

NBER WORKING PAPER SERIES

FACTOR MARKET FAILURES AND THE ADOPTION OF IRRIGATION IN RWANDA

Maria Jones
Florence Kondylis
John Loeser
Jeremy Magruder

Working Paper 26698
<http://www.nber.org/papers/w26698>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020

This draft benefited from comments from Chris Barrett, Paul Christian, Alain de Janvry, Simeon Djankov, Esther Duflo, Andrew Foster, Doug Gollin, Saahil Karpe, Elisabeth Sadoulet, John Strauss, Duncan Thomas, Chris Udry, and seminar audiences at Cornell University, Georgetown University, Michigan State University, North Carolina State University, World Bank, Northwestern University, University of California, Berkeley, University of Southern California and Y-RISE. 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. The views expressed in this manuscript do not reflect the views of the World Bank or the National Bureau of Economic Research. All errors are our own.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Maria Jones, Florence Kondylis, John Loeser, and Jeremy Magruder. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Factor Market Failures and the Adoption of Irrigation in Rwanda
Maria Jones, Florence Kondylis, John Loeser, and Jeremy Magruder
NBER Working Paper No. 26698
January 2020
JEL No. O1,O12,O13,Q12,Q15

ABSTRACT

We examine constraints to adoption of new technologies in the context of hillside irrigation schemes in Rwanda. We leverage a plot-level spatial regression discontinuity design to produce 3 key results. First, irrigation enables dry season horticultural production, which boosts on-farm cash profits by 70%. Second, adoption is constrained: access to irrigation causes farmers to substitute labor and inputs away from their other plots. Eliminating this substitution would increase adoption by at least 21%. Third, this substitution is largest for smaller households and wealthier households. This result can be explained by labor market failures in a standard agricultural household model.

Maria Jones
World Bank
1818 H Street NW
Washington, DC 20415
mjones5@worldbank.org

John Loeser
World Bank
1818 H Street NW
Washington, DC 20415
jloeser@worldbank.org

Florence Kondylis
World Bank
1818 H Street NW
Washington, DC 20415
fkondylis@worldbank.org

Jeremy Magruder
Department of Agricultural and Resource Economics
University of California at Berkeley
207 Giannini Hall
Berkeley, CA 94720-3310
and NBER
jmagruder@berkeley.edu

A data appendix is available at <http://www.nber.org/data-appendix/w26698>

A randomized controlled trials registry entry is available at <https://www.socialscienceregistry.org/trials/1323>

Contents

1	Introduction	3
2	Data and context	9
2.1	Irrigation in Rwanda	9
2.2	Data	11
2.2.1	Spatial sampling	11
2.2.2	Survey	12
2.3	Stylized facts	14
3	Impacts of irrigation	16
3.1	Empirical strategy	16
3.1.1	Balance	19
3.2	Estimating the effects of irrigation	20
3.2.1	Adoption Dynamics	20
3.2.2	Impacts of irrigation	21
3.2.3	Discussion of spillovers	23
4	Testing for binding constraint	25
4.1	Model	25
4.2	A test for separation failures	27
4.3	Separating constraints	30
5	Separation failures and adoption of irrigation	33
5.1	Empirical strategy	33
5.1.1	Balance	34
5.2	Results	35
5.2.1	A test for separation failures	35
5.2.2	Impacts of separation failures on adoption of irrigation	38
5.2.3	Separating constraints	39
6	Experimental evidence	40
6.1	Empirical strategy and results	42
7	Conclusions	43

1 Introduction

Limited adoption of productive technologies is a prominent explanation of low agricultural productivity in sub-Saharan Africa (World Bank, 2007). Existing productive technologies are underutilized due to inefficiencies in the markets faced by farmer households (Udry, 1997). A recent literature has provided robust evidence that these market failures limit 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 & 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 & Ranis, 1961; Benjamin, 1992). More recent empirical work has found that these inefficiencies are quantitatively important (Udry, 1997; Adamopoulos & Restuccia, 2014; Adamopoulos et al., 2017; Foster & Rosenzweig, 2017; Adamopoulos & Restuccia, 2018).

In this paper, we demonstrate that incomplete land and labor markets contribute to the productivity gap by limiting technology adoption.¹ We do so in the context of a potentially transformative technology: irrigation. Irrigation increases agricultural productivity in several ways: it adds additional agricultural seasons, enables cultivation of water-intensive crops, and reduces production uncertainty. However, irrigation is also costly: it requires large construction and maintenance costs, and is associated with increased usage of complementary inputs, such as labor, fertilizer, and improved seeds. Market failures, including in factor markets, therefore have the potential to cause inefficient levels of irrigation adoption

¹A related question is explored in papers which evaluate the effects of land titling and other formalized property rights on farm investment (Besley, 1995; Goldstein & Udry, 2008; Deininger & Feder, 2009; Besley & Ghatak, 2010; Ali et al., 2014; Goldstein et al., 2018). In our context, farmers have been assigned formal titles to our plots and so we identify the influence of factor market frictions on technology adoption in the presence of formalized rights. Our emphasis on the role of labor market frictions is also distinct.

as they induce a wedge between shadow prices and market prices of these inputs.

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 irrigation in Rwanda. We identify these returns using a plot-level spatial discontinuity design in newly constructed hillside irrigation schemes. We sample plots within 50 meters of gravity fed canals, which originate from a distant water source and must maintain a consistent gradient along the hillside. We survey 969 cultivators on 1,753 plots for 4 years.² 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 44% - 71% 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 & Pande \(2007\)](#) in India.³ 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.⁴

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

²These numbers are only for the sample of households whose sampled plot is within 50 meters of the associated discontinuity; in full we survey 1,695 cultivators on 3,332 plots.

³Existing 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 & Pande, 2007](#)), aquifer characteristics ([Sekhri, 2014](#); [Loeser, 2020](#)), 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 & Fishman \(2019\)](#) review the literature on the impacts of surface water irrigation infrastructure.

⁴This is distinct from the collective action failures discussed in ([Ostrom, 1990](#)). Low adoption of irrigation as a threat to sustainability has also been documented by [Attwood \(2005\)](#), who argues that cost recovery was a challenge for canal irrigation systems in nineteenth and early twentieth century India until the introduction of sugarcane.

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 44% - 71% to -12% - 38%. Valuing household labor at the market wage may not be appropriate: rural market wages are likely to be inefficiently 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 & Thomas, 2016). Heterogeneity in the shadow wage would then cause inefficient adoption of irrigation across households.⁵ 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 derive a test for inefficient adoption of irrigation caused by market failures. To produce this test, we build on seminal agricultural household models (Singh et al., 1986; Benjamin, 1992) and model households' production decisions, incorporating 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 as follows. 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 substitution of labor and inputs away from other plots.⁶ 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 caused by market failures, exploiting the

⁵This heterogeneity could only exist if there were frictions in at least one other market in addition to labor markets.

⁶The mechanism is straightforward: access to irrigation on one plot increases input use on that plot. That increase does not affect input demand on the farmers' other plots; however, if the farmer faces binding constraints in input, risk, or labor markets, that increase in input use must be associated with a decrease in input use on other plots.

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: for farmers with a plot in the command area, an additional irrigated plot caused by access to irrigation is associated with a 68 percentage point decrease in the probability of irrigating the second 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 21% higher. Moreover, the presence of this substitution implies current adoption of irrigation is inefficient: different households make different adoption decisions on technologically identical plots because of their access to irrigation on their other plots.⁷

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 inefficiently low adoption in this context. To produce this test, 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, while many patterns of differential substitution are possible, 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

⁷With sufficient time, these sites could reach an equilibrium in which this misallocation would have slowly been corrected by markets (Gollin & 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 & Ferrie, 2014; Smith, 2019).

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 62% - 86% less than average size households, while one standard deviation wealthier households substitute 40% - 80% 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: usage fees 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 al., 2018). We find no effects of receiving minikits on adoption of horticulture in our context, in contrast to existing work. A closer analysis indicates that the farmers who take up the minikits are the same farmers who would have been likely to cultivate horticulture absent the intervention. Combining this evidence with the model-based test above, we conclude that financial and informational constraints are unlikely to be a primary explanation for low and inefficient adoption of irrigation.

This paper demonstrates that frictions in land and labor markets cause inefficient adoption of hillside irrigation in Rwanda. This result integrates key findings from three large literatures in development economics. First, our result provides some ground-level evidence for the mechanisms underlying misallocation (Adamopoulos & Restuccia, 2014; Adamopoulos et al., 2017; Foster & Rosenzweig, 2017; Adamopoulos & Restuccia, 2018). We document that land misallocation hinders technology adoption, and that frictions in labor markets are

one reason why land market failures generate production inefficiencies. The intuition for our test expands on a deep literature on separation failures which empirically demonstrates that factor market failures affect the allocation of land and labor across households (Singh et al., 1986; Benjamin, 1992; LaFave & Thomas, 2016; Dillon & Barrett, 2017; Dillon et al., 2019).⁸ Our context allows us to innovate by demonstrating that separation failures induce differential adoption of irrigation on technologically identical plots. In doing so, we also contribute to a literature leveraging production function estimates to document misallocation of labor and inputs by inferring their marginal products from their allocations across plots or households (Jacoby, 1993; Skoufias, 1994; Udry, 1996; Restuccia & Santaaulalia-Llopis, 2017).⁹ Our test for inefficient technology adoption caused by land and labor market failures therefore complements this literature, by both imposing less structure and leveraging our plot-level discontinuity in access to irrigation as an exogenous labor- and input-complementing productivity shock.

This paper is organized as follows. Section 2 describes the context we study and our sources of data. Section 3 presents our estimates of the impacts of irrigation in Rwanda. Section 4 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 5, and experimental tests in Section 6. Section 7 concludes.

⁸The existing literature does so by testing whether households with different characteristics use different levels of inputs; however, this type of test stops short of showing that these allocations are inefficient (Udry, 1997). In particular, it can only conclude that one market has failed; because it can not conclude that at least two markets have failed, by Walras' Law it is insufficient to demonstrate an inefficiency.

⁹Although demonstrating heterogeneity in the marginal product of labor is sufficient to show that labor market failures generate inefficiencies, the methods employed by this literature are typically not robust to the presence of unobserved heterogeneity across plots or measurement error (Gollin & Udry, 2019).

2 Data and context

2.1 Irrigation in Rwanda

We study 3 hillside irrigation schemes, located in Karongi and Nyanza districts of Rwanda, that were constructed by the government in 2014; 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 perennial bananas which are sold commercially;¹⁰ most farmers adopt either a rotation of staples, fallowing land in the dry season, or cultivate bananas.

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, even with irrigation, for households that cultivate during the two rainy seasons. Instead, cultivating shorter cycle horticulture during the dry season becomes a possibility with the availability of irrigation. 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.¹¹ As 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. Absent irrigation, horticulture is familiar but uncommon around these areas; at baseline 3.2% of plots outside of the command area are planted with at least some

¹⁰Staple 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.

¹¹Kigali is less than a 3 hour drive from these markets, facilitating trade.

horticulture, primarily during the rainy seasons.

In this context, 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). The schemes in our study share some common features; a picture from one of the schemes is presented in Figure 2. In each site, land was terraced in preparation for the irrigation works (as hillside irrigation would be infeasible on non-terraced land). Construction and rehabilitation of terraces in these sites began in 2009 - 2010. The schemes are all gravity fed, and use surface water as the source.¹² From these water sources, a main canal (visible in Figure 2) was constructed along a contour of the hillside; engineering specifications required 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. Underground pipes run down the terraces from the canal every 200 meters. Farmers draw water from 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 the plots which are below the canal and located within 100 meters of one of these valves.

In all sites, sufficient water is available to enable irrigation year-round. To the extent that there is heterogeneity in plot-level water pressure, the plots nearest to the canal face the lowest pressure.¹³ The primary cost to farmers of irrigating a plot in this context is their 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, there are no fees associated with the use of irrigation water¹⁴

¹²In two sites, a river provides the water source, while in the third site, a dammed lake is the source.

¹³The lower pressure on these plots is attributable to the design of the pipes, which fill up with water before valves are opened; forces of gravity and the lower volume of water in the pipes above the highest valves generates somewhat weaker pressure than at the lower valves (though pressure is still sufficient for effective irrigation). This difference in pressure could become more serious if lower valves were opened at the same time as higher valves; in practice, schedules of water usage are agreed upon to prevent this from happening.

¹⁴The government does have an objective of developing the financial self-sufficiency of the schemes. To

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 terraces just above the canal fall outside the command area and do not have access to irrigation.

2.2 Data

2.2.1 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 constructed this sample of plots by dropping a uniform grid of points across the site at 2-meter resolution, and then randomly sampling points within the grid within 50m of the command area boundary.¹⁵ After each point was sampled, we excluded all points within 10m of that point (to avoid selecting multiple points too close together).

Enumerators were then given GPS devices with the locations of the points, and sent to do so, land taxes are intended to be applied to the plots in the command area, which (as land taxes) should not influence cultivation decisions. These taxes are intended to be small in magnitude compared to potential farmer yields as they are meant to fund only ongoing operations and maintenance costs rather than full cost recovery; the highest fees across the sites were 77,000 RwF/ha/year, while our dry season treatment on the treated estimates presented in Section 3 are 300,000 - 450,000 RwF/ha. The first attempts to collect these taxes were made in 2017 Rainy 1. The survey team engaged in an experiment to test whether these taxes were a barrier to use of the irrigation system by randomizing subsidies across farmers at up to 100%; we do not find any evidence that the taxes changed farming practices (results available from authors). This is perhaps unsurprising as tax compliance was very low, with 4% of scheduled taxes collected from farmers who did not receive full subsidies from the research team.

¹⁵In 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 below in Section 6. Additionally, only two of the three sites have a viable boundary of cultivable land both just inside and just outside the command area; we use only these sites for our analysis of the impacts of access to irrigation in Section 3 and Section 5.

each point with a key informant (often the village leader). For each point, they were asked to identify if the point was on cultivable land (this was to discard forest, swamps, thick bushes, bodies of water, or other terrain which would make cultivation impossible). When a point fell on cultivable land, they recorded the name of the cultivator of the plot, their contact information, as well as a sufficiently detailed description of the plot. In the rest of this paper, we refer to all plots thus identified as *sample plots*. Our main household sample was built from this sampling procedure: the data from this listing was used to construct a roster of all the unique names of cultivators, eliminating duplicate names. Finally, for each household with points falling on multiple plots, one of these points was randomly selected to be that household’s sample plot.

