovers. First, the across household spillovers we anticipate are general equilibrium effects, as the increased demand for labor and increased production of horticulture caused by the sample plot shock drives wages up and horticultural prices down. We examine wages and prices in Appendix Figures A5 and A6 and find no evidence that wages or staple prices changed over time, though prices of horticultural crops did decrease in one of the sites, reducing our estimated effects on revenues. Second, the within plot (across season) spillovers we anticipate are driven by the shift out of

Table 4: Access to irrigation enables transition to dry season horticulture from bananas, causes large increases in dry season labor and input use, yields, and sales; profitability depends on household's shadow wage

(a) Dry season

| | SP, Dry season, Discontinuity sample | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cultivated | Irrigated | Horticulture | Banana | HH labor/ha | Input exp./ha | Hired labor exp./ha | Yield | Sales /ha | Profits/ha Shadow wage = 0 | Profits/ha Shadow wage = 800 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) |
| RDD (Site-by-season FE, Specification 1) | | | | | | | | | | | |
| SP CA | 0.004 | 0.163 | 0.137 | -0.138 | 71.6 | 6.1 | 3.2 | 64.8 | 50.2 | 56.3 | 2.2 |
| | (0.041) | (0.024) | (0.024) | (0.037) | (18.2) | (1.5) | (1.9) | (23.0) | (14.3) | (20.9) | (16.5) |
| | [0.917] | [0.000] | [0.000] | [0.000] | [0.000] | [0.000] | [0.100] | [0.005] | [0.000] | [0.007] | [0.893] |
| SFE (Spatial FE, Specification 2) | | | | | | | | | | | |
| SP CA | 0.028 | 0.177 | 0.158 | -0.144 | 79.2 | 4.6 | 2.5 | 48.1 | 42.8 | 42.4 | -8.5 |
| | (0.043) | (0.030) | (0.028) | (0.034) | (21.2) | (1.8) | (2.4) | (26.9) | (17.4) | (24.4) | (20.2) |
| | [0.516] | [0.000] | [0.000] | [0.000] | [0.000] | [0.012] | [0.285] | [0.074] | [0.014] | [0.082] | [0.676] |
| # of observations | 2,439 | 2,439 | 2,438 | 2,438 | 2,428 | 2,431 | 2,431 | 2,307 | 2,431 | 2,307 | 2,305 |
| # of clusters | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 |
| Control mean | 0.383 | 0.051 | 0.058 | 0.244 | 60.1 | 2.4 | 3.1 | 80.5 | 47.3 | 75.2 | 31.6 |

(b) Rainy seasons

| | SP, Rainy seasons, Discontinuity sample | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | Cultivated | Irrigated | Horticulture | Banana | HH labor/ha | Input exp./ha | Hired labor exp./ha | Yield | Sales /ha | Profits/ha Shadow wage = 0 | Profits/ha Shadow wage = 800 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) |
| RDD (Site-by-season FE, Specification 1) | | | | | | | | | | | |
| SP CA | -0.085 | 0.038 | 0.025 | -0.164 | 15.5 | 2.0 | 3.5 | -18.4 | -10.7 | -22.7 | -33.4 |
| | (0.026) | (0.009) | (0.018) | (0.038) | (24.0) | (3.0) | (3.5) | (30.7) | (18.5) | (28.4) | (26.7) |
| | [0.001] | [0.000] | [0.168] | [0.000] | [0.520] | [0.503] | [0.313] | [0.550] | [0.563] | [0.425] | [0.211] |
| SFE (Spatial FE, Specification 2) | | | | | | | | | | | |
| SP CA | -0.052 | 0.062 | 0.053 | -0.170 | 10.4 | 2.2 | 3.1 | -19.2 | 5.8 | -23.2 | -31.3 |
| | (0.028) | (0.012) | (0.025) | (0.034) | (26.4) | (3.1) | (4.6) | (30.5) | (21.6) | (27.3) | (32.3) |
| | [0.057] | [0.000] | [0.031] | [0.000] | [0.695] | [0.475] | [0.500] | [0.530] | [0.787] | [0.395] | [0.333] |
| # of observations | 4,071 | 4,071 | 4,070 | 4,070 | 4,053 | 4,060 | 4,060 | 3,922 | 4,060 | 3,922 | 3,915 |
| # of clusters | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 | 173 |
| Control mean | 0.836 | 0.016 | 0.071 | 0.274 | 230.0 | 16.2 | 15.6 | 273.5 | 87.1 | 242.0 | 59.0 |

*Notes:* Regression coefficients on a sample plot command area indicator ("SP CA") are presented above. Specifications control for SP distance to the CA boundary and its interaction with SP CA, and SP log GPS area. RDD specification includes site-by-season fixed effects, and SFE specification includes spatial fixed effects. Cluster-robust standard errors (described in Section A) are in parentheses, and p-values in brackets.

bananas, which causes a change in patterns of cultivation during the rainy season, while adoption of irrigation is primarily during the dry season. However, we fail to find strong evidence of impacts on rainy season labor, inputs, yields, or profits.

Third, the within household (across plot) spillovers we anticipate are driven by the increase in demand for labor and inputs we observe on the sample plot, which may lead households to substitute labor and inputs away from their other plots. To address this spillover, in Section IV we model a household's agricultural production decisions and how they can generate substitution across plots, and in Section V we estimate these spillovers and quantify their implications for our estimates and for efficiency.

Taken together, these results suggest that irrigation leads to a large change in production practices for a minority of farmers. Those farmers cultivate horticulture in the dry season and a mix of horticulture, staples, and fallowing in the rainy seasons, they have substantially higher earnings in the dry season but similar earnings in the other seasons, and they invest more in inputs and much more in household labor in the dry seasons. Our estimates suggest that irrigation has the potential to be transformative in Africa, in light of the 43 - 62% increases in yields net of expenditures that we document from just three months of cultivation. At the same time, we observe a minority (30%) of farmers ultimately making use of the irrigation system. These results suggest that the shadow wage, and therefore labor market frictions, are likely to be important for the decision to cultivate horticulture. Building on this result, we next adapt the classical agricultural household model (Singh et al., 1986; Benjamin, 1992) to develop formal tests for the role of market failures in adoption of irrigation.

# IV Testing for binding constraint

In this section, we describe a model of a farmer household that chooses input and household labor allocations across multiple plots, subject to constraints in markets for labor, inputs, and risk. We use this model to generate testable predictions of the farmer household's responses to the sample plot shock for the presence of constraints and the nature of the constraints the farmer household faces. The formal setup of the model, motivating assumptions, and additional details are in Appendix C.

The farmer household has two plots: consistent with our data, we call these their sample plot and their largest other plot. Production on each plot is a function of input and household labor allocations on that plot, of the plot's productivity, and of a common productivity shock. While we abstract from the choice of crops or decision to irrigate in this framework, we interpret this production function as the envelope of production functions from cultivating different fractions of bananas and irrigated horticulture during the dry season, with cultivating horticulture as optimizing at a high labor and input intensity. The household maximizes expected utility over consumption and leisure, choosing its off-farm labor supply and its allocations of inputs and household labor on each plot.

We build on Benjamin (1992) and allow farmers to face three crucial constraints that cause deviations from expected profit maximization. First, access to insurance may be limited, so farmers may reduce labor and input use to avoid basis risk. Second, credit or access constraints may limit input use. Third, farmers' off-farm labor allocations may be constrained from above, resulting in overutilization of labor on the household farm. [13]

We model access to irrigation on the sample plot (the "sample plot shock") as a labor- and input-complementing increase in the productivity of the sample plot. This offers predictions consistent with our results in Section III: the sample plot shock increases production, labor allocations, and input allocations on the sample plot.

Next, we consider the impacts of the sample plot shock on production decisions on the largest other plot in this framework.

**Proposition 1.** *If no constraint binds, separation holds and input and labor use on the largest other plot does not respond to the sample plot shock.*

In the absence of the three constraints listed above, households maximize expected profits. As access to irrigation on the sample plot does not impact marginal products or prices on the largest other plot, labor and input allocations on the largest other plot do not respond to the sample plot shock.

**Proposition 2.** *If input, labor, or insurance constraints bind, then input and labor use are reduced on the largest other plot in response to the sample plot shock.*

The logic case-by-case is as follows. First, if input constraints bind, then the increase in inputs on the sample plot caused by access to irrigation must be associated with a reduction in inputs on the largest other plot. Second, if labor constraints bind, then the increase in labor on the sample plot caused by access to irrigation must be

[13] We model labor constraints as a constraint on off-farm labor allocations from above because this generates a shadow wage that is below the market wage, and we note this generalizes a model without an off-farm labor market. Agness et al. (2020) provide evidence in Kenya that households' shadow wages are almost always below the market wage, due to frictions in the labor market, while remaining agnostic as to what might cause these frictions. In addition, as we discuss in Section B, our results on profits are consistent with a shadow wage that is below the market wage.

associated with a reduction in the sum of labor on the largest other plot and leisure. Third, absent insurance, the increase in agricultural production caused by access to irrigation reduces the marginal utility from agricultural production relative to the marginal utility from consumption. In turn, this causes labor and input allocations on the largest other plot to fall.

An implicit assumption we make that generates this result is the absence of functioning land markets. With *perfectly* functioning land markets, shocks to the household's land endowment, such as the sample plot shock, should not affect productive decisions on the household's largest other plot. Instead, both the sample plot and the largest other plot would flow to the household with the highest willingness-to-pay for them. In practice, land transactions do occur; as discussed in Section 2, our survey tracks plots across transactions in land markets, so we are able to directly test the prediction that the sample plot shock does not affect the productive decisions on the largest other plot.[14]

Rejecting separation with the test suggested by Proposition 2 implies that the levels of irrigation adoption are inefficient and that land market failures contribute to this inefficiency. At the same time, this test does not allow us to test for which of the three constraints above interact with land market frictions to generate separation failures.

To shed light on which other constraints generate separation failures, we consider how households with different characteristics should *differentially* respond to the sample plot shock. We focus on two important household characteristics: household size, which determines the availability of household labor, and wealth, which determines the ability to purchase inputs.

**Proposition 3.** *If input or insurance constraints bind, then the input and labor allocations on the largest other plot of larger households (wealthier households) should be less (less) responsive to the sample plot shock.*

The intuition for this result is that both insurance and input constraints are ultimately financial constraints, which causes household size and wealth to enter the

[14]If households rent out their sample plot in response to the sample plot shock, then those households' input and labor use on sample plots will *fall* in response to the sample plot shock. In this case, we should expect average substitution away from largest other plots to be smaller than it would be if no households were able to endogenously rent out their sample plot; we are not able to evaluate this counterfactual, and instead acknowledge that imperfect but functioning land markets may partially reduce separation failures.

problem symmetrically. Under insurance constraints, both household size (by increasing the amount of labor income) and wealth increase household consumption. If we additionally assume that risk aversion is decreasing sufficiently quickly in consumption, then the allocations of wealthier and larger households will be closer to those maximizing expected profits, and therefore allocations on the largest other plot will be less responsive to the sample plot shock. Second, under input constraints, wealthier households are less likely to see the constraint bind. As the allocations on the largest other plot of unconstrained households do not respond to the sample plot shock, wealthier households should be less responsive. Similarly, if larger households could finance input purchases from labor income, larger households would be less likely to see the constraint bind. Therefore, their allocations on the largest other plot would be less responsive to the sample plot shock.

**Proposition 4.** *If labor constraints bind, then the relative responsiveness of input and labor allocations on the largest other plot of larger households (wealthier households) to the sample plot shock cannot be signed without further assumptions. If larger households and poorer households have more elastic on-farm labor supply schedules, and if on-farm labor supply exhibits sufficient curvature, then the input and labor allocations on the largest other plot of larger households (wealthier households) should be less (more) responsive to the sample plot shock.*

When labor constraints bind, the household responds to the sample plot shock by allocating additional labor to the sample plot, but may withdraw labor from either the largest other plot or from leisure. In general, the differential responses of wealthier and larger households cannot be signed. We focus on one particular case that builds on this intuition: where larger households have more elastic on-farm labor supply, while wealthier households have less elastic labor supply; a relationship has been posited as far back as Lewis (1954), and is discussed in depth in Sen (1966). These differences in on-farm labor supply generate the prediction that larger households should be less responsive to the sample plot shock, as they draw labor to the sample plot primarily from leisure, while wealthier households should be more responsive to the sample plot shock, as they draw labor primarily from the largest other plot.

These four propositions, summarized in Table 5, generate two sets of tests. First, Propositions 1 and 2 imply that substitution away from the largest other plot in response to the sample plot shock allows us to reject the absence of constraints.

Second, Propositions 3 and 4 produce a test of the absence of labor constraints. If the input and labor allocations of larger households are less responsive to the sample plot shock, while those of wealthier households are more responsive to the sample plot shock, then we would reject the absence of labor constraints.

We note that while it is unambiguous that constraints on selling labor would lead to inefficient adoption of irrigation, it is not clear whether they would result in adoption which is inefficiently high or low. This result depends on how wages respond to the removal of those constraints. We return to this point in our conclusions.

Table 5: Model predictions

| | $\frac{dL_2}{dA_1}$ | $\frac{d}{d\overline{L}}\frac{dL_2}{dA_1}$ | $\frac{d}{d\overline{M}}\frac{dL_2}{dA_1}$ |
|---|---|---|---|
| No constraints | 0 | 0 | 0 |
| Constraints | | | |
| Insurance | $-$ | $+$ | $+$ |
| Inputs | $-$ | $0/+$ | $+$ |
| Labor | $-$ | $+^*$ | $-^*$ |

*Notes:* Predicted signs from the model for key comparative statics of interest are presented in this table. $\frac{dL_2}{dA_1}$ is the effect of the sample plot shock on labor allocations on the largest other plot, and $\frac{d}{d\overline{L}}\frac{dL_2}{dA_1}$ and $\frac{d}{d\overline{M}}\frac{dL_2}{dA_1}$ are the impact of increased household size and wealth, respectively, on this effect. Predictions in the no constraints case correspond to Proposition 1. Predictions on $\frac{dL_2}{dA_1}$ correspond to Proposition 2. Predictions on $\frac{d}{d\overline{L}}\frac{dL_2}{dA_1}$ and $\frac{d}{d\overline{M}}\frac{dL_2}{dA_1}$ when insurance or input constraints bind correspond to Proposition 3, and when labor constraints bind correspond to Proposition 4. * is used to indicate predictions that hold when additional assumptions are made.

# V Separation failures and adoption of irrigation

## A Empirical strategy

When markets are complete, separation holds and access to irrigation on the sample plot does not affect production decisions on other plots. Our first specification mirrors Specification 1, which we use to estimate the impacts of irrigation. We still make use of the sample plot shock, but outcomes are now measured on the household's largest other plot (LOP) instead of the sample plot (SP).

$$y_{ist}^{\text{LOP}} = \beta_1 \text{CA}_{is}^{\text{SP}} + \beta_2 \text{Dist}_{is}^{\text{SP}} + \beta_3 \text{CA}_{is}^{\text{SP}} * \text{Dist}_{is}^{\text{SP}} + \beta_4 \text{CA}_{is}^{\text{LOP}} + \gamma_1 X_{is}^{\text{SP}} + \gamma_2 X_{is}^{\text{LOP}} + \alpha_{st} + \epsilon_{ist}^{\text{LOP}} \tag{3}$$

Specification 3 also includes controls $\text{CA}_{is}^{\text{LOP}}$, an indicator for whether the largest other plot is in the command area, and $X_{is}^{\text{SP}}$ and $X_{is}^{\text{LOP}}$, the log area of the sample plot and the largest other plot, respectively. We report $\beta_1$, the effect of the sample plot shock on outcomes on the largest other plot. In other specifications, we include $\text{CA}_{is}^{\text{SP}} * \text{CA}_{is}^{\text{LOP}}$ to account for potential heterogeneity with respect to the location of the largest other plot. In these specifications, we also report this difference in differences coefficient. For both this coefficient and $\beta_1$, in line with the model predictions in Table 5, we interpret negative coefficients on labor, inputs, irrigation use, and horticulture, as evidence of separation failures.