2.2.2 Survey

Our baseline survey was implemented in August - October 2015 and includes detailed agricultural production data (season-by-season) for seasons 2014 Dry through 2015 Rainy 2, that is, spanning the year from 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. Details of the construction of key variables we use for the analysis are presented in Appendix A. As mentioned above, this is not a “true” baseline as some farmers had already gained access to irrigation in 2014 Dry. However, relatively small parts of the site had access to irrigation at this point; in Section 3.2.1 we highlight that 2014 Dry adoption of irrigation is less than 25% of adoption in subsequent dry seasons, and in Section 3.1.1 we show balance across the command area boundary in household and plot characteristics. Production and input data are collected plot-by-plot; in the baseline we conducted this production data for up to four plots, although subsequent surveys maintain a panel of two plots. Each of these plots was also mapped using GPS devices during the baseline; we use this data to construct the area of plots and their locations. 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

next most important plot (defined at baseline; we refer to this as the “most important plot”). We also collected data on household characteristics, labor force behavior, and a short consumption and food security module. In analysis, we will focus on the sample plots to learn about the effects of the irrigation itself, and the most important plot to learn about how the presence of 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. In each survey, we asked 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 sample for the follow up surveys consists of all the baseline respondents. To build a panel of households and plots, we interviewed households from the baseline and recorded information on all their baseline plots. Whenever a household’s sample plot or most important plot was sold or rented out to another household, or a household stopped renting in that plot if it was not the owner (“transacted”), we ran a “tracking survey”. Specifically, we tracked and interviewed 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 Table [A11](#). Only 6.0% (6.4%) of plot-by-season observations for sample plots outside the command area in our primary analysis sample (defined in Section [3.1](#)) 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 6.4 - 10.2pp of plots to be rented out to the 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, and we discuss attrition further in Appendix G.

2.3 Stylized facts

To motivate our analysis of the impacts of hillside irrigation, we first introduce some stylized facts about irrigation in this context. Table 1 presents summary statistics for agricultural production from our four years of data, pooled across seasons; Figure 3 presents a subset of these statistics graphically.

Stylized Fact 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, 93% of plots cultivated with horticulture in the dry season use irrigation. This stylized fact makes agronomic sense as the rainfall in rainy seasons in this part of Rwanda is usually sufficient for either staple or horticultural production (and in wet years may be harmfully excessive for horticulture). Additionally, as staples do not have a sufficiently short cycle to permit cultivation during the relatively short dry season (while horticulture does), it is not agronomically feasible to use irrigation to cultivate staples during the dry season.

Stylized Fact 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 420 days/ha of household labor, 60 days/ha of hired labor, and 50,000 RwF/ha of inputs, regardless of the season in which it is planted.¹⁶ This contrasts to staple plots (260 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).

Stylized Fact 3. *Horticultural production produces much higher cash profits than other forms of agriculture.*

Horticultural production produces 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.

Stylized Fact 4. *Household labor is the primary input to production of any crop, and the economic profitability of horticulture depends critically on the shadow wage.*

A large existing literature examines separation failures in labor markets faced by agricultural households (e.g., [Singh et al. \(1986\)](#); [Benjamin \(1992\)](#); [LaFave & Thomas \(2016\)](#)). 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 (though both appear more profitable than cultivating staples).¹⁷ As a result, ultimately the economic profitability of horticulture relative to bananas will depend critically on the constraints on household labor supply decisions.

¹⁶For reference, in the study period, the exchange rate was approximately 800 RwF = 1 USD

¹⁷Both 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.

3 Impacts of irrigation

3.1 Empirical strategy

We start our analysis through a simple OLS framework, and we restrict this and subsequent analysis to sample plots within 50 meters of the discontinuity. If these nearby plots are sufficiently similar so that irrigation access can be taken as random within this sample, we can simply regress

$$y_{1ist} = \beta_0 + \beta_1 CA_{1is} + \alpha_{st} + \epsilon_{1ist} \quad (1)$$

Where y_{kist} is outcome y for plot k of household i located in site s in season t , CA_{kis} is an indicator for that plot being in the command area, and α_{st} are site-by-season fixed effects meant to control for any differences or trend differences across sites (including market access or prices). We use $k = 1$ to indicate the household's sample plot, as opposed to the household's most important plot.

Next, we consider two primary potential sources of omitted variable bias. First, plots that are positioned relatively higher on the hillside may have different agronomic characteristics, and accordingly farmers may differentially sort into these plots. As plots inside the command area are lower on the hillside (below the canal) and plots outside the command area are higher on the hillside (above the canal), the command area indicator will be correlated with position on the hillside and β_1 may be biased. Second, as the construction of the canal slices through plots on the hillside, this may differentially change the area of plots that are positioned higher or lower on the hillside. 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. As we anticipate this will cause plots to be relatively larger just inside the command area, and plots exhibit strong evidence of diminishing returns to scale in this context, this effect will likely bias β_1 downwards.

We account for these two potential sources of omitted variable bias by including controls. First, to account for position on the hillside, we control for distance of the plot from

the command area boundary, and distance of the plot from the command area boundary interacted with the command area indicator.¹⁸ This is a standard regression discontinuity specification, and as such compares sample plots that are just inside the command area to sample plots that are just outside the command area. Second, to account for differences in area of plots, we control for the log area of sample plots. Specifically, we estimate

$$y_{1ist} = \beta_0 + \beta_1 CA_{1is} + \beta_2 Dist_{1is} + \beta_3 Dist_{1is} * CA_{1is} + \alpha_{st} + \gamma X_{1is} + \epsilon_{1ist} \quad (2)$$

where $Dist_{1is}$ is the distance of plot 1 from the command area boundary (positive for plots within the command area, negative for plots outside the command area) and X_{1is} is the log plot area.

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.¹⁹ Second, marginal plots which would have been too unproductive to cultivate absent irrigation, and would 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 rental prices in our data are 150,000 RwF/ha, approximately 25% of annual yields.

We account for this potential source of bias using spatial fixed effects (SFE; see [Gold-](#)

¹⁸We calculate distance using the distance of the plot boundary to the command area boundary.

¹⁹Typically, forests were planted or preserved in areas of low productivity, where the slope of the hillside was relatively high and erosion was relatively common. Therefore, this amounts to selection out of our sample of low productivity plots outside the command area, which would bias β_1 downwards.

stein & Udry (2008); Conley & 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.²⁰

In practice, we define a set \mathcal{N}_{kist} to be the group of five closest plots to plot k observed in season t , including the plot itself. Then, for any variable z_{kist} , define $\bar{z}_{kist} = (1/|\mathcal{N}_{kist}|) \sum_{k' \in \mathcal{N}_{kist}} z_{k'ist}$. The SFE specification then estimates

$$y_{1ist} - \bar{y}_{1ist} = \beta_1(\text{CA}_{1is} - \overline{\text{CA}}_{1is}) + (V_{1is} - \bar{V}_{1is})'\gamma + (\epsilon_{1ist} - \bar{\epsilon}_{1ist}) \quad (3)$$

where V_{kis} includes all controls from Equation 2, 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, 49% of plots have 3 plots (self inclusive) within 50 meters, and 87% have 3 plots within 100m; 60% of plots have all 5 plots (self inclusive) within 100m, while 83% have all 5 plots within 150m. As reference, Conley & Udry (2010) use 500m as the bandwidth for their estimator, while Goldstein & 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 specifications (1) and (2), we cluster standard errors at the level of the

²⁰Formally, SFE estimators leverage the identification assumption $\lim_{\|k-k'\| \rightarrow 0} E[\epsilon_{kist}|X_{kist}] = E[\epsilon_{k'ist}|X_{k'ist}]$, where $\|k-k'\|$ represents the distance between plot k and plot k' .

nearest water user group, the group of plots that can source water from the same secondary pipe. When estimating specification (3), the spatial fixed effects generate correlation between the errors of close observations. To allow for this, we calculate Conley (1999) standard errors.²¹

3.1.1 Balance

We now use specifications (1), (2), and (3) to 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, 5, and 6), and our preferred specifications which we use in our analysis with key controls included (Columns 4, 7, and 8).

First, in specifications which control for distance to the boundary (Columns 5 through 8, Table 2), our sample plots are balanced in terms of ownership and rentals. Additionally, the vast majority 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 3.1, 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 3.2 on horticulture, input use, labor use, and yields, as all of these variables are larger in smaller plots in both the command area and outside the command area. Additionally, as suggested in Section 3.1, we find some additional imbalance on duration of plot ownership when the important control for distance to the boundary is omitted in Columns 3 and 4.²² We therefore present

²¹Specifically, we allow plot ℓ managed by household j and plot ℓ' managed by household j' to have correlated errors if there exists a plot k such that $\ell \in \mathcal{N}_{kist}$ or $k \in \mathcal{N}_{\ell jst}$, and $\ell' \in \mathcal{N}_{kist}$ or $k \in \mathcal{N}_{\ell' j'st}$.

²²We note that this imbalance goes the opposite direction suggested by the concern that the construction of the command area caused an increase in transactions before our baseline. This, combined with the coefficient dropping to 0 with the inclusion of controls, indicates that this imbalance is caused by relative position on the hillside and not by the command area. In fact, as shown in Table A11, we do find in follow up surveys that the command area causes an increase in rentals out to other farmers. However, as discussed in Appendix G, because we tracked plots across transactions, this did not lead to differential attrition and

results estimated using Equation (1), which does not control for distance to the boundary or log area, and using Equations (2) and (3), which do control for distance to the boundary and log area.

Following the ownership results, Table 3 examines the characteristics of households whose sample plots are just inside or just outside the command area. First, note that Column 1, which does not restrict to the discontinuity sample, performs poorly here; we find significant imbalance on half of our variables, and the omnibus test rejects the null of balance. However, we fail to reject balance for our preferred specifications (Columns 4, 7, and 8, Table 3) which restrict to the discontinuity sample; households with sample plots just inside the command area appear similar to households with sample plots just outside the command area. In Column 5, there are significant differences in whether the household head is female is the age of the household head, and in Column 7, there is a significant difference in whether the household head has completed primary schooling or not. We note that 1 out of 10 variables significant at the 10% level is what one would expect due to chance.

Lastly, in Section 5.1.1, we consider the characteristics of households' most important plots; we show that these appear similarly balanced.

3.2 Estimating the effects of irrigation

3.2.1 Adoption Dynamics

Figure 4 presents the share of plots irrigated by season for sample plots just inside the command area and sample plots outside the command area. First, as the irrigation sites were already partially online in our baseline, we already observe some increased adoption of irrigation in the command area in 2014 Dry: sample plots in the command area are approximately 5pp more likely to be irrigated than sample plots outside the command area. We present results from 2014 Dry and 2015 Rainy 1 and 2 in Appendix F; consistent with this low adoption, we do not find significant impacts of access to irrigation on inputs or

therefore does not bias our results.

output in these seasons. Second, starting with 2015, adoption of irrigation does not appear to trend, but exhibits meaningful seasonality. Differences remain around 3pp - 6pp in the rainy seasons, and 19pp - 26pp in the dry seasons.

Given the limited changes in adoption dynamics after 2014 and the stark differences in adoption across dry and rainy seasons, for the remainder of our analysis we estimate (1), (2), and (3) pooling across our three years of follow up surveys, splitting our results across dry and rainy seasons.

3.2.2 Impacts of irrigation

We now present our results on the impact of access to irrigation on crop choices, on input use, and on production. First, we present graphical evidence of the regression discontinuity in Figure 5; for parsimony, we do so only for the dry seasons (2016 Dry, 2017 Dry, and 2018 Dry).²³ In each of the regression discontinuity figures, distance to the canal in meters is represented on the x-axis, with a positive sign indicating that the plot is on the command area side of the boundary. Second, we present regression evidence in Tables 4, 5, and 6. In the discussion below, we focus on results from the tables, but we note that these results are consistent with visual intuition from Figure 5.

First, in line with results from Section 3.2.1, command area plots are 16pp - 20pp 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. This transition to dry season horticulture substitutes for cultivation of perennial bananas, a less productive but less input intensive commercial crop; we estimate a decrease of 13pp - 17pp in the command area, and as a consequence we observe no impacts on overall cultivation in the dry season.²⁴

²³Rainy season differences are always smaller and generally not visually noteworthy; we focus most of our discussion on the dry season results.

²⁴As bananas are perennials, plots cultivated with bananas typically have harvests in each season. In contrast, the rotations of staples and horticulture (or simply horticulture) that replace bananas may only

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 perennial bananas, which require little inputs and labor, into horticulture, which is highly input and labor intensive. 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 340 - 450 person-days/ha, 25,000 - 39,000 RwF/ha, and 19,000 - 28,000 RwF/ha, respectively; these numbers are similar to differences 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 280,000 - 360,000 RwF/ha, an order of magnitude larger than the effects on input expenditures or hired labor expenditures. Additionally, as reference, 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, consistent with our estimates of impacts on input use, we find large increases in dry season agricultural production. Treatment on the treated analysis suggests adoption of irrigation increases yields by 300,000 - 450,000 RwF/ha, 49 - 72% of annual agricultural production. As horticulture is primarily commercial: each 1 RwF/ha increase in yields is associated with a 0.76 - 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. Additionally, these impacts on yields are much larger than our estimates of impacts on input and hired labor expenditures; our results suggest irrigation increases yields net of expenditures by 250,000 - 390,000 RwF/ha, a 44 - 71% increase in annual yields net of expenditures. However, we should not interpret this as impacts on profits, as it implicitly places no value on the large increases in household labor. If we instead value household

involve two plantings and harvests, and we therefore see a modest decrease in cultivation during the rainy seasons of 5pp - 9pp on a baseline of 84%.

labor at 800 RwF/person-day, the median wage we observe, these impacts vanish completely. Therefore, the profitability of the transition to dry season horticulture enabled by irrigation depends crucially on the shadow wage at which household labor is valued.²⁵

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 44 - 71% increases in yields net of expenditures that we document from just three months of cultivation. At the same time, these results also 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 tests for the role of market failures in adoption of irrigation.

3.2.3 Discussion of spillovers

There are three categories of spillovers we anticipate in our context – across household spillovers, within plot (across season) spillovers, and within household (across plot) spillovers.

The across household spillovers we anticipate are general equilibrium effects. First, as access to irrigation increases demand for labor, we expect this to put upward pressure on wages. Second, as access to irrigation increases sales of horticulture, we expect the prices of horticultural crops to decrease.²⁶ Although our discontinuity design does not allow us

²⁵In Appendix B, 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.

²⁶We do not anticipate general equilibrium effects in other markets would meaningfully affect our results. First, agricultural inputs are significantly more tradable than horticultural output, and we therefore do not anticipate effects on agricultural input prices. Second, although we do also find an increase in land transactions caused by access to irrigation in Appendix G, only a small share of plots that receive access to irrigation are transacted. We therefore anticipate any general equilibrium effects through land markets to be small.

to estimate these general equilibrium effects, in Appendix C we plot wages and prices of agricultural output by season to see if we observe any increase in wages or decrease in the price of horticulture. While we find no evidence that wages or staple prices changed after the hillside irrigation schemes became operational, we find some suggestive evidence that prices of horticultural crops decreased in one of the sites.²⁷

To interpret the impact of these potential price effects, first note that, abstracting from within household spillovers, our estimates are the partial equilibrium effect of access to irrigation conditional on observed wages and prices. Second, note that essentially all dry season horticultural production in these sites is on irrigated plots. If horticulture prices decline in these sites in response to irrigation, we therefore interpret that effects of irrigation access on revenues and profits are smaller than they would be with frictionless trade.

The within plot (across season) spillovers we anticipate are driven by the shift out of perennial bananas, which causes a change in patterns of cultivation during the rainy season, while adoption of irrigation is primarily during the dry season. We estimate these spillovers and discuss further in Section 3.2.2; we do not find strong evidence of impacts on rainy season labor, inputs, yields, or profits.

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. If households face constraints, this spillover may be first order in our context and would generate inefficiency in technology adoption. To address this, in Section 4 we model a household's agricultural production decisions and how they can generate substitution across plots, and in Section 5 we estimate these spillovers and quantify their implications for our estimates and for efficiency.

²⁷In section 5 we will provide evidence that farmers are optimizing labor inputs with respect to shadow wages rather than market wages. This may explain why market wages do not respond to the increased labor demand from irrigation.

4 Testing for binding constraint

4.1 Model

Farmers have 2 plots, indexed by k : $k = 1$ indicates the sample plot, while $k = 2$ indicates the most important plot. On each plot k , they have access to a simple production technology $\sigma A_k F_k(M_k, L_k)$ where A_k is plot productivity, M_k is the inputs applied to plot k and L_k is the household labor applied to plot k . The common production shock σ is a random variable such that $\sigma \sim \Psi(\sigma)$, $E[\sigma] = 1$.²⁸ While this specification assumes a single production function on each plot, we can think of $F_k(M_k, L_k)$ as the envelope of production functions from cultivating different fractions of bananas and horticulture on the dry season; thus we will think of cultivating bananas as optimizing at a low input intensity. Utilizing subscripts to indicate partial derivatives and subsuming arguments we assume $F_{kM} > 0$, $F_{kL} > 0$, $F_{kML} > 0$, $F_{kMM} < 0$, $F_{kLL} < 0$.²⁹ Farmers have a budget of \bar{M} which, if not utilized for inputs, can be invested in a risk-free asset which appreciates at rate r . In this context, farmers maximize expected utility over consumption and leisure l , considering their budget constraint and a labor constraint \bar{L} which is allocated to labor on each plot, leisure, and up to \bar{L}^O units of off farm labor L^O .³⁰ Finally, we model irrigation access as an increase in A_1 . As we consider the role of each different constraint, we develop the necessary assumptions to produce the results from Section 3: that this increase in A_1 generates an increase in demand for inputs and labor on plot A_1 .

²⁸While we refer to σ as a production shock, this incorporates general uncertainty in the value of production which includes joint price and production risk.

²⁹Among these, $F_{kML} > 0$ is the most controversial. Existing evidence on F_{kML} in developing country agriculture is mixed (see Heisey & Norton (2007) for discussion). In our context, we expect $F_{kML} > 0$ primarily because $F_k(\cdot, \cdot)$ encompasses the transition from bananas to horticulture, which should be associated with increased input demands according to Stylized Fact 2.

³⁰We follow Benjamin (1992) in modeling incomplete labor markets as driven by an off farm labor constraint. As in Benjamin (1992), we do so to match the observation that rural wages appear to be higher than the productivity of on-farm labor. However, for the predictions that follow it is sufficient that households face a downward sloping labor demand curve.

Farmers maximize expected utility

$$\max_{M_1, M_2, L_1, L_2, l, L^O} E[u(c, l)]$$

subject to the constraints enumerated above

$$\begin{aligned} \sigma A_1 F(M_1, L_1) + \sigma A_2 F(M_2, L_2) + wL^O + r(\bar{M} - M_1 - M_2) &= c \\ M_1 + M_2 &\leq \bar{M} \\ L_1 + L_2 + l + L^O &= \bar{L} \\ L^O &\leq \bar{L}^O \end{aligned}$$

In this framework, there are three crucial constraints farmers may face that cause deviations from expected profit maximization: access to insurance may be limited, reducing input use to avoid basis risk; credit or access constraints may limit input use; and farmers' off farm labor allocations may be constrained from above, resulting in overutilization of labor on the household farm. In analyzing model predictions we discuss the cases in which each of these constraints do or do not bind.³¹

After substituting in the constraints which bind with equality, we derive the following first order conditions³²

$$(M_k) \quad \left(1 + \frac{\text{cov}(\sigma, u_c)}{\mathbf{E}[u_c]}\right) A_k F_{kM} = (1 + \lambda_M)r \quad (4)$$

$$(L_k) \quad \left(1 + \frac{\text{cov}(\sigma, u_c)}{\mathbf{E}[u_c]}\right) A_k F_{kL} = (1 - \lambda_L)w \quad (5)$$

$$(\ell) \quad \frac{\mathbf{E}[u_\ell]}{\mathbf{E}[u_c]} = (1 - \lambda_L)w \quad (6)$$

Intuitively, the first order conditions for inputs and labor include three parts. First, each

³¹These constraints correspond with those most commonly cited by farmers in focus groups as driving crop choice. In particular, farmers frequently cite *imbaraga*, or strength, of the household head (corresponding to labor market constraints), *igishoro*, or access to capital (corresponding to credit or input market constraints), and *isoko*, or access to markets (corresponding to price risk, or insurance market constraints).

³²The derivation is in Appendix D.

contains the marginal product of the factor, $A_k F_{kM}$ and $A_k F_{kL}$ respectively, on the left hand side, and the market price of the factor, r and w respectively, on the right hand side. The second piece, $1 + \frac{\text{cov}(\sigma, u_c)}{\mathbf{E}[u_c]}$, is the ratio of the marginal utility from agricultural production to the marginal utility from certain consumption. This ratio scales down the marginal product of the factor. It is less than 1 because agricultural production is uncertain, and higher in periods in which marginal utility is lower, so $\text{cov}(\sigma, u_1) < 0$. With perfect insurance, $\text{cov}(\sigma, u_1) = 0$, and this piece disappears. Without it, however, farmers will underinvest in both inputs and labor relative to the perfect insurance optimum.³³ Third, there are the Lagrange multipliers associated with the input constraint λ_M and with the labor constraint λ_L , which scale the associated factor prices up and down, respectively.

When these constraints do not bind, and with perfect insurance, we have the familiar result that marginal products equal marginal prices. However, if any of these constraints bind, then separation fails: farmer characteristics which are related to λ_L , λ_M , or $\text{cov}(\sigma, u_1)$ will be correlated with inefficient input allocation on all plots (inefficiently low in the case of inputs and inefficiently high in the case of labor).

4.2 A test for separation failures

In this context, we consider a new test of separation: the effect of a change in access to irrigation on the sample plot on production decisions on the most important plot. Much of the literature that tests for separation, building on Benjamin (1992), has focused on tests built around the assumption that household characteristics should not affect the household's optimal production decisions under perfect markets. We instead leverage the assumption that access to irrigation on the sample plot (the "sample plot shock") should not affect the optimal production decisions on the household's most important plot.

Following our model, we show how these market failures in insurance, labor, or input

³³This result does not generically hold in models of agricultural households, as when consumption is separately modeled, households that are net buyers of an agricultural good may overinvest in inputs and labor relative to the perfect insurance optimum (Barrett, 1996). This is unlikely to be first order in our context, as we sampled cultivators and our results are driven by production of commercial crops.

markets generate a separation failure between production decisions on the sample plot and production decisions on the most important plot. First, we derive the classic separation result from Singh et al. (1986) in our framework when there are no market imperfections.

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

Showing this result is straightforward: with perfect markets for inputs, labor, and insurance, $\frac{\text{cov}(\sigma, u_c)}{\mathbf{E}[u_c]} = 0$, $\lambda_L = 0$, and $\lambda_M = 0$, respectively. The first order conditions then simplify to

$$\begin{aligned} (M_k) \quad A_k F_{kM} &= r \\ (L_k) \quad A_k F_{kL} &= w \\ (\ell) \quad \frac{\mathbf{E}[u_\ell]}{\mathbf{E}[u_c]} &= w \end{aligned}$$

The household's labor and input allocations on plot 2 depend only on plot 2 productivity A_2 , the price of inputs r , and the wage w , and not on access to irrigation on plot 1 (A_1).

In contrast to the case with perfect markets, in the presence of market failures, the sample plot shock can affect the households allocations on its most important plot. Roughly speaking, the sample plot shock increases the household's agricultural production, and increases its labor and input demands on the sample plot. When markets fail, this reduces the value the household places on agricultural production, and increases its opportunity costs of labor and inputs, and the household reduces its labor and input allocations on its most important plot. The following propositions require additional assumptions on the shape of the utility function or on the distribution of σ ; we highlight those in the text below each proposition.

Proposition 2. *If input, labor, or insurance constraints bind, then input and labor use are reduced on the most important plot in response to the sample plot shock.*³⁴

³⁴See proof in Appendix D.

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 most important plot. As inputs and labor are complements, this causes labor allocations on the most important plot to fall as well. Second, if labor constraints bind, then the increase in labor on the sample plot caused by access to irrigation must be associated with a reduction in the sum of leisure and labor on the most important plot. Under standard restrictions on the household's on farm labor supply, this must be associated with a reduction in labor on the most important plot.³⁵ As inputs and labor are complements, this causes input allocations on the most important plot to fall as well. Third, absent insurance, then the increase in agricultural production caused by access to irrigation reduces the marginal utility from agricultural production relative to the marginal utility from consumption.³⁶ In turn, this causes labor and input allocations to the most important 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 most important plot. Instead, both the sample plot and most important plot would flow to the household with the highest willingness-to-pay for them. In practice, land transactions do occur; as discussed in Section 2.2.2, 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 most important plot itself.

Proposition 2 produces a test of separation. Rejecting separation with this test 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 other constraints interact with land market frictions to generate separation failures. This is because

³⁵Specifically, we assume that leisure demand is increasing in consumption; this assumption is not necessary but is sufficient.

³⁶This does not generically hold; however, restrictions on the distribution of σ are sufficient to imply that marginal utility from agricultural production relative to the marginal utility from consumption is falling in agricultural production. Details are in Appendix D.

the presence of any set of constraints that generate separation failures yields the same prediction: the sample plot shock should cause input and labor allocations on the most important plot to fall. In particular, the intuition that observing changes in input allocations, labor allocations, or cropping decisions on the most important plot might suggest the presence of input constraints, labor constraints, or insurance constraints, respectively, fails, because inputs, labor, and horticulture are all complements in the production function.

4.3 Separating constraints

To shed light on which other constraints generate separation failures, we leverage the fact that our model offers predictions about how households with different characteristics should *differentially* respond to the sample plot shock. Roughly speaking, depending on which constraint binds, changes in different household characteristics may slacken or tighten the binding constraint. We focus on two important household characteristics in our model: we use household size to shift \bar{L} , the household's total available labor, and wealth to shift \bar{M} , the household's exogenous income available for input expenditures. We present these predictions below.

Proposition 3. *If input constraints or insurance constraints bind, then the input and labor allocations on the most important plot of larger households (wealthier households) should be less (less) responsive to the sample plot shock.³⁷*

Under insurance constraints, both wealth and household size enter the model symmetrically by increasing consumption; therefore, in all cases, wealthier and larger households will respond similarly to the sample plot shock. 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 most important plot will be less responsive to the sample plot shock.

³⁷See proof in Appendix D.

Under input constraints, wealthier households are less likely to see the constraint bind. As the allocations on the most important plot of unconstrained households do not respond to the sample plot shock, wealthier households should be less responsive. Now, note that in this model, farmers cannot use labor income to purchase additional inputs. In a more general model with borrowing, they may be able to; in that case, both wealthier households and larger households are less likely to see the constraint bind, and therefore will both be less responsive to the sample plot shock on their most important plots.³⁸

Proposition 4. *If labor constraints bind, then the relative responsiveness of input and labor allocations on the most important 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 most important plot of larger households (wealthier households) should be less (more) responsive to the sample plot shock.*³⁹

When labor constraints bind, the household responds to the sample plot shock by allocating additional labor to the sample plot, but they may withdraw that labor from either the most important plot or from leisure. Whether wealthier or larger households withdraw relatively more labor from the most important plot depends on the higher order derivatives of the utility and production functions; in general, these differential responses can not be signed.⁴⁰ Additionally, one key difference from the insurance case and input case is that household size and wealth no longer enter the model symmetrically. In one sense, household size and wealth instead enter the model as opposing forces: wealthier households allocate less labor to their plots, as they value leisure relatively more than consumption, while larger

³⁸If *all* households are input constrained, then the effect of the sample plot shock on input allocations on the most important plot depends on characteristics of the production function. Note that in this case, larger households will still exhibit a response in the same direction as wealthier households as both effects enter only through the wealth channel.

³⁹See proof in Appendix D.

⁴⁰Of course, the potential for ambiguous responses is heightened further if other forms of labor constraints, for example on hiring labor, are also considered.

households allocate more labor to their plots.

We focus on one particular case that builds on this intuition, presented in Figure 6. When on farm labor supply exhibits sufficient curvature, then changes in responsiveness to the sample plot shock of allocations on the most important plot are dominated by changes in the elasticity of on farm labor supply; suppose this to be the case, and further suppose that the elasticity of on farm labor supply is decreasing in the shadow wage. As we can think of household size as shifting out on farm labor supply (by increasing \bar{L}), and wealth as shifting in on farm labor supply (by increasing the marginal utility of leisure relative to the marginal utility of consumption), then larger households are located on a more elastic portion of their on farm labor supply schedule, while wealthier households are located on a less elastic portion of their on farm labor supply schedule.⁴¹ As a result, larger households will be less responsive to the sample plot shock, as they will primarily draw labor on the sample plot from leisure, while wealthier households will be more responsive to the sample plot shock, as they will primarily draw labor on the sample plot from the most important plot.

These predictions of the model, summarized in Table 7, generate a test that allows us to reject the absence of labor constraints. In particular, note that while insurance constraints or input constraints can rationalize the allocations of wealthier households to their most important plot as less responsive to the sample plot shock, only the presence of labor constraints can rationalize them as *more* responsive to the sample plot shock. Additionally, note that the model would struggle to rationalize larger households as more responsive to the sample plot shock, although it is possible to do so in the presence of labor constraints. In sum, we would interpret observing larger households as (weakly) less responsive and richer households as less responsive to the sample plot shock as most consistent with the presence of either input or insurance constraints, observing larger households as less responsive and richer households as more responsive as evidence for the presence of labor constraints, and

⁴¹This relationship between household size, wealth, and on farm labor supply elasticity has been posited as far back as Lewis (1954), and is discussed in depth in Sen (1966).

observing larger households as more responsive as inconsistent with our model.