As in Section III, we include a specification with spatial fixed effects. Specifically, we estimate

$$y_{ist}^{\text{LOP}} - \overline{y_{ist}^{\text{LOP}}} = \beta_1 (\text{CA}_{is}^{\text{SP}} - \overline{\text{CA}_{is}^{\text{SP}}}) + (V_{is}^{\text{SP}} - \overline{V_{is}^{\text{SP}}})' \gamma_1 + (V_{is}^{\text{LOP}} - \overline{V_{is}^{\text{LOP}}})' \gamma_2 + (\epsilon_{ist}^{\text{LOP}} - \overline{\epsilon_{ist}^{\text{LOP}}}) \tag{4}$$

We then test for which constraints drive separation failures, leveraging our predictions in Section IV for heterogeneous effects of access to irrigation on the sample plot on production decisions on other plots with respect to household characteristics. Our benchmark specification to test for which constraints drive the separation failures is similar to Specifications 3 and 4, but also includes the interaction of household characteristics with the sample plot shock. We estimate

$$\begin{aligned} y_{ist}^{\text{LOP}} = \beta_1 \text{CA}_{is}^{\text{SP}} + W_{is}' \beta_2 + \text{CA}_{is}^{\text{SP}} * W_{is}' \beta_3 + \beta_4 \text{Dist}_{is}^{\text{SP}} + \beta_5 \text{CA}_{is}^{\text{SP}} * \text{Dist}_{is}^{\text{SP}} \\ + \beta_6 \text{CA}_{is}^{\text{LOP}} + \gamma_1 X_{is}^{\text{SP}} + \gamma_2 X_{is}^{\text{LOP}} + \alpha_{st} + \epsilon_{ist}^{\text{LOP}} \end{aligned} \tag{5}$$

where $W_{is}$ is a vector of household characteristics, which includes household size and an asset index in our primary specifications.[15] We focus on $\beta_3$: the heterogeneity,

[15]Through the lens of our model, household size and the asset index act as shifters of the household's availability of labor and the household's ability to purchase inputs, respectively. First, this requires that we have meaningful variation in household size and wealth in the cross-section. In our sample, household size has a standard deviation of 2.2 members (see Table 3), and a one standard deviation increase in the asset index corresponds to an additional 30,000 RwF of liquid assets (goats and chickens). For comparison, our estimates in Section B suggest that 1.3 members worth

with respect to household characteristics, of the impacts of the sample plot shock on outcomes on the largest other plot. The signs on $\beta_3$ produce our main test of which market failures cause separation failures; Table 5 presents which signs map to which market failures.

We present results only for the dry seasons (2016 Dry, 2017 Dry, and 2018 Dry), because these are the primary seasons for irrigation use, during which we anticipate substitution effects. Additionally, we present results only on cultivation decisions and input use, because we expect these substitution effects to be smaller than the direct effects and therefore we do not anticipate being able to detect effects on output. In Appendix Tables A9 and A10 we find no evidence of impacts on irrigation, labor, or input use during the rainy seasons, consistent with the lack of irrigation use on sample plots during those seasons.

## 1 Balance

Specifications 3 and 4 generate a credible test for inefficient irrigation use if LOPs are otherwise similar regardless of whether the SP is in the command area or just outside it. If the placement of the SP inside or outside the boundary is exogenous as suggested by our tests in Section 1, we should anticipate that LOPs are similar regardless of the location of the SP. We test whether largest other plots are comparable for households whose sample plot is just inside or just outside the command area using Specifications 3 and 4. As in Section 1, for each of these specifications, we show balance both with key controls omitted (Columns 3 and 4), and our preferred specifications which we use in our analysis with key controls included (Columns 5 and 6). Balance tests for largest other plots are reported in Table 6. First, note that specifications that do not restrict to the discontinuity sample perform particularly poorly here. Most notably, largest other plots are more likely to be located in the command area when sample plots are also located in the command area, as households' plots tend to be located near

of labor and 10,000 RwF of inputs are necessary to cultivate median household landholdings with horticulture during the dry season. Second, it requires that household size and the asset index are not correlated with other household or plot characteristics that might affect patterns of substitution. In Appendix Table A12, we justify our choice of shifters by showing that household size and asset index explain agricultural production decisions in the expected manner, and that this correlation is unaffected by the inclusion of key household covariates (results not reported). In addition, we verify that including other household and plot characteristics in the interaction – number of plots and number of command area plots (if households with more land are more price risk averse), and plot area (as a proxy for quality, if smaller plots are typically higher quality) – does not affect the significance of our results (not reported).

each other, and have much higher NDVI pre-construction. In contrast, our preferred specifications (Columns 5 and 6) which restrict to the discontinuity sample correct for these imbalances. For both specifications, the omnibus test fails to reject the null of balance.[16]

As an additional check, in Appendix Tables A7 and A8, we estimate for 2014 Dry Specifications 3 and 4. As the command area had not yet caused a large increase in demand for labor and inputs, or caused large increases in agricultural production, we would not anticipate any effects on LOPs in that baseline year. In line with this prediction, we do not find any consistent significant effects on LOPs.

## B Results

### 1 A test for separation failures

We now visually present our approach to documenting the impact of the sample plot shock on production decisions on the household's largest other plot, in order to test for separation failures. In Figure 4, as in earlier figures, we plot average outcomes against distance of the sample plot to the command area boundary in meters, with a positive sign indicating that the sample plot is in the command area. However, we now plot outcomes on both the sample plot and the largest other plot. In Figure 4, we observe substitution: access to irrigation on the sample plot causes irrigation use to sharply increase on the sample plot, and in turn causes irrigation use to sharply *decrease* on the largest other plot. As described in Section IV, this represents a separation failure – absent constraints, the household's production decisions on the largest other plot should not be affected by access to irrigation on the sample plot.

We now present our full regression evidence on the impacts of access to irrigation on the sample plot on production decisions on the household's largest other plot in Tables 7 and 8, with corresponding graphical evidence of these impacts presented in Figure 4 and Appendix Figure A4.

First, consistent with the presence of separation failures, we find households substitute labor and inputs away from their largest other plot. Households decrease allocations of household labor (36 person-days/ha) and inputs (5,800 - 6,600 RwF/ha)

[16] In results not reported, we verify that LOP characteristics demonstrate similar balance for heterogeneity with respect to the location of the LOP, household size, and household wealth when we estimate Specification 5.