5 Separation failures and adoption of irrigation

5.1 Empirical strategy

Our first specification to test for separation failures mirrors Equation (1), which we use to estimate the impacts of irrigation. We still make use of the discontinuity across the command area boundary, but outcomes are now on the household’s most important plot (plot 2) instead of the sample plot (plot 1).

$$y_{2ist} = \beta_0 + \beta_1 CA_{1is} + \alpha_{st} + \epsilon_{2ist} \quad (7)$$

We report β_1 , the effect of the sample plot shock on outcomes on the most important plot. In other specifications, we also consider heterogeneity with respect to the location of the most important plot, and include $CA_{1is} * CA_{2is}$ to test for this. In these specifications, we also report this difference in differences coefficient. For both this coefficient and β_1 , in line with the model predictions in Table 7, we interpret negative coefficients on labor, inputs, irrigation use, and horticulture, as evidence of separation failures.

As in Section 3, we include specifications with progressively more controls. Specifically, we also estimate

$$y_{2ist} = \beta_0 + \beta_1 CA_{1is} + \beta_2 Dist_{1is} + \beta_3 CA_{1is} * Dist_{1is} + \beta_4 CA_{2is} + \gamma_1 X_{1is} + \gamma_2 X_{2is} + \alpha_{st} + \epsilon_{2ist} \quad (8)$$

$$y_{2ist} - \bar{y}_{2ist} = \beta_1 (CA_{1is} - \overline{CA}_{1is}) + (V_{1is} - \overline{V}_{1is})' \gamma_1 + (V_{2is} - \overline{V}_{2is})' \gamma_2 + (\epsilon_{2ist} - \bar{\epsilon}_{2ist}) \quad (9)$$

Equation 8 includes controls CA_{2is} , an indicator for whether the most important plot is in the command area, and X_{1is} and X_{2is} , the log area of the sample plot and the most important

plot, respectively. Equation 9 uses spatial fixed effects, as described in Section 3.1.⁴²

Our benchmark specification to test for which constraints drive the separation failures is similar, but also includes the interaction of households characteristics with the sample plot shock. For parsimony, we only present the specification of this interaction for a specification similar to Equation 8; all tables present results with interactions included in Equation 7 and Equation 9 similarly.

$$y_{2ist} = \beta_0 + \beta_1 CA_{1is} + W_i' \beta_2 + CA_{1is} * W_i' \beta_3 + \beta_4 Dist_{1is} + \beta_5 CA_{1is} * Dist_{1is} \\ + \beta_6 CA_{2is} + X'_{1is} \gamma_1 + X'_{2is} \gamma_2 + \alpha_{st} + \epsilon_{2ist} \quad (10)$$

where W_i is a vector of household characteristics, which includes household size and an asset index in our primary specifications. We focus on β_3 : the heterogeneity, with respect to household characteristics, of the impacts of the sample plot shock on outcomes on the most important plot. The signs on β_3 give our main test of which market failures cause separation failures; Table 7 presents which signs map to which market failures.

5.1.1 Balance

We now use specifications (7), (8), and (9) to examine whether the most important plots in our sample are comparable for households whose sample plot is just inside or just outside the command area. As in Section 3.1.1, for each of these specifications, we show balance both with key controls omitted (Columns 3, 5, and 6), and our preferred specifications which we use in our analysis with key controls included (Columns 4, 7, and 8). Balance tests for most important plots are reported in Table 8. First, note that specifications that do not restrict to the discontinuity sample perform particularly poorly here. Most notably, most important plots are more likely to be located in the command area when sample plots are

⁴²Note that all differencing in this specification is done using the location of sample plots; in other words, most important plots whose associated sample plots are near each other are compared, as opposed to most important plots which are near each other.

also located in the command area, as households' plots tend to be located near each other. In contrast, our preferred specifications (Columns 4, 7, and 8, Table 8) which restrict to the discontinuity sample correct for this imbalance. Otherwise, we have a p-value of less than 0.1 for one variable in Column 4 (an indicator for owning the plot); for all three specifications, the omnibus test fails to reject the null of balance.

As an additional check, in Appendix F, we estimate for 2014 Dry specifications (7), (8), and (9), and specifications with heterogeneity following Equation (10). As the command area, as of the baseline, had not yet caused a large increase in demand for labor or inputs, or caused large increases in agricultural production, we would not anticipate any effects on MIPs. In line with this prediction, we fail to find any consistent significant effects on MIPs, either in our main specifications or for heterogeneity.

5.2 Results

5.2.1 A test for separation failures

We now present results on separation failures, demonstrating that the sample plot shock causes farmers to substitute away from their most important plot.⁴³ First, we present graphical evidence of this substitution in Figure 7. As in earlier figures, distance of the sample plot to the canal in meters is represented on the x-axis, with a positive sign indicating that the plot is on the command area side of the boundary. However, we now plot outcomes on both the sample plot and the most important plot. In this figure, substitution will manifest as decreases in input and labor use on the most important plot when the sample plot is in the command area, while input and labor use increase on the sample plot. Second, we present regression evidence in Tables 9, 10, and 11. In the discussion below, we focus on results from the tables, but we note that these results are consistent with visual intuition

⁴³We 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. These additional results are available upon request.

from Figure 7.

First, consistent with the presence of separation failures, in Columns 3 through 5 we find households substitute labor and inputs away from their most important plot. Households decrease allocations of household labor (11 - 33 person-days/ha) and inputs (2,100 - 6,700 RwF/ha) on their most important 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 (1.9 - 4.4pp) and cultivation of horticulture (1.6 - 3.8pp), while increasing cultivation of bananas (6.5 - 9.2pp).⁴⁴

Next, we expect the results above to be driven primarily by most important 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, in Columns 6 through 8, we find our results on irrigation, horticulture, and inputs are all driven by plots located in the command area. When the most important plot is located in the command area, the 16 - 20pp increase in irrigation use on sample plots in the command area coincides with a 8 - 10pp decrease in irrigation use on the most important 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.

As discussed in Section 3, the direct effects of the command area appear driven by enabling the transition to dry season horticultural cultivation and substitution away from lower value banana cultivation. However, the model in Section 4 is agnostic about whether decreases in labor and input allocations on the most important plot are driven by extensive margin responses (i.e., decreases in horticulture) or intensive margin responses (i.e., decreases in labor and input allocations conditional on crop choice). To test this, in Tables 12 and 13,

⁴⁴While these results are not consistently statistically significant, the specifications used lose power by including most important plots outside the command area, which are almost never irrigated and have small allocations of labor and inputs during the dry season. As discussed in the next paragraph, specifications which include the interaction of the sample plot command area indicator with a most important plot command area indicator are more precise for irrigation use, horticulture cultivation, and labor and input use.

we present results of the sample plot shock on labor and input use on sample plots and most important plots, controlling for cultivation and crop choice.⁴⁵ Table 12 confirms that the effects we document in Section 3 are driven by the shift to dry season horticulture, as effects on sample plots all but disappear controlling for crop choice. However, Table 13 suggests that much of the effect of the sample plot shock on labor and input use on most important plots is driven by intensive margin responses, as coefficients on household labor and inputs fall by only 18% - 36%. Combined with our results on irrigation use and horticulture, this suggests that households respond to the sample plot shock on both the intensive and extensive margins on their most important plot.

These results on separation failures imply the existence of a within household negative spillover, as they show that having one additional plot in the command area causes a household to substitute away from their other plots, reducing their use of irrigation, labor, and inputs on those 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. We would be particularly concerned about the bias generated by these reallocations if inputs were being shifted out of production on non-irrigated plots: in that case, our estimated impacts of access to irrigation would include reduced farming intensity on non-irrigated plots. However, we don't observe large reallocations of inputs away from non-CA plots: when the MIP is out of the command area, the estimated effect of an additional command area plot on inputs applied to the MIP is usually about 1/4 the magnitude of the effect on inputs in a command area MIP⁴⁶ and generally not statistically different from zero. We therefore conclude that the dominant within-household spillover is a reduced intensity of cultivation on irrigated plots,

⁴⁵As crop fixed effects are a "bad control" (Angrist & Pischke, 2008), which introduces selection bias, we interpret these results as suggestive. However, we anticipate that selection conditional on crop choice should bias us towards finding no intensive margin effect on most important plots, as the particularly constrained households switching out of horticulture in response to the sample plot shock are likely to be the households who used less labor and inputs.

⁴⁶The coefficient on sample plot in CA is usually about 1/3 the magnitude of the coefficient on sample plot in CA * MIP in CA in tables 9, 10, and 11 and MIPs in the CA experience the sum of both effects

suggesting that any bias in our ToT estimates above may be small and render those estimates conservative⁴⁷.

5.2.2 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 two or more plots in the command area only had one plot 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 most important plot, which we can interpret as the effect of a household’s second plot (the sample plot) being moved to the command area on adoption of irrigation on its first plot in the command area (the most important plot).

Specifically, we calculate

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

First, $(\beta_1 + \beta_{3,CA})$ is the total effect of the sample plot shock on adoption of irrigation on most important plots in the command area. Second, in the denominator, we count the total number of command area plots among households’ sample plots and most important plots.⁴⁸ Third, in the numerator, we apply the estimated substitution caused by the sample plot shock to both the sample plot and the most important plot, as households are also substituting away from their sample plot when the most important plot is in the command area.

We find adoption of irrigation would be 21% - 24% higher under this counterfactual. 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 by

⁴⁷This view is further supported by table 12, which indicates that conditional on crop choice sample plots in the CA are cultivated with very similar intensity to sample plots outside of the CA

⁴⁸We implicitly ignores households’ other plots; we do so because our research design has little to say about the impacts of additional command area plots, or on households’ behavior on these plots, so we interpret this exercise as estimating a lower bound on the impact of reallocation on adoption of irrigation.

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.

5.2.3 Separating constraints

We now provide evidence on the source of the separation failure by estimating heterogeneous impacts, with respect to household size and wealth, of the sample plot shock on outcomes on the most important plot. Recall that for this analysis, the key predictions of the model were 1) if only insurance or input constraints bind, wealthier households and larger households should be less responsive, and 2) if only labor constraints bind, differential responsiveness of wealthier and larger households is ambiguous, but under reasonable assumptions wealthier households should be more responsive 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 Tables 14 and 15. 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 most important plot by 50% - 94% for irrigation use, 73% - 102% for horticulture, 63% - 75% for household labor, and 20% - 21% for inputs, with all but the input coefficient statistically significant and robust across specifications.⁴⁹ In contrast,

⁴⁹These percentages, and the remainder of percentages in this paragraph, are expressed relative to the estimated impact of the sample plot shock on the most important plot, and are calculated only for Specifications

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 most important plot by 41% - 97% for irrigation use, 39% - 81% for horticulture, 39% - 72% for household labor, and 42% - 58% 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 most important plot in response to the sample plot shock. As discussed in Section 4.3, 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.

6 Experimental evidence

Our results leveraging the discontinuity suggest that land and labor market frictions combine to constrain the adoption of hillside irrigation in Rwanda. We design and run three experiments to test for the presence of other constraints to adoption of irrigation – specifically, we focus on operations and maintenance of irrigation schemes, and financial and informational constraints. These experimental results corroborate that labor market failures are a primary constraint to adoption of irrigation in this context. Additional details on the motivation, treatment assignment protocols, and logistics of implementation of each of these experiments are presented in Appendix E.

First, we test whether failures of operations and maintenance impose a constraint that limits farmers’ adoption of irrigation. The government implementing agency designed a centralized O&M system to establish and enforce water usage schedules to ensure farmers’ access to water. If farmers faced limited access to water due to problems in the operations and maintenance system, this could constrain adoption of irrigation. We sought to alleviate

(8) and (9).

this potential constraint by randomizing empowerment of local monitors to assist system operators and report maintenance needs. We find no evidence this experiment changed cultivation practices. This result is likely because very few farmers report any challenges related to operations and maintenance over the four years of survey data collection. Second, the government planned to charge farmers in the command area land taxes, which were unconditional on cultivation decisions, to fund operations and maintenance in the schemes. To test whether these fees would limit farmers adoption of irrigation, we randomized subsidies of farmers' fees. We find no evidence this experiment changed cultivation practices. This result is likely because compliance with the fees was extremely low (4%), so collected fees were too low to plausibly constrain farmers. We discuss these experiments further in Appendix E.2, and conclude here that these issues were not relevant in this context.

Third, we test whether financial and informational constraints limit adoption of irrigation. To do so, we assigned horticultural minikits to randomly selected farmers 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. Lastly, they should facilitate experimentation and increase adoption 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., 2018). To test for spillovers, water user groups were randomly assigned to 20%, 60%, or 100% minikit saturation, with rerandomization for balance on Zone and O&M treatment status. Minikits were offered to assigned individuals prior to 2017 Rainy 1 and 2017 Dry.⁵⁰

⁵⁰Each of these three interventions exist only in the command area. As such, the effects of irrigation estimated throughout this paper are averages across the experimental treatments. Overall, this concern is mitigated by the fact that all three experimental treatments had very limited impacts on cultivation practices. In addition, the first two of these treatments (fee subsidies and monitoring systems) vary characteristics which would be heterogeneous across different irrigation systems; we are therefore comfortable with the interpretation that estimates above exist for the average of these treatments. Readers may be most concerned about interpretations of treatment effects in the presence of the minikit treatment; in addition to the modest effects on cultivation described below, we have also conducted analysis excluding minikit winners

6.1 Empirical strategy and results

We estimate the impact of minikits using the specification

$$y_{1ist} = \beta_0 + \beta_1 \text{Assigned minikit}_i + \beta_2 \text{Minikit saturation}_i + X'_{1is} \gamma + \epsilon_{1ist} \quad (11)$$

Assigned minikit_{*i*} is a dummy for whether household *i* was randomly assigned to receive a minikit, Minikit saturation_{*i*} is the probability of receiving a minikit for households in the water user group of household *i*'s sample plot, and X_{1is} includes the stratification variables (Zone fixed effects and O&M treatment status), as well as indicator variables reflecting the probability that a household would receive a minikit⁵¹ and in some specifications 2016 Dry horticulture adoption. As minikit saturation is assigned at the water user groups level, robust standard errors are clustered at the water user group level.

For our primary outcomes *y*, we focus on whether households used a minikit (in 2017 Rainy 1 or in 2017 Dry) and adoption of horticulture. Impacts on minikit use are our first stage and impacts on adoption of horticulture are our measure of learning from the minikits. For precision, we restrict to command area plots, and for plot level outcomes we focus on 2017 Dry and 2018 Dry; these are the plots and seasons in which we expect households to adopt horticulture in response to being assigned a minikit.

We present the results of this analysis in Table 16. First, we find a strong first stage; households assigned to receive a minikit are 40pp more likely to use a minikit than households not assigned to receive a minikit. Almost all non-compliance is driven by households who were assigned to receive a minikit but did not pick it up – 4.8% of households not assigned to receive a minikit used one, while 43.8% of households assigned to receive a minikit used one. Second, we find no effects of minikits on horticulture use, and we have sufficient precision to

and conclusions are qualitatively unaffected.

⁵¹After matching names from the lists of water user group members to our baseline survey, we found that 32% of households either had multiple household members on the lists of water user group members or had a single household member listed multiple times; these households are more likely to be assigned to receive a minikit and may differ from other households

reject estimates from other contexts of the effect of minikits on technology adoption (Emerick et al., 2016; Jones et al., 2018). Third, consistent with this null effect on horticulture use, we find no effects of minikit saturation, although these estimates are less precise than those of the impacts of assignment to receive a minikit; we note that we also fail to reject that the sum of the coefficients on assigned minikit and minikit saturation (the effect on adoption in a fully treated compared to an untreated waater user group) is zero. Fourth, we find strong positive selection into using a minikit: farmers who grew horticulture in 2016 Dry, who are 30.6pp more likely to grow horticulture in 2017 and 2018 Dry, are 13.1pp more likely to use a minikit in response to assignment to receive a minikit receipt.

We interpret these results as corroborating evidence that information and financial constraints are not dominant constraints to adoption of irrigation. Most farmers assigned to receive a minikit do not pick it up and use it, and the farmers who do pick it up typically would have grown horticulture even if not assigned to receive a minikit. We similarly find no evidence that saturation of minikits lead to increased adoption, as we might expect if learning was important.⁵² Our experimental evidence therefore supports the conclusion that, in this context, financial and informational frictions are not the primary explanations for the low and inefficient irrigation use we observe.

7 Conclusions

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 document that the construction of an irrigation system leads to a 44% - 71% increase in cash profits. These profits are generated by a switch in cropping patterns from perennial bananas towards a rotation of dry-season horticulture and rainy-season staples, which itself is associated with an increase in input intensity. In our context,

⁵²That information is not a binding constraint is also consistent with the stability in levels of irrigation adoption that we observe over time, in contrast to an S-curve of adoption which would be consistent with learning.

the primary increase in input demands is for household labor, which is used intensively on horticulture and minimally on banana cultivation.

These results suggest that irrigation may have similar potential in Africa to the transformative role it played in South Asia, where other studies have documented similar impacts of irrigation on farmer revenues and yields. In some ways, this is surprising: other evidence on the use of inputs in Africa and the returns to those inputs often finds lower usage and technological returns in the African context. These two facts together suggest that expanding irrigation access in Africa may be a necessary contributor to shrinking the yield gap.

At the same time, even with access to a new, highly productive technology provided freely by the government we observe a minority of farmers adopting this technology four years after introduction. Given the returns identified above, we take this as evidence that the existence of a productive technology is not itself sufficient to generate majority adoption in all agricultural contexts. We further document that frictions in land and labor markets contribute to low utilization of irrigation systems by examining farmers' input utilization on other plots in response to irrigation investments. 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.

These results highlight 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 generate limits to the effectiveness of financial and information interventions in improving agricultural productivity.

References

- Adamopoulos, T., Brandt, L., Leight, J., & Restuccia, D. (2017). *Misallocation, Selection and Productivity: A Quantitative Analysis with Panel Data from China*. Working Paper 23039, National Bureau of Economic Research.
- Adamopoulos, T. & Restuccia, D. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, 104(6), 1667–97.
- Adamopoulos, T. & Restuccia, D. (2018). *Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data*. Working Paper 24532, National Bureau of Economic Research.
- Ali, D. A., Deininger, K., & Goldstein, M. (2014). Environmental and gender impacts of land tenure regularization in africa: Pilot evidence from rwanda. *Journal of Development Economics*, 110, 262–275.
- Anderson, M. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American Statistical Association*, 103(484), 1481–1495.
- Angrist, J. D. & Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.
- Attwood, D. W. (2005). Big is ugly? how large-scale institutions prevent famines in western india. *World Development*, 33(12), 2067–2083.
- Barrett, C. B. (1996). On price risk and the inverse farm size-productivity relationship. *Journal of Development Economics*, 51(2), 193–215.
- Benjamin, D. (1992). Household composition, labor markets, and labor demand: testing for separation in agricultural household models. *Econometrica: Journal of the Econometric Society*, (pp. 287–322).

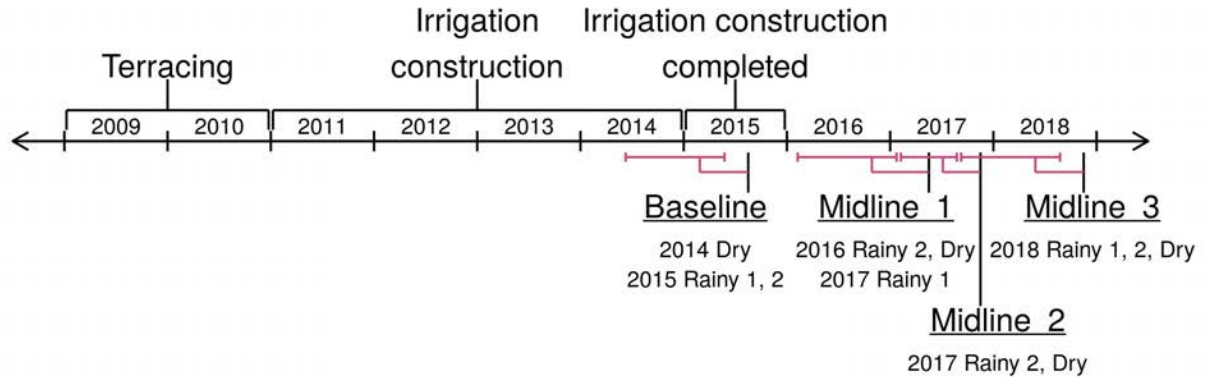
- Besley, T. (1995). Property rights and investment incentives: Theory and evidence from ghana. *Journal of Political Economy*, 103(5), 903–937.
- Besley, T. & Ghatak, M. (2010). Property rights and economic development. In *Handbook of development economics*, volume 5 (pp. 4525–4595). Elsevier.
- Bleakley, H. & Ferrie, J. (2014). Land openings on the georgia frontier and the coase theorem in the short-and long-run.
- Breza, E., Krishnaswamy, N., & Kaur, S. (2018). Scabs: The social suppression of labor supply.
- Conley, T. G. (1999). Gmm estimation with cross-sectional dependence. *Journal of Econometrics*, 92(1), 1–45.
- Conley, T. G. & Udry, C. R. (2010). Learning about a new technology: Pineapple in ghana. *The American Economic Review*, (pp. 35–69).
- De Janvry, A., Sadoulet, E., & Suri, T. (2017). Field experiments in developing country agriculture. In *Handbook of Economic Field Experiments*, volume 2 (pp. 427–466). Elsevier.
- Deininger, K. & Feder, G. (2009). Land registration, governance, and development: Evidence and implications for policy. *The World Bank Research Observer*, 24(2), 233–266.
- Dillon, A. (2011). The effect of irrigation on poverty reduction, asset accumulation, and informal insurance: Evidence from northern mali. *World Development*, 39(12), 2165–2175.
- Dillon, A. & Fishman, R. (2019). Dams: Effects of hydrological infrastructure on development. *Annual Review of Resource Economics*, 11.
- Dillon, B. & Barrett, C. B. (2017). Agricultural factor markets in sub-saharan africa: An updated view with formal tests for market failure. *Food policy*, 67, 64–77.

- Dillon, B., Brummund, P., & Mwabu, G. (2019). Asymmetric non-separation and rural labor markets. *Journal of Development Economics*, 139, 78–96.
- Duflo, E. & Pande, R. (2007). Dams. *The Quarterly Journal of Economics*, 122(2), 601–646.
- Emerick, K., de Janvry, A., Sadoulet, E., & Dar, M. H. (2016). Technological innovations, downside risk, and the modernization of agriculture. *American Economic Review*, 106(6), 1537–61.
- Fafchamps, M. (1993). Sequential labor decisions under uncertainty: An estimable household model of west-african farmers. *Econometrica*, 61, 1173–1197.
- Fei, J. C. & Ranis, G. (1961). A theory of economic development. *The American Economic Review*, 51(4), 533–565.
- Foster, A. & Rosenzweig, M. (2017). Are there too many farms in the world? labor-market transaction costs, machine capacities, and optimal farm size. *NBER Working Paper*, (No. 23909).
- Goldstein, M., Hounghbedji, K., Kondylis, F., O’Sullivan, M., & Selod, H. (2018). Formalization without certification? experimental evidence on property rights and investment. *Journal of Development Economics*, 132, 57–74.
- Goldstein, M. & Udry, C. (2008). The profits of power: Land rights and agricultural investment in ghana. *Journal of political Economy*, 116(6), 981–1022.
- Gollin, D. & Udry, C. (2019). Heterogeneity, measurement error and misallocation: Evidence from african agriculture. *NBER Working Paper 25440*.
- Heisey, P. W. & Norton, G. W. (2007). Fertilizers and other farm chemicals. *Handbook of agricultural economics*, 3, 2741–2777.
- Jacoby, H. G. (1993). Shadow wages and peasant family labour supply: an econometric application to the peruvian sierra. *The Review of Economic Studies*, 60(4), 903–921.

- Jacoby, H. G. (2017). “well-fare” economics of groundwater in south asia. *The World Bank Research Observer*, 32(1), 1–20.
- Jones, M., Kondylis, F., Mobarak, A. M., & Stein, D. (2018). Evaluating the integrated agriculture productivity project in bangladesh.
- Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics*, 129(2), 597–652.
- Kaur, S. (2014). Nominal wage rigidity in village labor markets. *NBER Working Paper 20770*.
- Key, N., Sadoulet, E., & Janvry, A. D. (2000). Transactions costs and agricultural household supply response. *American journal of agricultural economics*, 82(2), 245–259.
- LaFave, D. & Thomas, D. (2016). Farms, families, and markets: New evidence on completeness of markets in agricultural settings. *Econometrica*, 84(5), 1917–1960.
- Lewis, W. A. (1954). Economic development with unlimited supplies of labour. *The manchester school*, 22(2), 139–191.
- Loeser, J. (2020). India’s groundwater revolution and agricultural productivity.
- Magruder, J. (2012). High unemployment yet few small firms: The role of centralized bargaining in south africa. *American Economic Journal: Applied Economics*, 4(3), 138–66.
- Magruder, J. (2013). Can minimum wages cause a big push? evidence from indonesia. *Journal of Development Economics*, 100(1), 48–62.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.

- Restuccia, D. & Santaaulalia-Llopis, R. (2017). Land misallocation and productivity. *NBER Working Paper 23128*.
- Sekhri, S. (2014). Wells, water, and welfare: the impact of access to groundwater on rural poverty and conflict. *American Economic Journal: Applied Economics*, 6(3), 76–102.
- Sen, A. K. (1966). Peasants and dualism with or without surplus labor. *Journal of political Economy*, 74(5), 425–450.
- Singh, I., Squire, L., & Strauss, J. (1986). *Agricultural household models: Extensions, applications, and policy*. The Johns Hopkins University Press.
- Skoufias, E. (1994). Using shadow wages to estimate labor supply of agricultural households. *American journal of agricultural economics*, 76(2), 215–227.
- Smith, C. (2019). Land concentration and long-run development: Evidence from the frontier united states.
- Suri, T. (2011). Selection and comparative advantage in technology adoption. *Econometrica*, 79(1), 159–209.
- Udry, C. (1996). Gender, agricultural production, and the theory of the household. *Journal of political Economy*, 104(5), 1010–1046.
- Udry, C. (1997). Efficiency and market structure: testing for profit maximization in african agriculture.
- World Bank (2007). World development report 2008: Agriculture for development.

Figure 1: Timeline



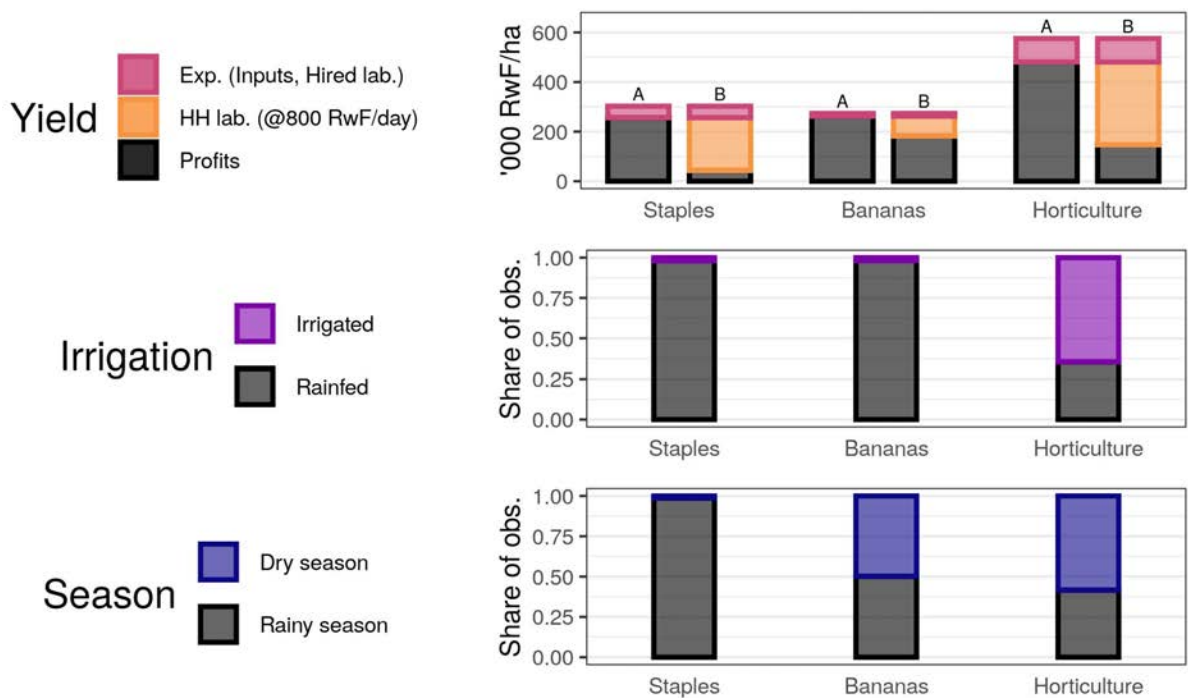
Notes: A timeline of events on the 3 hillside irrigation schemes we study is presented in this figure. Black lines are used to indicate when (or the period during which) events took place, while pink lines are used to indicate survey recall periods.