Table 6: Balance: Largest other plot characteristics

| | Full sample | Discontinuity sample | | | | |
|---|---|---|---|---|---|---|
| | SP CA Coef. (SE) [p-value] | Dep. var. mean (Dep. var. SD) # of obs. | SP CA Coef. (SE) [p-value] | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Dep. var. (LOP, Landsat) | | | | | | |
| 100 * NDVI (pre-2009) | 0.990 | 24.976 | 0.242 | 0.204 | 0.314 | 0.277 |
| | (0.194) | (2.288) | (0.263) | (0.269) | (0.272) | (0.282) |
| | [0.000] | 757 | [0.358] | [0.448] | [0.248] | [0.326] |
| 100 * NDVI (pre-2009, Dry) | 0.883 | 24.568 | -0.120 | -0.043 | 0.002 | 0.042 |
| | (0.213) | (2.478) | (0.267) | (0.285) | (0.274) | (0.295) |
| | [0.000] | 757 | [0.652] | [0.879] | [0.993] | [0.887] |
| 100 * NDVI (pre-2009, Rainy) | 1.032 | 25.156 | 0.391 | 0.311 | 0.442 | 0.381 |
| | (0.196) | (2.400) | (0.285) | (0.278) | (0.294) | (0.291) |
| | [0.000] | 757 | [0.171] | [0.263] | [0.133] | [0.191] |
| Dep. var. (LOP, Baseline) | | | | | | |
| log GPS area | -0.108 | -2.373 | 0.125 | 0.078 | | |
| | (0.068) | (1.040) | (0.128) | (0.134) | | |
| | [0.114] | 757 | [0.326] | [0.559] | | |
| Own plot | 0.025 | 0.878 | 0.035 | 0.035 | 0.033 | 0.031 |
| | (0.019) | (0.328) | (0.033) | (0.039) | (0.033) | (0.038) |
| | [0.174] | 757 | [0.291] | [0.370] | [0.313] | [0.404] |
| Owned plot >5 years | 0.005 | 0.960 | 0.017 | 0.035 | 0.016 | 0.032 |
| | (0.014) | (0.197) | (0.024) | (0.024) | (0.022) | (0.025) |
| | [0.728] | 563 | [0.469] | [0.147] | [0.489] | [0.196] |
| Rented out, farmer | 0.013 | 0.033 | -0.026 | -0.040 | -0.028 | -0.040 |
| | (0.010) | (0.180) | (0.022) | (0.025) | (0.023) | (0.026) |
| | [0.224] | 757 | [0.234] | [0.113] | [0.226] | [0.121] |
| Slope | 0.024 | 0.269 | 0.009 | -0.007 | 0.009 | -0.008 |
| | (0.009) | (0.144) | (0.015) | (0.018) | (0.015) | (0.018) |
| | [0.012] | 757 | [0.568] | [0.687] | [0.529] | [0.643] |
| Command area | 0.187 | 0.395 | -0.046 | -0.074 | | |
| | (0.032) | (0.490) | (0.058) | (0.059) | | |
| | [0.000] | 757 | [0.430] | [0.210] | | |
| Omnibus F-stat [p] | 5.9 | | 0.9 | 1.5 | 0.9 | 1.3 |
| | [0.000] | | [0.518] | [0.147] | [0.480] | [0.227] |
| Controls | | | | | | |
| Site FE | | | X | | X | |
| SP distance to boundary | | | X | X | X | X |
| SP log GPS area | | | | | X | X |
| LOP log GPS area | | | | | X | X |
| LOP CA | | | | | X | X |
| Spatial FE | | | | X | | X |

*Notes:* Column 2 presents the mean of the dependent variable and the standard deviation of the dependent variable in parentheses, for sample plots in the discontinuity sample that are outside the command area, and the total number of observations. Columns 1 and 3 through 6 present regression coefficients on a command area indicator, with cluster-robust standard errors (described in Section A) in parentheses, and p-values in brackets. Column 1 uses the full sample, while Columns 2 through 6 use the discontinuity sample. Columns 5 and 6 use Specifications 3 and 4, respectively.

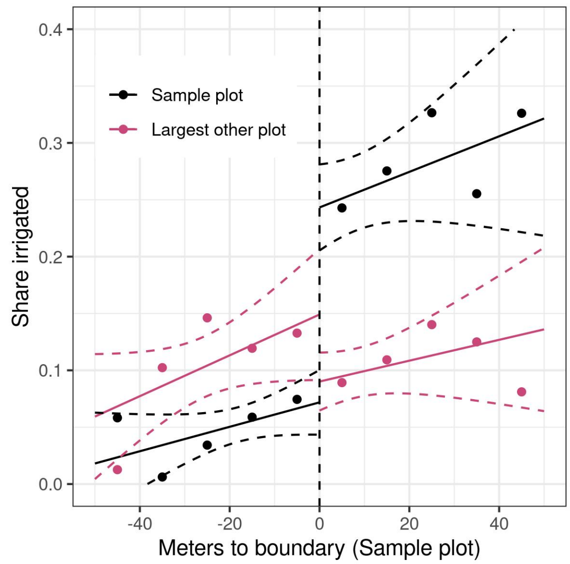


Figure 4: Separation fails, as access to irrigation on the sample plot causes substitution of irrigation use away from the largest other plot

on their largest other plot in response to the sample plot shock. Additionally, they substitute away from labor and input intensive technologies, consistent with our interpretation of the production function as the envelope of production functions across crop choices. Households decrease use of irrigation (4.1pp - 4.9pp) and cultivation of horticulture (4.2pp - 4.4pp), while increasing cultivation of bananas (6.7pp - 8.6pp). In Appendix Table A2, we confirm the sign of these results leveraging measures of vegetation from satellite imagery – we find that access to irrigation on sample plots causes increases in cultivation of bananas on largest other plots.

Next, we expect the results above to be driven primarily by largest other plots located in the command area for most outcomes. This is because there is limited irrigation, and therefore input use or horticulture during the dry season, on plots that cannot be irrigated. Consistent with this, we find our results on irrigation, horticulture, and inputs are all driven by plots located in the command area. When the largest other plot is located in the command area, the 16 - 18pp increase in irrigation use on sample plots in the command area coincides with an 11pp decrease in irrigation use on the largest other plot; these relative magnitudes suggest that separation failures cause few households to be able to use irrigation on more than one plot in the command area.

These results on separation failures imply the existence of a within-household

Table 7: Sample plot shock causes households to substitute labor and input intensive irrigated horticulture away from largest other plot

| | LOP, Dry season, Discontinuity sample | | | | | | |
|---|---|---|---|---|---|---|---|
| | Culti-vated | Irri-gated | Horti-culture | Banana | HH labor/ ha | Input exp./ha | Hired labor exp./ha |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| RDD (Site-by-season FE, Specification 3) | | | | | | | |
| SP CA | 0.028 | -0.049 | -0.044 | 0.086 | -36.4 | -5.8 | -1.9 |
| | (0.040) | (0.026) | (0.023) | (0.032) | (20.6) | (2.7) | (2.1) |
| | [0.486] | [0.055] | [0.058] | [0.007] | [0.078] | [0.032] | [0.365] |
| Sample plot effect | 0.004 | 0.163 | 0.137 | -0.138 | 71.6 | 6.1 | 3.2 |
| SFE (Spatial FE, Specification 4) | | | | | | | |
| SP CA | 0.002 | -0.041 | -0.042 | 0.067 | -36.2 | -6.6 | -0.1 |
| | (0.048) | (0.032) | (0.028) | (0.036) | (24.2) | (2.8) | (2.2) |
| | [0.973] | [0.206] | [0.137] | [0.065] | [0.134] | [0.017] | [0.956] |
| Sample plot effect | 0.028 | 0.177 | 0.158 | -0.144 | 79.2 | 4.6 | 2.5 |
| # of observations | 2,107 | 2,107 | 2,107 | 2,107 | 2,094 | 2,097 | 2,097 |
| # of clusters | 165 | 165 | 165 | 165 | 165 | 165 | 165 |
| Control mean | 0.368 | 0.114 | 0.107 | 0.201 | 68.1 | 5.4 | 3.7 |

*Notes:* Regression coefficients on a sample plot command area indicator ("SP CA") are presented above. Specifications control for SP distance to the command area boundary, its interaction with SP CA, SP and LOP log GPS areas, and LOP CA. RDD specification includes site-by-season fixed effects, and SFE specification includes spatial fixed effects. Cluster-robust standard errors (described in Section A) are in parentheses, and p-values in brackets. "Sample plot effect" estimates are from Table 4.

negative spillover, as they show that having one additional plot in the command area causes a household to reduce their use of irrigation, labor, and inputs on their other plots. In principle, this means that our estimates of the impacts of irrigation are the impacts of irrigating one of a farmers' plots, gross of any input reallocations made by the farmers across plots in response to that irrigation. Table 8 indicates that by far the largest spillovers induce substitution away from largest other plots within the command area.[17] We therefore conclude that the dominant within-household spillover is a reduced intensity of cultivation on irrigated plots. This suggests any bias in our estimates caused by spillovers across plots render estimates in Section III conservative.

[17] The substitution of inputs we estimate from largest other plots outside the command area is generally not significantly different from zero, and the largest point estimate implies it is 26% as large as substitution away from command area plots.

Table 8: Sample plot shock causes households to substitute labor and input intensive irrigated horticulture away from largest other plot

| | LOP, Dry season, Discontinuity sample | | | | | | |
|---|---|---|---|---|---|---|---|
| | Culti-vated | Irri-gated | Horti-culture | Banana | HH labor/ ha | Input exp./ha | Hired labor exp./ha |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| RDD (Site-by-season FE, Specification 3) | | | | | | | |
| SP CA | 0.072 | -0.002 | -0.002 | 0.085 | -11.2 | -2.7 | 0.5 |
| | (0.044) | (0.020) | (0.018) | (0.041) | (14.7) | (1.8) | (2.1) |
| | [0.101] | [0.921] | [0.903] | [0.038] | [0.447] | [0.137] | [0.820] |
| SP CA * LOP CA | -0.103 | -0.112 | -0.101 | 0.003 | -59.8 | -7.5 | -5.7 |
| | (0.054) | (0.035) | (0.035) | (0.042) | (24.4) | (3.2) | (2.4) |
| | [0.055] | [0.002] | [0.004] | [0.940] | [0.015] | [0.020] | [0.020] |
| Joint F-stat [p] | 2.3 | 5.0 | 4.1 | 4.0 | 3.0 | 2.9 | 2.8 |
| | [0.106] | [0.008] | [0.017] | [0.020] | [0.053] | [0.059] | [0.066] |
| Sample plot effect | 0.004 | 0.163 | 0.137 | -0.138 | 71.6 | 6.1 | 3.2 |
| Average effect | 0.028 | -0.049 | -0.044 | 0.086 | -36.4 | -5.8 | -1.9 |
| SFE (Spatial FE, Specification 4) | | | | | | | |
| SP CA | 0.065 | 0.018 | 0.002 | 0.082 | -9.8 | -2.7 | 2.6 |
| | (0.048) | (0.025) | (0.021) | (0.044) | (19.1) | (1.9) | (2.4) |
| | [0.175] | [0.464] | [0.936] | [0.064] | [0.609] | [0.149] | [0.283] |
| SP CA * LOP CA | -0.140 | -0.131 | -0.097 | -0.031 | -58.6 | -8.7 | -6.0 |
| | (0.059) | (0.045) | (0.043) | (0.044) | (32.3) | (3.9) | (3.0) |
| | [0.018] | [0.003] | [0.025] | [0.475] | [0.070] | [0.027] | [0.047] |
| Joint F-stat [p] | 3.1 | 4.4 | 2.5 | 1.8 | 1.8 | 3.2 | 2.0 |
| | [0.045] | [0.012] | [0.081] | [0.164] | [0.172] | [0.040] | [0.135] |
| Sample plot effect | 0.028 | 0.177 | 0.158 | -0.144 | 79.2 | 4.6 | 2.5 |
| Average effect | 0.002 | -0.041 | -0.042 | 0.067 | -36.2 | -6.6 | -0.1 |
| # of observations | 2,107 | 2,107 | 2,107 | 2,107 | 2,094 | 2,097 | 2,097 |
| # of clusters | 165 | 165 | 165 | 165 | 165 | 165 | 165 |
| Control mean | 0.368 | 0.114 | 0.107 | 0.201 | 68.1 | 5.4 | 3.7 |

*Notes:* Regression coefficients on a sample plot command area indicator ("SP CA") and its interaction with a largest other plot command area indicator ("SP CA * LOP CA") are presented above. Specifications control for SP distance to the command area boundary, its interaction with SP CA, SP and LOP log GPS areas, and LOP CA. RDD specification includes site-by-season fixed effects, and SFE specification includes spatial fixed effects. Cluster-robust standard errors (described in Section A) are in parentheses, and p-values in brackets. "Sample plot effect" estimates are from Table 4, and "Average effect" estimates are from Table 7.

## 2 Impacts of separation failures on adoption of irrigation

We now quantify the impact of separation failures on adoption of irrigation. We ask what would happen to adoption of irrigation if all households with their sample plot and largest other plot in the command area only had one of the two plots in the command area. This counterfactual follows naturally from our estimates of the effect of the sample plot shock on adoption of irrigation on the largest other plot.

Specifically, we calculate

$$\frac{2 * (\#\text{ of HH with 2 CA plots})}{2 * (\#\text{ of HH with 2 CA plots}) + (\#\text{ of HH with 1 CA plot})}(\beta_1 + \beta_{3,\text{CA}}) \qquad (6)$$

First, $(\beta_1 + \beta_{3,\text{CA}})$ (from Specification 5) is the total effect of the sample plot shock on adoption of irrigation on largest other plots in the command area. Second, in the denominator, we count the total number of command area plots among households' sample plots and largest other plots. As we ignore other plots, we interpret this exercise as estimating a lower bound of the impact of reallocation of all of the household's plots on adoption of irrigation. Third, in the numerator, we apply the estimated substitution caused by the sample plot shock to both the sample plot and the largest other plot, as households are also substituting away from their sample plot when the largest other plot is in the command area. This leverages an implicit assumption that spillovers across plots are symmetric; in other words, the impact of sample plot access to irrigation on largest other plot irrigation we report in Tables 8 and the impact of largest other plot access to irrigation on sample plot irrigation are equal. In Appendix Table A6, we demonstrate that these spillovers are remarkably similar; we estimate impacts of largest other plot access to irrigation by restricting to largest other plots within 50 meters of the command area boundary and flipping the sample plot and largest other plot in Specifications 1, 2, 3, and 4.

We find adoption of irrigation would be 5.5pp higher under this counterfactual, which represents a 34% increase.[18] This counterfactual relates to land market frictions – absent these frictions, we would expect that the increased adoption of irrigation caused by this reallocation would be achieved through land markets. Intuitively, under perfect land markets, characteristics of the household that manages a particular command area plot at baseline, including the number of other command area plots that household managed at baseline, should not affect equilibrium adoption of irrigation on that plot. Relatedly, as shown in the model, this would also be true if all markets (except potentially land markets) were frictionless.

[18]The p-value on this estimate is 0.035, which we calculate using block bootstrapped standard errors at the nearest Water User Group level to account for uncertainty in both the numerator and denominator of Equation 6.

## 3 Separating constraints

We now provide evidence on the source of the separation failure by estimating heterogeneous impacts of the sample plot shock on outcomes on the largest other plot. Recall that for this analysis, the model makes two key predictions. First, if only insurance or input constraints bind, wealthier households and larger households should be less responsive. Second, if only labor constraints bind, the differential responsiveness of wealthier households and larger households cannot in general be signed. However, under additional assumptions, households with more elastic on-farm labor supply (likely poorer households and larger households) should be less responsive. Note that this test does not allow us to reject a null that a particular constraint exists; any pattern of differential responses is consistent with all constraints binding. However, if we observe that wealthier households are more responsive, we can reject the null of no labor constraints. Additionally, we would interpret observing wealthier households to be more responsive and larger households to be less responsive as the strongest evidence of the presence of labor constraints from this test.

We present the results of this test in Table 9. First, larger households are less responsive to the sample plot shock across every outcome. A household with 2 additional members, approximately one standard deviation of household size, is less responsive to the sample plot shock on its largest other plot by 37% - 78% for irrigation use, 59% - 86% for horticulture, 42% - 55% for household labor, and 17% - 18% for inputs, with all but the input coefficient statistically significant. In contrast, wealthier households are more responsive to the sample shock across these same outcomes. A household with a one standard deviation higher asset index is more responsive to the sample plot shock on its largest other plot by 45% - 90% for irrigation use, 39% - 72% for horticulture, 35% - 62% for household labor, and 48% - 60% for input use; however, these results are less precise. In effect, these results suggest that our estimates of separation failures are driven by the behavior of small, rich households, while large, poor households do not change their allocations on their largest other plot in response to the sample plot shock. As discussed in Section IV, these results are very difficult to reconcile with a model that does not feature labor market failures.