Figure 2: Hillside irrigation scheme




Notes: A photograph of Karongi 12, one of the hillside irrigation schemes in this study, is presented in this figure.

Figure 3: Irrigation used for dry season labor intensive horticulture, profitability depends on household's shadow wage



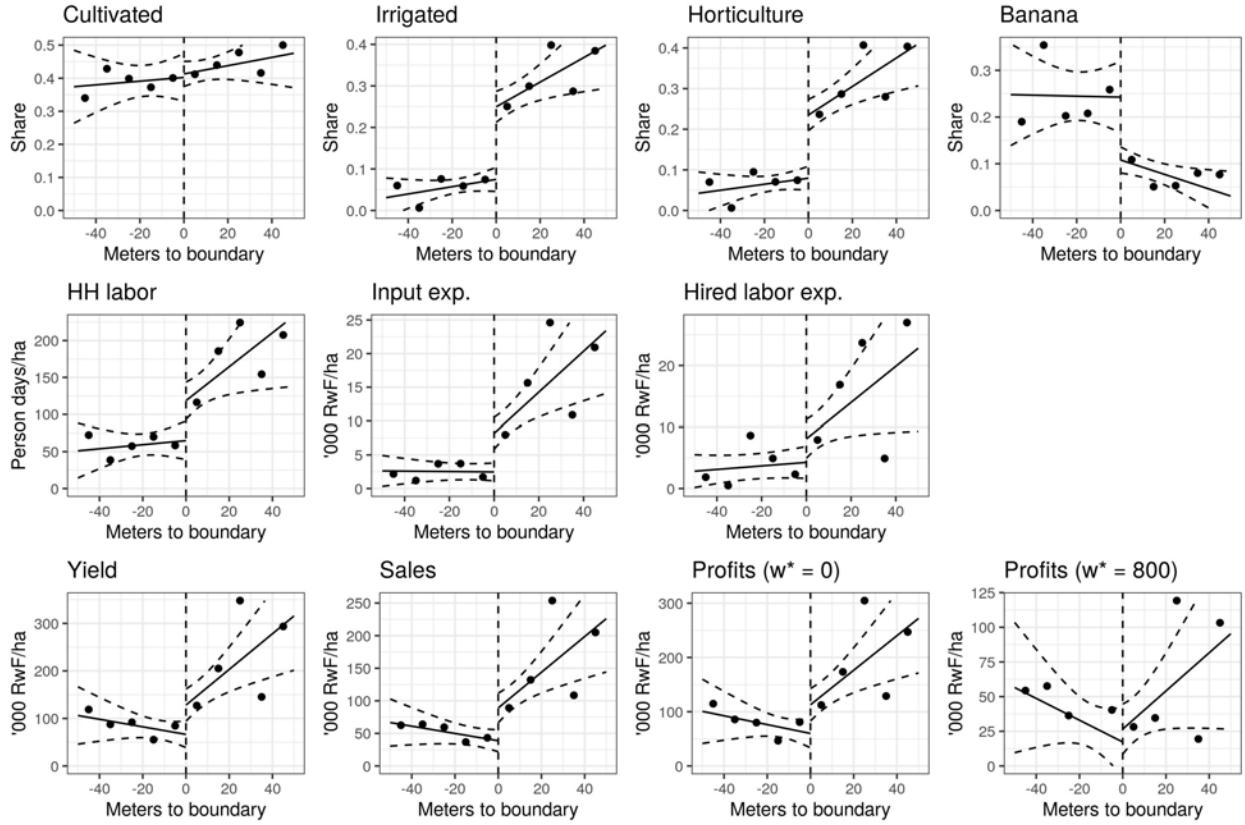
Notes: Sample averages of outcomes by crop per agricultural season are presented in this figure. In the top panel, the height of bars is yield. Columns A show profits calculated valuing household labor at 0 RwF/person-day, while Columns B show profits calculated valuing household labor at 800 RwF/person-day (the median wage in our data). In the middle and bottom panel, bars represent, in our data, the share of observations of each set of crops that are irrigated compared to rainfed and are during the dry season compared to the rainy season, respectively.

Figure 4: Adoption dynamics




Notes: Average adoption of irrigation by season on sample plots in the main discontinuity sample, inside and outside the command area, is presented in this figure. Averages outside the command area are in black, while averages inside the command area and 95% confidence intervals for the difference are in pink. Robust standard errors are clustered at the nearest water user group level.

Figure 5: Regression discontinuity estimates of impacts of irrigation



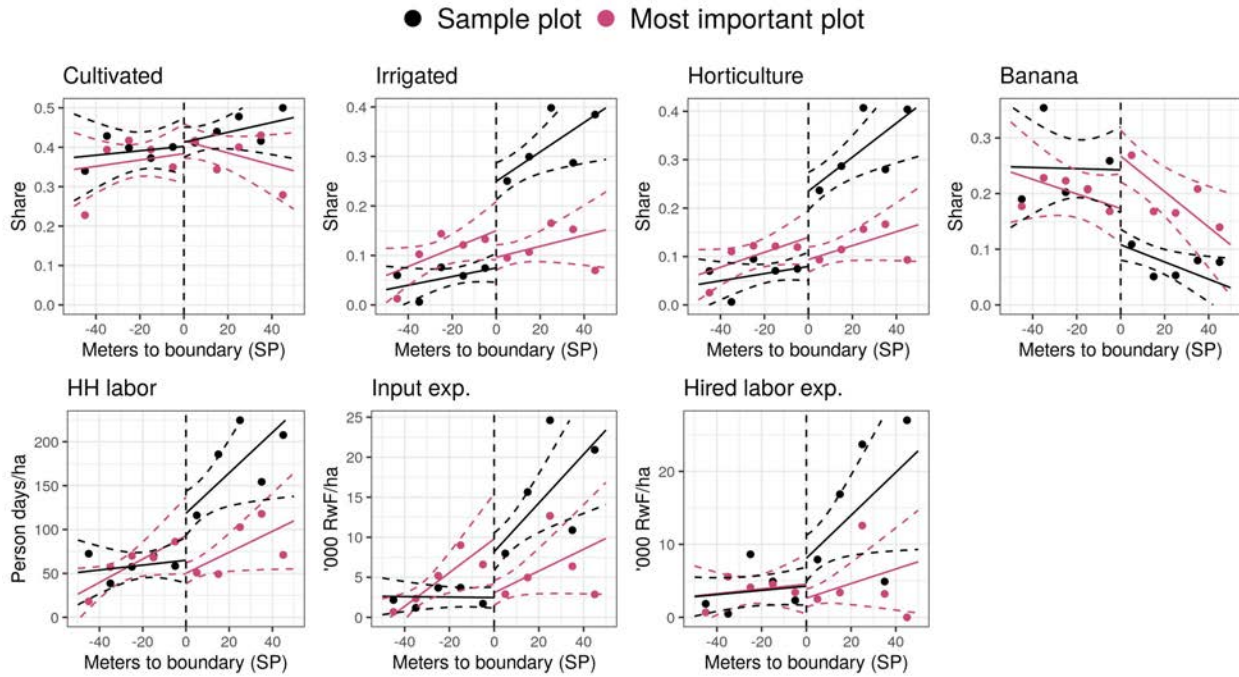
Notes: Visual regression discontinuity analysis on sample plots in the main discontinuity sample during the dry season is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes. Predicted outcomes from regressions of outcomes on distance to the command area boundary, a command area dummy, and their interaction are presented with 95% confidence intervals on the prediction. Robust standard errors are clustered at the nearest water user group level.

Figure 6: Differential responses to sample plot shock under labor constraints



Notes: Households' labor allocations under a binding off farm labor constraint are presented in this figure. L_k and l are the household's labor allocation on plot k and choice of leisure, respectively, as a function of the shadow wage, with the argument suppressed. $L_1 + L_2$ is total household on farm labor demand; if the household's sample plot ($k = 1$) is in the command area ("sample plot shock"), on farm labor demand shifts out to $L'_1 + L_2$. $\bar{L}^{\text{SM}} - \bar{L}^{\text{O}} - l$ is household on farm labor supply; for large households, on farm labor supply is shifted out to $\bar{L}^{\text{BIG}} - \bar{L}^{\text{O}} - l$. The shadow wage is determined by the intersection of on farm labor demand and on farm labor supply, and labor allocations on the most important plot are L_2 evaluated at this shadow wage. In this figure, larger households are on a more elastic portion of their on farm labor supply schedule; as a result, the sample plot shock causes a smaller increase in the shadow wage, and in turn a smaller decrease in labor allocations on the most important plot (smaller in magnitude dL_2/dA_1).

Figure 7: Regression discontinuity estimates of most important plot responses to sample plot shock



Notes: Visual regression discontinuity analysis on sample plots and associated most important plots during the dry season, for sample plots in the main discontinuity sample, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes. Predicted outcomes from regressions of outcomes on distance to the command area boundary, a command area dummy, and their interaction are presented with 95% confidence intervals on the prediction. Robust standard errors are clustered at the nearest water user group level.

Table 1: Summary statistics on agricultural production

	Staples			Bananas	Horticulture		
	Staples (1)	Maize (2)	Beans (3)		All (5)	Rainy (6)	Dry (7)
Yield	302	318	285	273	575	588	566
Hired labor (days)	37	37	37	9	61	66	57
HH labor (days)	266	248	260	101	417	414	420
Inputs	19	35	16	3	50	50	50
Profits							
Shadow wage = 0 RwF/day	256	255	241	263	481	489	475
Shadow wage = 800 RwF/day	43	56	34	182	147	158	139
Sales share	0.19	0.30	0.14	0.46	0.62	0.60	0.63
Irrigated	0.02	0.02	0.02	0.02	0.65	0.25	0.93
Rainy	0.99	1.00	1.00	0.50	0.42	1.00	0.00
log area	-2.44	-2.26	-2.47	-2.10	-2.71	-2.83	-2.62
Share of obs.	0.65	0.13	0.42	0.19	0.12	0.05	0.07

Notes: Sample averages of outcomes by crop per agricultural season are presented in this table. Yield, inputs, 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.

Table 2: Balance: Sample plot characteristics

	Full sample		RD sample					
	Coef. (SE) [p]	Dep. var.	Coef. (SE) [p]					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log area	0.045 (0.077) [0.554]	-2.515 (1.179) 969	0.219 (0.087) [0.012]	0.285 (0.087) [0.001]	0.425 (0.121) [0.000]	0.200 (0.128) [0.118]		
Own plot	-0.012 (0.020) [0.535]	0.894 (0.309) 969	0.003 (0.023) [0.897]	-0.001 (0.024) [0.966]	0.004 (0.032) [0.907]	-0.004 (0.038) [0.921]	-0.001 (0.032) [0.972]	-0.006 (0.038) [0.877]
Owned plot >5 years	0.045 (0.019) [0.020]	0.880 (0.326) 686	0.070 (0.026) [0.006]	0.072 (0.025) [0.004]	0.019 (0.037) [0.613]	0.012 (0.035) [0.723]	0.007 (0.036) [0.834]	0.010 (0.034) [0.767]
Rented out, farmer	0.027 (0.012) [0.022]	0.032 (0.177) 969	0.018 (0.014) [0.197]	0.019 (0.014) [0.182]	-0.003 (0.023) [0.884]	0.009 (0.027) [0.726]	-0.009 (0.023) [0.699]	0.007 (0.027) [0.796]
Omnibus F-stat [p]	2.6 [0.038]		3.4 [0.010]	4.9 [0.001]	3.2 [0.013]	0.6 [0.639]	0.1 [0.979]	0.1 [0.984]
Site FE				X	X		X	
Distance to boundary					X	X	X	X
log area							X	X
Spatial FE						X		X

Notes: Balance for sample plot characteristics is presented in this table. Column 2 presents, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 1 and 3 through 8 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Controls are listed below. The final row of each column presents the Omnibus F-stat for the null of balance on all outcomes, with the p-value for the associated test in brackets. Column 1 uses the full sample, while Columns 2 through 8 use the discontinuity sample. Column 4 uses the specification in Equation (1), Column 7 uses the specification in Equation (2), and Column 8 uses the specification in Equation (3).

Table 3: Balance: Household characteristics

	Full sample		RD sample					
	Coef. (SE) [p]	Dep. var.	Coef. (SE) [p]					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
HHH female	0.041 (0.025) [0.094]	0.221 (0.416) 969	0.057 (0.029) [0.054]	0.055 (0.029) [0.063]	0.045 (0.046) [0.326]	0.044 (0.050) [0.378]	0.043 (0.046) [0.345]	0.041 (0.050) [0.412]
HHH age	0.5 (0.8) [0.497]	47.5 (14.5) 967	1.5 (0.9) [0.096]	1.5 (0.9) [0.087]	2.1 (1.4) [0.127]	0.7 (1.8) [0.694]	1.4 (1.4) [0.298]	0.3 (1.9) [0.863]
HHH completed primary	0.069 (0.025) [0.005]	0.287 (0.453) 966	0.044 (0.031) [0.159]	0.052 (0.032) [0.106]	0.128 (0.047) [0.006]	0.102 (0.062) [0.097]	0.119 (0.047) [0.012]	0.099 (0.062) [0.111]
HHH worked off farm	0.023 (0.027) [0.392]	0.410 (0.493) 969	-0.023 (0.035) [0.516]	-0.033 (0.035) [0.350]	-0.039 (0.051) [0.441]	-0.019 (0.064) [0.763]	-0.024 (0.050) [0.631]	-0.011 (0.064) [0.868]
# of plots	0.61 (0.18) [0.001]	5.19 (3.38) 969	0.37 (0.22) [0.099]	0.16 (0.21) [0.442]	0.20 (0.36) [0.582]	0.35 (0.46) [0.448]	0.36 (0.36) [0.319]	0.40 (0.46) [0.382]
# of HH members	0.17 (0.11) [0.104]	4.89 (2.16) 969	0.04 (0.15) [0.799]	0.02 (0.15) [0.916]	-0.00 (0.21) [0.985]	-0.03 (0.25) [0.917]	-0.01 (0.22) [0.971]	-0.03 (0.25) [0.908]
# who worked off farm	0.10 (0.05) [0.039]	0.77 (0.85) 969	0.04 (0.06) [0.523]	0.02 (0.06) [0.771]	0.01 (0.08) [0.909]	0.03 (0.10) [0.799]	0.01 (0.08) [0.906]	0.04 (0.10) [0.722]
Housing expenditures	-2.3 (6.9) [0.739]	49.2 (127.4) 962	3.5 (9.0) [0.700]	3.3 (9.0) [0.717]	-5.6 (14.9) [0.707]	-16.7 (19.0) [0.380]	-6.5 (14.7) [0.658]	-18.6 (19.1) [0.328]
Asset index	0.11 (0.05) [0.034]	-0.12 (0.99) 967	0.06 (0.07) [0.372]	0.07 (0.07) [0.303]	0.15 (0.12) [0.215]	0.06 (0.13) [0.647]	0.13 (0.12) [0.291]	0.04 (0.13) [0.738]
Omnibus F-stat [p]	3.6 [0.000]		1.6 [0.122]	1.6 [0.118]	1.8 [0.080]	0.8 [0.571]	1.5 [0.158]	0.9 [0.507]
Site FE				X	X		X	
Distance to boundary					X	X	X	X
log area							X	X
Spatial FE						X		X

Notes: Balance for household characteristics is presented in this table. Column 2 presents, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 1 and 3 through 8 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Controls are listed below. The final row of each column presents the Omnibus F-stat for the null of balance on all outcomes, with the p-value for the associated test in brackets. Column 1 uses the full sample, while Columns 2 through 8 use the discontinuity sample. Column 4 uses the specification in Equation (1), Column 7 uses the specification in Equation (2), and Column 8 uses the specification in Equation (3).

Table 4: Access to irrigation enables transition to dry season horticulture from perennial bananas

	Dry season				Rainy seasons			
	Dep. var.	Coef. (SE) [p]			Dep. var.	Coef. (SE) [p]		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cultivated	0.391 (0.488) 2,537	0.033 (0.031) [0.289]	0.005 (0.041) [0.909]	0.022 (0.044) [0.610]	0.838 (0.369) 4,236	-0.054 (0.020) [0.006]	-0.092 (0.025) [0.000]	-0.053 (0.027) [0.051]
Irrigated	0.058 (0.233) 2,537	0.202 (0.019) [0.000]	0.162 (0.024) [0.000]	0.171 (0.030) [0.000]	0.016 (0.127) 4,236	0.044 (0.007) [0.000]	0.035 (0.009) [0.000]	0.059 (0.012) [0.000]
Horticulture	0.065 (0.246) 2,536	0.180 (0.020) [0.000]	0.137 (0.024) [0.000]	0.156 (0.029) [0.000]	0.073 (0.260) 4,235	0.044 (0.011) [0.000]	0.016 (0.018) [0.371]	0.048 (0.025) [0.056]
Banana	0.245 (0.430) 2,536	-0.134 (0.024) [0.000]	-0.133 (0.037) [0.000]	-0.142 (0.035) [0.000]	0.274 (0.446) 4,235	-0.149 (0.024) [0.000]	-0.158 (0.038) [0.000]	-0.168 (0.034) [0.000]
Site-by-season FE		X	X			X	X	
Distance to boundary			X	X			X	X
log area			X	X			X	X
Spatial FE				X				X

Notes: Regression analysis is presented in this table. Columns 1 through 4 restrict to observations during the dry season, while columns 5 through 8 restrict to observations during the rainy season. Columns 1 and 5 present, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 2 through 4 and 6 through 8 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Columns 2 and 6 use the specification in Equation (1). Columns 3 and 7 use the regression discontinuity specification in Equation (2). Columns 4 and 8 use the spatial fixed effects specification in Equation (3).

Table 5: Access to irrigation causes large increases in dry season labor and input use

	Dry season				Rainy seasons			
	Dep. var.	Coef. (SE) [p]			Dep. var.	Coef. (SE) [p]		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
HH labor/ha	59.5 (201.4) 2,523	69.6 (14.7) [0.000]	70.8 (17.5) [0.000]	76.9 (20.7) [0.000]	226.7 (316.7) 4,215	-7.7 (18.3) [0.671]	8.5 (23.1) [0.714]	9.9 (24.7) [0.689]
Input exp./ha	2.5 (17.4) 2,527	7.4 (1.3) [0.000]	6.3 (1.5) [0.000]	4.3 (1.8) [0.019]	16.1 (40.9) 4,223	2.5 (2.0) [0.205]	1.1 (2.9) [0.710]	2.1 (3.1) [0.511]
Hired labor exp./ha	3.7 (25.6) 2,527	5.6 (1.9) [0.003]	3.7 (2.1) [0.082]	3.2 (2.6) [0.221]	15.9 (47.1) 4,223	7.1 (2.4) [0.003]	3.7 (3.4) [0.276]	3.1 (4.5) [0.490]
Site-by-season FE		X	X			X	X	
Distance to boundary			X	X			X	X
log area			X	X			X	X
Spatial FE				X				X

Notes: Regression analysis is presented in this table. Columns 1 through 4 restrict to observations during the dry season, while columns 5 through 8 restrict to observations during the rainy season. Columns 1 and 5 present, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 2 through 4 and 6 through 8 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Columns 2 and 6 use the specification in Equation (1). Columns 3 and 7 use the regression discontinuity specification in Equation (2). Columns 4 and 8 use the spatial fixed effects specification in Equation (3).

Table 6: Access to irrigation causes large increases in dry season yields and sales, profitability depends on household's shadow wage

	Dry season				Rainy seasons			
	Dep. var.	Coef. (SE) [p]			Dep. var.	Coef. (SE) [p]		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Yield	82.3 (278.7) 2,402	61.2 (20.7) [0.003]	73.1 (23.2) [0.002]	55.0 (28.5) [0.054]	271.5 (457.0) 4,085	-45.1 (22.0) [0.041]	-22.6 (30.8) [0.462]	-15.4 (30.8) [0.617]
Sales/ha	49.7 (180.8) 2,527	52.3 (13.3) [0.000]	55.5 (14.5) [0.000]	49.3 (19.2) [0.010]	85.1 (229.1) 4,223	-4.8 (10.8) [0.660]	-13.3 (18.5) [0.472]	5.6 (21.6) [0.793]
Profits/ha								
Shadow wage = 0	76.1 (265.9) 2,402	49.6 (18.6) [0.008]	63.9 (21.0) [0.002]	49.1 (25.7) [0.057]	239.8 (432.5) 4,085	-53.4 (20.5) [0.009]	-26.4 (28.5) [0.354]	-19.4 (27.5) [0.480]
Shadow wage = 800	32.8 (224.1) 2,400	-0.3 (12.0) [0.978]	9.3 (16.5) [0.573]	-3.0 (20.8) [0.886]	59.5 (364.5) 4,078	-47.2 (16.5) [0.004]	-31.8 (26.4) [0.228]	-27.3 (31.9) [0.393]
Site-by-season FE		X	X			X	X	
Distance to boundary			X	X			X	X
log area			X	X			X	X
Spatial FE				X				X

Notes: Regression analysis is presented in this table. Columns 1 through 4 restrict to observations during the dry season, while columns 5 through 8 restrict to observations during the rainy season. Columns 1 and 5 present, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 2 through 4 and 6 through 8 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and Conley (1999) standard errors are used in specifications with Spatial FE. Columns 2 and 6 use the specification in Equation (1). Columns 3 and 7 use the regression discontinuity specification in Equation (2). Columns 4 and 8 use the spatial fixed effects specification in Equation (3).

Table 7: Model predictions

	$\frac{dL_2}{dA_1}$	$\frac{d}{d\bar{L}} \frac{dL_2}{dA_1}$	$\frac{d}{d\bar{M}} \frac{dL_2}{dA_1}$
No constraints	0	0	0
<hr/>			
Constraints			
Insurance	-	+	+
Inputs	-	0/+	+
Labor	-	+*	-*

Notes: Predicted signs from the model for key comparative statics of interest are presented in this table. Predictions in the no constraints case correspond to Proposition 1. Predictions on $\frac{dL_2}{dA_1}$ correspond to Proposition 2. Predictions on $\frac{d}{d\bar{L}} \frac{dL_2}{dA_1}$ and $\frac{d}{d\bar{M}} \frac{dL_2}{dA_1}$ 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.

Table 8: Balance: Most important plot characteristics

	Full sample		RD sample					
	Coef. (SE) [p]	Dep. var.	Coef. (SE) [p]					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log area	-0.108 (0.068) [0.114]	-2.381 (1.041) 784	0.043 (0.083) [0.603]	0.089 (0.082) [0.275]	0.094 (0.128) [0.460]	0.074 (0.136) [0.588]		
Own plot	0.025 (0.019) [0.174]	0.875 (0.331) 784	0.048 (0.023) [0.037]	0.043 (0.023) [0.064]	0.040 (0.033) [0.226]	0.033 (0.039) [0.392]	0.039 (0.032) [0.232]	0.029 (0.037) [0.436]
Owned plot >5 years	0.005 (0.014) [0.728]	0.960 (0.197) 585	-0.004 (0.016) [0.811]	-0.003 (0.016) [0.853]	0.012 (0.024) [0.617]	0.033 (0.024) [0.175]	0.011 (0.023) [0.617]	0.030 (0.025) [0.233]
Rented out, farmer	0.013 (0.010) [0.224]	0.033 (0.179) 784	-0.006 (0.013) [0.664]	-0.006 (0.013) [0.645]	-0.026 (0.022) [0.249]	-0.040 (0.025) [0.114]	-0.029 (0.023) [0.222]	-0.041 (0.026) [0.116]
Command area	0.187 (0.032) [0.000]	0.399 (0.491) 784	0.074 (0.039) [0.059]	0.045 (0.037) [0.219]	-0.053 (0.058) [0.360]	-0.079 (0.059) [0.183]		
Terraced	0.017 (0.028) [0.539]	0.626 (0.485) 784	-0.030 (0.035) [0.403]	-0.043 (0.035) [0.225]	-0.099 (0.053) [0.063]	-0.091 (0.055) [0.099]	-0.076 (0.051) [0.134]	-0.058 (0.052) [0.260]
Rented out, comm. farmer	0.035 (0.018) [0.054]	0.081 (0.273) 784	0.017 (0.025) [0.486]	0.008 (0.023) [0.735]	-0.042 (0.040) [0.292]	-0.016 (0.034) [0.638]	-0.036 (0.036) [0.324]	-0.004 (0.031) [0.895]
Omnibus F-stat [p]	5.6 [0.000]		1.8 [0.093]	1.6 [0.132]	1.3 [0.278]	1.5 [0.153]	1.4 [0.209]	1.2 [0.292]
Site FE				X	X		X	
Distance to boundary					X	X	X	X
log area							X	X
MIP log area							X	X
MIP CA							X	X
Spatial FE						X		X

Notes: Balance for most important plot characteristics is presented in this table. Column 2 presents, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 1 and 3 through 8 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Controls are listed below. The final row of each column presents the Omnibus F-stat for the null of balance on all outcomes, with the p-value for the associated test in brackets. Column 1 uses the full sample, while Columns 2 through 8 use the discontinuity sample. Column 4 uses the specification in Equation (7), Column 7 uses the specification in Equation (8), and Column 8 uses the specification in Equation (9).

Table 9: Sample plot shock causes households to substitute labor and input intensive irrigated horticulture away from most important plot

	Sample plot		MIP						
	Coef. (SE) [p]		Dep. var.	Coef. (SE) [p]					
	(1)	(2)		(3)	(4)	(5)	(6)	(7)	(8)
Cultivated									
CA	0.033 (0.031) [0.289]	0.368 (0.483) 2,179	0.049 (0.023) [0.035]	0.038 (0.040) [0.344]	0.004 (0.049) [0.930]	0.085 (0.030) [0.005]	0.076 (0.043) [0.079]	0.059 (0.048) [0.215]	
CA * MIP CA						-0.094 (0.053) [0.078]	-0.089 (0.052) [0.089]	-0.121 (0.056) [0.030]	
Joint F-stat [p]						3.9 [0.021]	2.1 [0.122]	2.7 [0.070]	
Irrigated									
CA	0.202 (0.019) [0.000]	0.114 (0.319) 2,179	-0.019 (0.017) [0.251]	-0.044 (0.026) [0.087]	-0.036 (0.033) [0.270]	0.013 (0.008) [0.123]	-0.004 (0.020) [0.836]	0.010 (0.026) [0.686]	
CA * MIP CA						-0.097 (0.035) [0.006]	-0.094 (0.035) [0.007]	-0.103 (0.045) [0.021]	
Joint F-stat [p]						4.1 [0.019]	3.6 [0.028]	2.7 [0.069]	
Site-by-season FE	X		X	X		X	X		
Distance to boundary				X	X		X	X	
log area				X	X		X	X	
Spatial FE					X			X	
MIP log area				X	X		X	X	
MIP CA				X	X	X	X	X	

Notes: Regression analysis is presented in this table. Column 1 uses outcomes on the sample plot (and replicates analysis in Table 4), while Columns 3 through 8 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 2 presents, for the most important plot associated with sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. For Columns 1 and 3 through 8, Rows “CA” present coefficients on a command area indicator for the sample plot, while Rows “CA * MIP in CA” present coefficients on the interaction of a command area indicator for the sample plot with a command area indicator for the most important plot; standard errors are in parentheses, and p-values are in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Column 3 uses the specification in Equation (7), Column 4 uses the specification in Equation (8), and Column 5 uses the specification in Equation (9). Columns 6 through 8 uses analogous specifications building on Equation (10).

Table 10: Sample plot shock causes households to substitute labor and input intensive irrigated horticulture away from most important plot

	Sample plot		MIP					
	Coef. (SE) [p]	Dep. var.	Coef. (SE) [p]					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Horticulture								
CA	0.180 (0.020) [0.000]	0.109 (0.312) 2,179	-0.016 (0.016) [0.323]	-0.038 (0.024) [0.110]	-0.037 (0.029) [0.206]	0.010 (0.009) [0.244]	-0.004 (0.018) [0.813]	-0.007 (0.023) [0.771]
CA * MIP CA						-0.082 (0.035) [0.020]	-0.080 (0.035) [0.021]	-0.066 (0.044) [0.133]
Joint F-stat [p]						2.9 [0.060]	2.7 [0.070]	1.3 [0.286]
Banana								
CA	-0.134 (0.024) [0.000]	0.199 (0.399) 2,179	0.066 (0.023) [0.004]	0.092 (0.032) [0.004]	0.065 (0.036) [0.072]	0.077 (0.033) [0.021]	0.096 (0.041) [0.019]	0.087 (0.044) [0.047]
CA * MIP CA						-0.013 (0.043) [0.766]	-0.009 (0.042) [0.824]	-0.048 (0.044) [0.275]
Joint F-stat [p]						5.9 [0.003]	4.5 [0.013]	2.0 [0.139]
Site-by-season FE	X		X	X		X	X	
Distance to boundary				X	X		X	X
log area				X	X		X	X
Spatial FE					X			X
MIP log area				X	X		X	X
MIP CA				X	X	X	X	X

Notes: Regression analysis is presented in this table. Column 1 uses outcomes on the sample plot (and replicates analysis in Table 4), while Columns 3 through 8 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 2 presents, for the most important plot associated with sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. For Columns 1 and 3 through 8, Rows “CA” present coefficients on a command area indicator for the sample plot, while Rows “CA * MIP in CA” present coefficients on the interaction of a command area indicator for the sample plot with a command area indicator for the most important plot; standard errors are in parentheses, and p-values are in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and Conley (1999) standard errors are used in specifications with Spatial FE. Column 3 uses the specification in Equation (7), Column 4 uses the specification in Equation (8), and Column 5 uses the specification in Equation (9). Columns 6 through 8 uses analogous specifications building on Equation (10).