In sum, these results provide strong evidence for the existence of labor market failures that generate separation failures, which in turn cause inefficient adoption of irrigation.

Table 9: Larger and poorer households do not substitute away from largest other plot in response to sample plot shock

| | LOP, Dry season, Discontinuity sample | | | | | | |
|---|---|---|---|---|---|---|---|
| | Culti-vated | Irri-gated | Horti-culture | Banana | HH labor/ ha | Input exp./ha | Hired labor exp./ha |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| RDD (Site-by-season FE, Specification 5) | | | | | | | |
| SP CA | -0.063 | -0.094 | -0.109 | 0.050 | -73.3 | -8.0 | -6.1 |
| | (0.086) | (0.048) | (0.046) | (0.064) | (34.6) | (4.2) | (3.2) |
| | [0.466] | [0.050] | [0.018] | [0.440] | [0.034] | [0.054] | [0.059] |
| SP CA * # of HH members | 0.018 | 0.009 | 0.013 | 0.007 | 7.6 | 0.5 | 0.8 |
| | (0.013) | (0.007) | (0.007) | (0.011) | (4.3) | (0.5) | (0.4) |
| | [0.176] | [0.204] | [0.079] | [0.513] | [0.081] | [0.376] | [0.032] |
| SP CA * Asset index | -0.008 | -0.022 | -0.017 | 0.001 | -12.8 | -2.8 | -0.2 |
| | (0.026) | (0.016) | (0.016) | (0.023) | (9.9) | (1.5) | (1.4) |
| | [0.760] | [0.167] | [0.296] | [0.963] | [0.198] | [0.061] | [0.898] |
| Joint F-stat [p] | 1.0 | 1.8 | 2.0 | 2.7 | 1.6 | 2.1 | 1.7 |
| | [0.383] | [0.153] | [0.115] | [0.047] | [0.190] | [0.100] | [0.169] |
| Average effect | 0.028 | -0.049 | -0.044 | 0.086 | -36.4 | -5.8 | -1.9 |
| SFE (Spatial FE, Specification 5) | | | | | | | |
| SP CA | -0.183 | -0.117 | -0.130 | -0.058 | -83.6 | -9.3 | -4.8 |
| | (0.099) | (0.051) | (0.046) | (0.084) | (39.9) | (4.2) | (3.2) |
| | [0.065] | [0.021] | [0.005] | [0.489] | [0.036] | [0.026] | [0.138] |
| SP CA * # of HH members | 0.038 | 0.016 | 0.018 | 0.025 | 10.0 | 0.6 | 0.9 |
| | (0.015) | (0.008) | (0.008) | (0.015) | (4.7) | (0.5) | (0.4) |
| | [0.010] | [0.049] | [0.016] | [0.088] | [0.032] | [0.269] | [0.019] |
| SP CA * Asset index | -0.038 | -0.037 | -0.030 | -0.009 | -22.6 | -4.0 | -0.5 |
| | (0.032) | (0.018) | (0.020) | (0.027) | (12.3) | (1.6) | (1.4) |
| | [0.232] | [0.044] | [0.139] | [0.737] | [0.067] | [0.016] | [0.734] |
| Joint F-stat [p] | 3.0 | 2.4 | 2.7 | 2.3 | 2.0 | 2.5 | 2.0 |
| | [0.031] | [0.069] | [0.045] | [0.072] | [0.110] | [0.055] | [0.115] |
| Average effect | 0.002 | -0.041 | -0.042 | 0.067 | -36.2 | -6.6 | -0.1 |
| # of observations | 2,104 | 2,104 | 2,104 | 2,104 | 2,091 | 2,094 | 2,094 |
| # of clusters | 165 | 165 | 165 | 165 | 165 | 165 | 165 |
| Control mean | 0.368 | 0.114 | 0.107 | 0.201 | 68.1 | 5.4 | 3.7 |

*Notes:* Regression coefficients on a sample plot command area indicator ("SP CA") and its interaction with W ("SP CA * W") are presented above. Specifications control for SP distance to the command area boundary, its interaction with SP CA, SP and LOP log GPS areas, LOP CA, and all W. RDD specification includes site-by-season fixed effects, and SFE specification includes spatial fixed effects. Cluster-robust standard errors (described in Section A) are in parentheses, and p-values in brackets. "Average effect" estimates are from Table 7.

# VI Experimental evidence

Our results leveraging the discontinuity suggest that land and labor market frictions constrain the adoption of hillside irrigation in Rwanda. In this section, we present evidence from three randomized controlled trials testing other classical constraints to adoption of irrigation: the empowerment of local monitors to address operation and management challenges of irrigation schemes, subsidies for command area land taxes to reduce potential financial burdens, and the provision of horticultural minikits to alleviate financial and informational constraints. We find no impacts of these interventions on access to irrigation, and we combine this result with descriptive statistics and impacts on intermediate outcomes to provide evidence that the three targeted constraints do not explain low adoption of irrigation in our context.

## A O&M monitors

First, we test whether failures of scheme operations and maintenance (O&M) limit farmers' adoption of irrigation. Qualitative work raised concerns that the Water User Groups as established may not be sufficient to enforce water usage schedules and that routine maintenance tasks would not be performed adequately, as has been documented by Ostrom (1990). If farmers faced limited access to water due to problems in the centralized O&M system, this could constrain adoption of irrigation. To alleviate this potential constraint, we randomly assigned Water User Groups to different approaches to O&M monitoring. 251 Water User Groups across three irrigation sites were randomized, stratified across the 33 Zones these irrigation sites are divided into, into three arms. These treatment arms were as follows: 40% were assigned to a status quo arm, where "irrigator/operators" employed by the site were responsible for enforcing water usage schedules and reporting O&M problems to the local Water User Association. 30% were assigned to an arm where the Water User Group elected a monitor who was tasked with these responsibilities, trained in implementing them, and given worksheets to fill and return to the Water User Association reporting challenges with enforcement of the water usage schedule and any O&M concerns. In an additional 30%, the elected monitor was required to have a plot near the top of the Water User Group, where the flow of water is most negatively impacted when too many farmers try to irrigate at once. Monitors were trained just before the 2016 Dry season, with refresher trainings during 2016 Dry and 2017 Rainy 1.

For O&M, and each of the two other experiments, we estimate the impacts of assignment to treatment controlling for stratifying covariates on outcomes on the household's sample plot unless otherwise stated. We restrict to command area sample plots for precision, and cluster robust standard errors at the Water User Group level.

We regress

$$y_{ist}^{\mathrm{SP}} = \beta \mathrm{Treated}_{is}^{\mathrm{SP}} + (X_{is}^{\mathrm{SP}})'\gamma + \epsilon_{ist}^{\mathrm{SP}} \tag{7}$$

where $\mathrm{Treated}_{is}^{\mathrm{SP}}$ measures assignment to treatment for the sample plot unless otherwise stated, and $X_{is}^{\mathrm{SP}}$ includes stratifying covariates. For O&M, we define treatment as assignment to either of the two "farmer monitor" arms, and include Zone fixed effects and minikit saturation as controls. We present estimated treatment effects of each of the three interventions, and control group means, in Figure 5 and equivalently in Appendix Table A13.