Table 11: Sample plot shock causes households to substitute labor and input intensive irrigated horticulture away from most important plot

	Sample plot		MIP						
	Coef. (SE) [p]		Dep. var.	Coef. (SE) [p]					
	(1)	(2)		(3)	(4)	(5)	(6)	(7)	(8)
<hr/>									
HH labor/ha									
CA	69.6 (14.7) [0.000]	66.8 (210.5) 2,166	-11.2 (11.9) [0.351]	-32.2 (20.0) [0.107]	-33.2 (23.8) [0.162]	3.2 (6.2) [0.609]	-13.6 (14.1) [0.338]	-15.4 (19.2) [0.422]	
CA * MIP CA						-41.7 (26.8) [0.120]	-44.1 (23.5) [0.060]	-39.7 (31.2) [0.204]	
Joint F-stat [p]						1.2 [0.290]	1.8 [0.164]	1.1 [0.324]	
<hr/>									
Input exp./ha									
CA	7.4 (1.3) [0.000]	5.6 (28.2) 2,169	-2.1 (1.5) [0.158]	-6.0 (2.7) [0.028]	-6.7 (2.8) [0.017]	0.2 (0.7) [0.805]	-3.3 (1.8) [0.070]	-3.8 (2.1) [0.076]	
CA * MIP CA						-6.3 (3.4) [0.067]	-6.3 (3.2) [0.044]	-6.5 (3.7) [0.079]	
Joint F-stat [p]						1.7 [0.190]	2.6 [0.078]	3.0 [0.050]	
<hr/>									
Hired labor exp./ha									
CA	5.6 (1.9) [0.003]	3.9 (24.6) 2,169	-0.9 (1.3) [0.506]	-1.8 (2.1) [0.404]	-0.5 (2.3) [0.825]	0.8 (1.2) [0.477]	0.2 (2.1) [0.922]	1.5 (2.6) [0.546]	
CA * MIP CA						-4.4 (2.7) [0.099]	-4.7 (2.5) [0.066]	-4.5 (3.1) [0.146]	
Joint F-stat [p]						1.4 [0.255]	1.8 [0.175]	1.1 [0.345]	
Site-by-season FE	X		X	X		X	X		
Distance to boundary				X	X		X	X	
log area				X	X		X	X	
Spatial FE					X			X	
MIP log area				X	X		X	X	
MIP CA				X	X	X	X	X	

Notes: Regression analysis is presented in this table. Column 1 uses outcomes on the sample plot (and replicates analysis in Table 5), while Columns 3 through 8 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 2 presents, for the most important plot associated with sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. For Columns 1 and 3 through 8, Rows “CA” present coefficients on a command area indicator for the sample plot, while Rows “CA * MIP in CA” present coefficients on the interaction of a command area indicator for the sample plot with a command area indicator for the most important plot; standard errors are in parentheses, and p-values are in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Column 3 uses the specification in Equation (7), Column 4 uses the specification in Equation (8), and Column 5 uses the specification in Equation (9). Columns 6 though 8 uses analogous specifications building on Equation (10).

Table 12: Impacts of access to irrigation are explained by transition to horticulture from bananas

	Sample plot						
	Dep. var.	Coef. (SE) [p]					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
HH labor/ha	59.5 (201.4) 2,523	69.6 (14.7) [0.000]	70.8 (17.5) [0.000]	76.9 (20.7) [0.000]	12.0 (10.2) [0.239]	30.0 (13.2) [0.023]	29.0 (14.8) [0.051]
Input exp./ha	2.5 (17.4) 2,527	7.4 (1.3) [0.000]	6.3 (1.5) [0.000]	4.3 (1.8) [0.019]	0.6 (0.9) [0.509]	1.2 (1.2) [0.330]	-1.5 (1.4) [0.302]
Hired labor exp./ha	3.7 (25.6) 2,527	5.6 (1.9) [0.003]	3.7 (2.1) [0.082]	3.2 (2.6) [0.221]	1.7 (1.5) [0.275]	0.8 (2.0) [0.681]	0.3 (2.5) [0.894]
Site-by-season FE		X	X		X	X	
Distance to boundary			X	X		X	X
log area			X	X		X	X
Spatial FE				X			X
Crop					X	X	X

Notes: Regression analysis is presented in this table. All columns restrict to observations during the dry season. Column 1 presents, for sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 2 through 7 present regression coefficients on a command area indicator, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Columns 2 and 5 use the specification in Equation (1). Columns 3 and 6 use the regression discontinuity specification in Equation (2). Columns 4 and 7 use the spatial fixed effects specification in Equation (3). Columns 5, 6, and 7 also include controls for cultivation, horticulture, and bananas.

Table 13: Impacts of sample plot shock on most important plot are on both extensive and intensive margins

	MIP						
	Dep. var.	Coef. (SE) [p]					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
HH labor/ha	66.8 (210.5) 2,166	-11.2 (11.9) [0.351]	-32.2 (20.0) [0.107]	-33.2 (23.8) [0.162]	-9.2 (9.1) [0.312]	-25.0 (14.1) [0.077]	-21.4 (16.8) [0.203]
Input exp./ha	5.6 (28.2) 2,169	-2.1 (1.5) [0.158]	-6.0 (2.7) [0.028]	-6.7 (2.8) [0.017]	-1.5 (1.2) [0.227]	-4.6 (2.1) [0.032]	-4.9 (2.2) [0.023]
Hired labor exp./ha	3.9 (24.6) 2,169	-0.9 (1.3) [0.506]	-1.8 (2.1) [0.404]	-0.5 (2.3) [0.825]	-0.8 (1.2) [0.483]	-1.4 (1.9) [0.482]	0.2 (2.1) [0.915]
Site-by-season FE		X	X		X	X	
Distance to boundary			X	X		X	X
log area			X	X		X	X
MIP log area			X	X		X	X
MIP CA			X	X		X	X
Spatial FE				X			X
Crop					X	X	X

Notes: Regression analysis is presented in this table. All columns restrict to observations during the dry season. Column 1 presents, for the most important plot associated with sample plots in the main discontinuity sample that are outside the command area, the mean of the dependent variable, the standard deviation of the dependent variable in parentheses, and the total number of observations. Columns 2 through 7 present regression coefficients on a command area indicator for the sample plot, with standard errors in parentheses, and p-values in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. Columns 2 and 5 use the specification in Equation (7). Columns 3 and 6 use the regression discontinuity specification in Equation (8). Columns 4 and 7 use the spatial fixed effects specification in Equation (9). Columns 5, 6, and 7 also include controls for cultivation, horticulture, and bananas.

Table 14: Larger and poorer households do not substitute away from most important plot in response to sample plot shock

	MIP				MIP		
	Coef. (SE) [p]				Coef. (SE) [p]		
	(1)	(2)	(3)		(1)	(2)	(3)
Cultivated				Horticulture			
CA	-0.046 (0.073) [0.526]	-0.069 (0.086) [0.424]	-0.188 (0.098) [0.056]	CA	-0.082 (0.041) [0.045]	-0.107 (0.046) [0.020]	-0.129 (0.047) [0.006]
CA * # of HH members	0.019 (0.013) [0.158]	0.021 (0.013) [0.112]	0.039 (0.014) [0.007]	CA * # of HH members	0.013 (0.008) [0.109]	0.014 (0.007) [0.061]	0.019 (0.007) [0.012]
CA * Asset index	-0.007 (0.028) [0.814]	-0.013 (0.027) [0.620]	-0.043 (0.032) [0.181]	CA * Asset index	-0.019 (0.020) [0.353]	-0.015 (0.017) [0.384]	-0.030 (0.021) [0.156]
Joint F-stat [p]	2.4 [0.072]	1.5 [0.213]	3.5 [0.015]	Joint F-stat [p]	1.5 [0.225]	1.8 [0.147]	2.6 [0.050]
Irrigated				Banana			
CA	-0.071 (0.043) [0.098]	-0.097 (0.048) [0.046]	-0.121 (0.051) [0.017]	CA	0.036 (0.064) [0.573]	0.052 (0.065) [0.418]	-0.052 (0.083) [0.532]
CA * # of HH members	0.010 (0.009) [0.227]	0.011 (0.007) [0.155]	0.017 (0.008) [0.030]	CA * # of HH members	0.006 (0.011) [0.596]	0.008 (0.011) [0.485]	0.023 (0.015) [0.110]
CA * Asset index	-0.022 (0.020) [0.256]	-0.018 (0.016) [0.277]	-0.035 (0.020) [0.077]	CA * Asset index	0.009 (0.024) [0.696]	-0.003 (0.023) [0.900]	-0.012 (0.026) [0.661]
Joint F-stat [p]	1.3 [0.284]	1.5 [0.214]	2.3 [0.079]	Joint F-stat [p]	3.0 [0.031]	3.1 [0.026]	2.1 [0.094]
# of HH members	X	X	X	# of HH members	X	X	X
Asset index	X	X	X	Asset index	X	X	X
Site-by-season FE	X	X		Site-by-season FE	X	X	
Distance to boundary		X	X	Distance to boundary		X	X
log area		X	X	log area		X	X
MIP log area		X	X	MIP log area		X	X
MIP CA		X	X	MIP CA		X	X
Spatial FE			X	Spatial FE			X

Notes: Regression analysis is presented in this table. All columns use outcomes on most important plots and restrict to observations during the dry season.. Rows “CA” present coefficients on a command area indicator for the sample plot, while Rows “CA * W” present coefficients on the interaction of a command area indicator for the sample plot with a household characteristic W; standard errors are in parentheses, and p-values are in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. The Row “Joint F-stat [p]” presents F-statistics for the null that all 3 coefficients are 0, with the p-value for the associated test in brackets. Columns 1, 2, and 3 use regression specifications building on Equation (10) following Equations (7), (8), and (9), respectively.

Table 15: Larger and poorer households do not substitute away from most important plot in response to sample plot shock

	MIP				MIP		
	Coef. (SE) [p]				Coef. (SE) [p]		
	(1)	(2)	(3)		(1)	(2)	(3)
<hr/>				<hr/>			
HH labor/ha				Hired labor exp./ha			
CA	-78.5 (32.0) [0.014]	-82.7 (34.8) [0.018]	-94.6 (38.9) [0.015]	CA	-4.1 (2.9) [0.155]	-4.5 (3.6) [0.216]	-2.1 (3.5) [0.551]
CA * # of HH members	13.4 (5.5) [0.015]	10.1 (4.5) [0.025]	12.5 (4.5) [0.006]	CA * # of HH members	0.6 (0.5) [0.201]	0.5 (0.5) [0.273]	0.3 (0.5) [0.539]
CA * Asset index	-22.7 (12.7) [0.074]	-12.4 (10.2) [0.226]	-24.0 (12.7) [0.060]	CA * Asset index	-0.1 (1.4) [0.968]	0.3 (1.4) [0.856]	-0.3 (1.4) [0.813]
Joint F-stat [p]	2.1 [0.099]	2.0 [0.122]	2.9 [0.033]	Joint F-stat [p]	0.8 [0.488]	0.6 [0.592]	0.1 [0.937]
<hr/>				<hr/>			
Input exp./ha				# of HH members	X	X	X
CA	-6.2 (3.5) [0.075]	-9.1 (4.2) [0.031]	-10.3 (4.1) [0.013]	Asset index	X	X	X
CA * # of HH members	0.8 (0.6) [0.161]	0.6 (0.5) [0.237]	0.7 (0.5) [0.185]	Site-by-season FE	X	X	X
CA * Asset index	-3.2 (1.8) [0.076]	-2.5 (1.6) [0.117]	-3.9 (1.7) [0.025]	Distance to boundary		X	X
Joint F-stat [p]	1.4 [0.239]	1.9 [0.128]	2.6 [0.051]	log area		X	X
<hr/>				<hr/>			
# of HH members	X	X	X	MIP log area		X	X
Asset index	X	X	X	MIP CA		X	X
Site-by-season FE	X	X	X	Spatial FE			X
Distance to boundary		X	X				
log area		X	X				
MIP log area		X	X				
MIP CA		X	X				
Spatial FE			X				

Notes: Regression analysis is presented in this table. All columns use outcomes on most important plots and restrict to observations during the dry season.. Rows “CA” present coefficients on a command area indicator for the sample plot, while Rows “CA * W” present coefficients on the interaction of a command area indicator for the sample plot with a household characteristic W; standard errors are in parentheses, and p-values are in brackets. Robust standard errors are clustered at the nearest water user group level in specifications without Spatial FE, and [Conley \(1999\)](#) standard errors are used in specifications with Spatial FE. The Row “Joint F-stat [p]” presents F-statistics for the null that all 3 coefficients are 0, with the p-value for the associated test in brackets. Columns 1, 2, and 3 use regression specifications building on Equation (10) following Equations (7), (8), and (9), respectively.

Table 16: Minikits do not cause increased adoption of horticulture, strong positive selection into minikit takeup

	Minikit takeup		Horticulture	
	(1)	(2)	(3)	(4)
Assigned minikit	0.398 (0.038) [0.000]	0.395 (0.044) [0.000]	0.035 (0.041) [0.396]	0.052 (0.042) [0.221]
Minikit saturation	-0.047 (0.056) [0.394]	-0.064 (0.057) [0.260]	-0.078 (0.054) [0.149]	-0.067 (0.054) [0.218]
Horticulture (2016 Dry)		0.046 (0.049) [0.345]		0.306 (0.053) [0.000]
Assigned minikit * Horticulture (2016 Dry)		0.131 (0.068) [0.052]		-0.019 (0.070) [0.788]
# of lotteries entered	X	X	X	X
O&M treatment	X	X	X	X
Zone FE	X	X	X	X
# of observations	910	762	838	727
# of clusters	187	170	182	167

Notes: Regression analysis is presented in this table. All columns use outcomes on sample plots. Each row presents coefficients, with robust standard errors clustered at the water user group level in parentheses, and p-values in brackets. “Assigned minikit” is an indicator for whether the household was assigned to receive a minikit, “Minikit saturation” is the probability of minikit assignment that was assigned to the water user group of the household’s sample plot, and “Horticulture (2016 Dry)” is an indicator that the household planted horticulture on their sample plot in 2016 Dry.