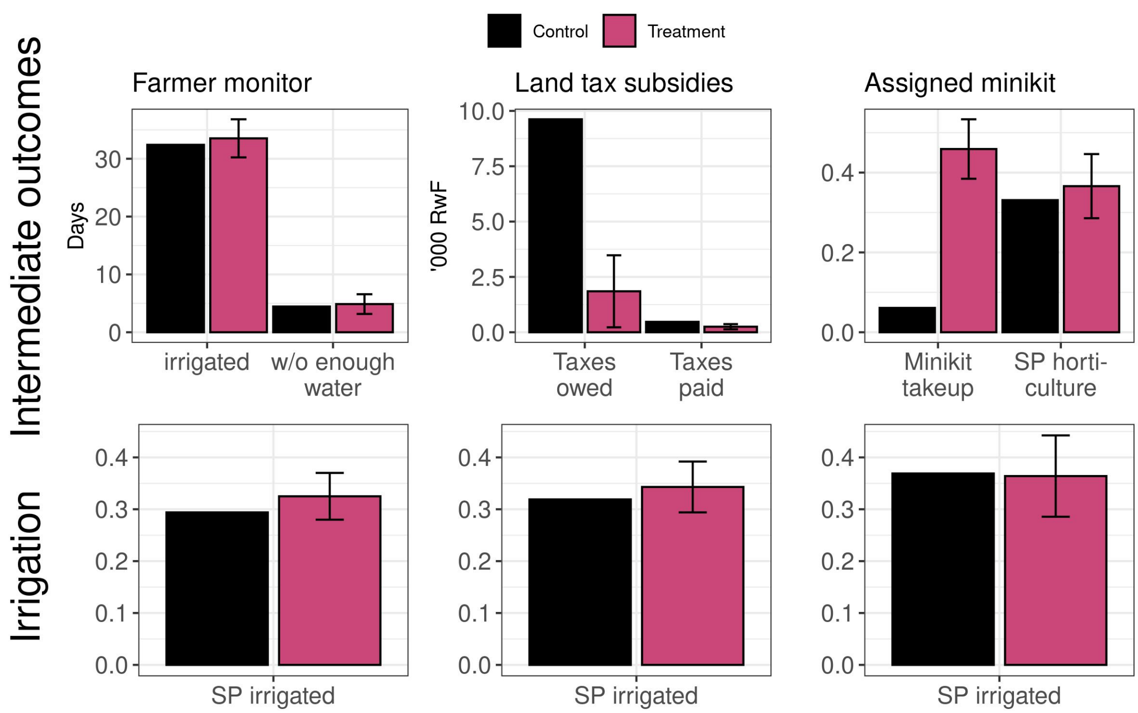


Figure 5: Interventions targeting operations and maintenance, land taxes, and access to inputs did not increase adoption of irrigation

*Notes:* Control mean outcomes are presented in black, and control mean outcomes plus estimated treatment effects from Equation 7 are presented in pink with 95% confidence intervals on treatment effects. Outcomes variables are on the horizontal axis, and treatment variables are titles of each column of graphs.

We evaluate the effectiveness of these O&M monitoring interventions in the first column of Figure 5. The top panel presents water availability for farmers who chose to irrigate their plots. We do not observe impacts on the number of days successfully irrigated or the number of days without sufficient water between these two groups.

Moreover, the number of days without water is small both in absolute terms (4.5 days) and as a proportion of the number of days irrigated (32.4). Perhaps unsurprisingly then, the bottom panel demonstrates that these treatments did not induce a change in the number of farmers who irrigate their sample plots; plots in monitored Water User Groups are irrigated 3.1pp more frequently, with a 95% confidence interval of (-1.4pp, 7.6pp). We conclude that the O&M monitoring interventions did not meaningfully affect constraints to adoption of irrigation. As these recently constructed irrigation systems had relatively few maintenance issues, it seems likely that O&M challenges did not present a binding constraint to adoption of irrigation over our study period.

## B Land taxes

Second, the government charged farmers with plots in the command area land taxes to fund O&M in the schemes. These taxes were as high as 77,000 RwF/ha/year, unconditional on cultivation decisions. For reference, this is roughly 20% of our dry season treatment on the treated estimates on yields, and roughly 60% of median land rental prices. As land taxes are assessed unconditionally, these taxes should not have influenced optimal cultivation decisions. However, if farmers believed that they were more likely to be required to pay the taxes if they used the irrigation infrastructure, then these taxes could influence farmers' production decisions, even though they are small relative to potential yield gains from irrigation use. To test whether these taxes limit farmers' adoption of irrigation, we randomized tax subsidies for farmers. Public lotteries for subsidies were conducted at the Zone level; 80% of households in our sample were entered as intended into just one lottery. At these public lotteries, 40% of farmers received no subsidy, 20% received a 50% subsidy for one season, 20% received a 100% subsidy for one season, and 20% received a 100% subsidy for two seasons. The lotteries took place at the start of the 2017 Rainy 1, and subsidies were for 2017 Rainy 1 and 2017 Rainy 2; at the time the Water User Associations did not plan to collect the land tax during the 2017 Dry season.

We estimate the impacts of land taxes using Equation 7; we define treatment as the average subsidy across the two seasons that the household was assigned to receive in its highest subsidy lottery (0.25, 0.5, and 1 across the 3 subsidy arms), and normalize treatment to be mean one for treated households to facilitate comparison to the two other interventions. We include Zone fixed effects, minikit saturation, and dummies for number of land tax subsidy lotteries the household was entered into.

The second column of Figure 5 presents results from this experiment, using the site's administrative records on taxes owed and paid. Control farmers owe, on average 9,600 RwF in taxes; subsidies reduce taxes owed by 7,800 RwF, an 81% reduction. However, almost no farmers in either the subsidy treatment or control actually pay their full taxes owed. Control farmers average 500 RwF in taxes paid, which subsidies reduce by 200 RwF. At such low levels of tax compliance, it seems unlikely that these land taxes impose a binding constraint on usage. However, if farmers are uncertain whether taxes would be enforced, then beliefs about enforcement may still constrain use. In the bottom panel, we find they did not: subsidies increase irrigation use by 2.4pp, with a 95% confidence interval of (-2.5pp, 7.3pp), suggesting either farmers knew effective tax rates would be low or that fully enforced land taxes would not have constrained adoption of irrigation.

## C Minikits

Third, we test whether financial and informational constraints limit adoption of irrigation. We assigned horticultural minikits to farmers randomly selected from Water User Group member lists. Each minikit included horticultural seeds, chemical fertilizer, and insecticide, in sufficient quantities to cultivate 0.02 ha. In principle, these minikits should resolve constraints related to input access, including credit constraints. In addition, they should reduce basis risk which may resolve insurance constraints and facilitate experimentation if information is a constraint. In other contexts, minikits of similar size relative to median landholdings have been shown to increase adoption of new crop varieties or varieties with low levels of adoption (Emerick et al., 2016; Jones et al., 2021). To test for spillovers, such as social learning, Water User Groups were randomly assigned to 20%, 60%, or 100% minikit saturation, with re-randomization for balance on Zone and O&M treatment status. Minikits were offered to assigned individuals prior to 2017 Rainy 1 and 2017 Dry.

We estimate the impacts of minikits using Equation 7; we define treatment as assignment to receive a minikit, and include Zone fixed effects, O&M treatment assignment, minikit saturation, and dummies for the number of minikit lotteries the household was entered into as controls.[19]

[19] After matching names from the lists of Water User Group members used for the minikit randomization to our baseline survey, we found 32% of households were listed multiple times.

Column 3 of Figure 5 presents the primary results of this experiment. In the top panel, we see that minikit distribution was quite successful: households assigned a minikit are 40pp more likely to take up the minikit than households that are not. However, the bottom panel indicates that households assigned a minikit are not more likely to adopt horticulture or irrigate: minikits increase horticulture adoption by 3.5pp and decrease irrigation use by 0.5pp, with 95% confidence intervals of (-4.5pp, 11.5pp) and (-8.3pp, 7.3pp), respectively. These confidence intervals reflect sufficient precision to reject estimates from other contexts of the effect of minikits on technology adoption (Emerick et al., 2016; Jones et al., 2021). We also find no evidence that saturating minikits increase irrigation: in fact, the point estimate suggests that farmers in densely saturated zones irrigate their farms less. We conclude that neither providing minikits nor saturating minikits lead to additional irrigation use.

# VII Conclusion

This paper provides evidence that irrigation has the potential to be a transformative technology in sub-Saharan Africa. Using data from very proximate plots which receive differential access to irrigation, we show the construction of an irrigation system leads to a 43% - 62% increase in cash profits. These profits are generated by a switch in cropping patterns from bananas towards a more input-intensive rotation of dry-season horticulture and rainy-season staples.

At the same time, four years after introduction we observe only a minority of farmers adopting this technology that is provided to them for free. We further document that frictions in land and labor markets cause inefficient adoption of irrigation. This result provides novel evidence that separation failures in agricultural household production lead to land misallocation and inefficient adoption of a new technology in Rwanda. This result has stark policy relevance: without greater adoption, these irrigation systems will not be able to generate sufficient revenue to be sustainable.

While our results highlight the presence of constraints on land and labor markets and demonstrate those constraints generate inefficiencies in an important technology adoption context, they can not provide evidence as to whether those inefficiencies lead to too much or too little adoption of irrigation. If farmers faced no constraints in the amount of labor they could sell at the market wage, the model would suggest they would irrigate even less; indeed, based on the descriptive statistics in Section

C, the model would predict nearly all of these farmers would either cultivate low-intensity bananas or exit farming altogether. Of course, a labor supply shift of this magnitude is likely to put downward pressure on the wage. If the ability to sell labor without frictions led to a substantial reduction in the wage rate, we may see many more farmers hiring labor and cultivating horticulture on more of their plots: this shift could allow irrigation systems to realize their transformative potential.

These results underscore the need for more evidence on both the role of factor markets in technology adoption, and the identification of particular institutions which contribute to or which can smooth those market failures. In some cases, these market failures may pose a competing constraint which coexists with other, more conventional constraints to production: if frictions in factor markets similarly constrain adoption of new technologies in other environments, then incomplete factor markets may limit the effectiveness of infrastructure, financial, or information interventions in improving agricultural productivity. This is a fruitful area for future research.

# References


**Adamopoulos, Tasso and Diego Restuccia**, "The size distribution of farms and international productivity differences," *American Economic Review*, 2014, *104* (6), 1667–97.

_ **and** _ , "Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data," Working Paper 24532, National Bureau of Economic Research April 2018.

_ **, Loren Brandt, Jessica Leight, and Diego Restuccia**, "Misallocation, Selection and Productivity: A Quantitative Analysis with Panel Data from China," Working Paper 23039, National Bureau of Economic Research January 2017.

**Agness, Daniel, Travis Baseler, Sylvain Chassang, Pascaline Dupas, and Erik Snowberg**, "Valuing the Time of the Self-Employed," 2020.

**Benjamin, Dwayne**, "Household composition, labor markets, and labor demand: testing for separation in agricultural household models," *Econometrica: Journal of the Econometric Society*, 1992, pp. 287–322.

**Bleakley, Hoyt and Joseph Ferrie**, "Land openings on the georgia frontier and the coase theorem in the short-and long-run," 2014.

**Breza, Emily, Nandita Krishnaswamy, and Supreet Kaur**, "Scabs: The Social Suppression of Labor Supply," 2018.

**Conley, Timothy G**, "GMM Estimation with Cross-Sectional Dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

_ **and Christopher R Udry**, "Learning about a new technology: Pineapple in Ghana," *The American Economic Review*, 2010, pp. 35–69.

**Dillon, Andrew**, "The effect of irrigation on poverty reduction, asset accumulation, and informal insurance: Evidence from Northern Mali," *World Development*, 2011, *39* (12), 2165–2175.

_ **and Ram Fishman**, "Dams: Effects of Hydrological Infrastructure on Development," *Annual Review of Resource Economics*, 2019, *11.*

**Dillon, Brian and Christopher B Barrett**, "Agricultural factor markets in Sub-Saharan Africa: An updated view with formal tests for market failure," *Food policy*, 2017, *67*, 64–77.

_ **, Peter Brummund, and Germano Mwabu**, "Asymmetric non-separation and rural labor markets," *Journal of Development Economics*, 2019, *139*, 78–96.

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

**Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H Dar**, "Technological innovations, downside risk, and the modernization of agriculture," *American Economic Review*, 2016, *106* (6), 1537–61.

**Fafchamps, Marcel**, "Sequential Labor Decisions Under Uncertainty: An Estimable Household Model of West-African Farmers," *Econometrica*, 1993, *61*, 1173–1197.

**Fei, John C.H. and Gustav Ranis**, "A Theory of Economic Development," *The American Economic Review*, 1961, *51* (4), 533–565.

**Foster, Andrew and Mark Rosenzweig**, "Are there too many farms in the world? Labor-Market Transaction Costs, Machine Capacities, and Optimal Farm Size," *NBER Working Paper*, 2017, (No. 23909).

**Foster, Andrew D and Mark R Rosenzweig**, "Technical change and human-capital returns and investments: evidence from the green revolution," *The American economic review*, 1996, pp. 931–953.

**Goldstein, Markus and Christopher Udry**, "The profits of power: Land rights and agricultural investment in Ghana," *Journal of political Economy*, 2008, *116* (6), 981–1022.

**Gollin, Douglas and Christopher Udry**, "Heterogeneity, measurement error and misallocation: Evidence from african agriculture," *NBER Working Paper 25440*, 2019.

**Jacoby, Hanan G**, "“Well-fare” Economics of Groundwater in South Asia," *The World Bank Research Observer*, 2017, *32* (1), 1–20.

**Janvry, Alain De, Elisabeth Sadoulet, and Tavneet Suri**, "Field experiments in developing country agriculture," in "Handbook of Economic Field Experiments," Vol. 2, Elsevier, 2017, pp. 427–466.

**Jones, M., F. Kondylis, J. Loeser, A. M. Mobarak, and D. Stein**, "Evaluating the Integrated Agriculture Productivity Project in Bangladesh," 2021.

**Jones, Maria, Florence Kondylis, Loeser John, and Jeremy Magruder**, "Data and Code for: Factor Market Failures and the Adoption of Irrigation in Rwanda," Inter-university Consortium of Political and Social Research [distributor] 2022.

**Kaur, Supreet**, "Nominal wage rigidity in village labor markets," *NBER Working Paper 20770*, 2014.

**Key, Nigel, Elisabeth Sadoulet, and Alain De Janvry**, "Transactions costs and agricultural household supply response," *American journal of agricultural economics*, 2000, *82* (2), 245–259.

**LaFave, Daniel and Duncan Thomas**, "Farms, families, and markets: New evidence on completeness of markets in agricultural settings," *Econometrica*, 2016, *84* (5), 1917–1960.

**Lewis, W Arthur**, "Economic development with unlimited supplies of labour," *The manchester school*, 1954, *22* (2), 139–191.

**Magruder, Jeremy**, "High unemployment yet few small firms: The role of centralized bargaining in South Africa," *American Economic Journal: Applied Economics*, 2012, *4* (3), 138–66.

_ , "Can minimum wages cause a big push? Evidence from Indonesia," *Journal of Development Economics*, 2013, *100* (1), 48–62.

**Ostrom, Elinor**, *Governing the commons: The evolution of institutions for collective action*, Cambridge university press, 1990.

**Sekhri, Sheetal**, "Wells, water, and welfare: the impact of access to groundwater on rural poverty and conflict," *American Economic Journal: Applied Economics*, 2014, *6* (3), 76–102.

**Sen, Amartya K**, "Peasants and Dualism with or without Surplus Labor," *Journal of political Economy*, 1966, *74* (5), 425–450.

**Singh, Inderjit, Lyn Squire, and John Strauss**, *Agricultural household models: Extensions, applications, and policy*, The Johns Hopkins University Press, 1986.

**Smith, Cory**, "Land Concentration and Long-Run Development: Evidence from the Frontier United States," 2019.

**Suri, Tavneet**, "Selection and comparative advantage in technology adoption," *Econometrica*, 2011, *79* (1), 159–209.

**Udry, Christopher**, "Efficiency and market structure: testing for profit maximization in African agriculture," 1997.

**World Bank**, "World Development Report 2008: Agriculture for Development," 2007.