Essays on Irrigation

By

John Ashton Loeser

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jeremy Magruder, Chair
Professor Benjamin Faber
Professor Aprajit Mahajan
Professor Elisabeth Sadoulet

Spring 2019

Essays on Irrigation

Abstract

Essays on Irrigation

by

John Ashton Loeser

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Jeremy Magruder, Chair

In the first chapter, I estimate an elasticity of irrigation adoption to its gross returns in rural India. Many approaches to estimating this elasticity fail when agents select into adopting irrigation on heterogeneous gross returns and costs. I develop a novel approach to correct for selection using two instrumental variable estimators that can be implemented with aggregate data on gross revenue and adoption of irrigation. I use climate and soil characteristics as an instrument for gross returns to irrigation, and hydrogeology as an instrument for irrigation to correct for selection. I estimate that a 1% increase in the gross returns to irrigation causes a 0.7% increase in adoption of irrigation. I use this elasticity to infer changes in profits from changes in adoption of irrigation caused by shocks to its profitability, and to conduct counterfactuals. First, groundwater depletion from 2000-2010 in northwestern India permanently reduced economic surplus by 1.2% of gross agricultural revenue. Second, I evaluate a policy that optimally reduces relative subsidies for groundwater irrigation in districts with large negative pumping externalities, while holding total subsidies fixed. Under the policy, depletion caused by subsidies decreases by 16%, but farmer surplus increases by only 0.07% of gross agricultural revenue.

In the second chapter, co-authored with Maria Jones, Florence Kondylis, and Jeremy Magruder, we examine the returns to newly-constructed hillside irrigation schemes in Rwanda using a very granular spatial regression discontinuity design. We find that irrigation enables dry season horticultural production which is associated with large increases in labor and input usage and boosts on-farm yields and cash profits by 70%. At the same time, irrigation use remains limited after 4 years. We leverage the spatial discontinuity in access to irrigation to develop a test for separation failures based on farmer behavior on other plots and conclude that separation failures restrain technology adoption. Unlike existing separation tests, our test allows us to distinguish the role of labor constraints from credit and insurance constraints; we find robust evidence that labor constraints limit adoption.

In the third chapter, I develop a new approach to quantify the welfare gains from risky technologies for intertemporal substitution, ranging from agricultural technologies to financial products. Traditionally, these welfare gains are measured either by the technology's effect on a welfare proxy or by estimating a structural model. Using a welfare proxy may be suboptimal due to noise in measurement and the challenge of converting estimated effects into a money metric, while structural approaches may require strong functional form assumptions and depend on unexpected moments of the data. In contrast, despite some drawbacks, Marshallian consumer surplus is frequently used as a metric for the welfare gains from access to a new product in a static setting, and with sufficient variation in prices may be relatively easy to precisely estimate. I show that under a broad class of models of dynamic optimization which nest Deaton (1991), Marshallian consumer surplus is a reasonable welfare metric for access to an intertemporal substitution technology. I demonstrate how to calculate it, and apply the approach to three experiments which randomly varied either interest rates or prices: I compare the consumer surplus from grants of index insurance in Ghana to their actuarially fair value, I calculate the surplus to households from access to a leading MFI in Mexico, and I bound the foregone consumer surplus due to inattention to the Savers' Credit among households in the United States. In all cases, the calculation is straightforward, transparent, and can be represented graphically as a "welfare triangle".

Acknowledgements

This dissertation was possible because of the support of my advisors, dissertation committee members, coauthors, colleagues, friends, and family. Jeremy provided uncountably many hours of advice for this dissertation, and almost all of my research interests have sprung from our work together. Aprajit provided helpful support in identifying contributions and set a high standard for attention to detail in the theory. Betty taught me the importance of the clarity and precision of my work and continually inspires me to improve along these dimensions, and her optimism has kept me striving forward. Ben provided valuable guidance on this dissertation, and his teaching inspired my interest in many subfields of economics.

Through various projects in Rwanda, I have had the opportunity to collaborate with innumerable colleagues who have informed this dissertation, to say nothing of the many friends with whom I've had conversations that have inspired my research more broadly. In particular, thank you Esdras Byiringiro, Maria Jones, Saahil Karpe, Innocent Musabyimana, and Christophe Ndahimana, and I look forward to our continued work together (*turi kumwe*). Special thanks to Florence Kondylis for valued mentorship and collaborations on these projects, and for invaluable advice and support during critical months.

Many colleagues provided emotional and technical support during my time in Berkeley. I received particularly helpful feedback and encouragement from Thibault Fally, Fred Finan, Marco Gonzalez-Navarro, Supreet Kaur, Ted Miguel, Sofia Villas-Boas, and Chris Walters. I have had the privilege of a wonderful set of peers during this time; in particular, thank you Ceren Baysan, Tamma Carleton, Narayan Das, Kwabena Donkor, Alessandra Fenizia, Erin Kelley, Greg Lane, Yotam Shem-Tov, Deirdre Sutula, Itai Trilnick, and Shaoda Wang. Special thanks to Carmen Karahalios and Diana Lazo for being there during all the most difficult times.

My family has always been supportive of my decision to move across the country to pursue research in development economics. Thank you Mom, Dad, and Edward for accepting my less than frequent communications and yet encouraging me to push forward.

Contents

Chapter 1

The treatment effect elasticity of demand: Estimating the welfare losses from groundwater depletion in India

1.1 Introduction

A common parameter of interest in economics is the elasticity of adoption of a binary treatment to its treatment effect. In the classic Roy (1951) model, workers' relative potential wages across sectors determine their sectoral choice. Similarly, the effect of the skill premium on high school graduates' decisions to attend college is an important input to models of directed technical change (Acemoglu, 1998), as is the effect of firms' potential profits on entry decisions to many models in industrial organization and trade (Melitz, 2003). An estimate of this elasticity is useful both for counterfactuals, such as agents' responses to a tax, or for welfare analysis, such as inferring lost surplus from behavioral responses to a shock.

Selection complicates consistent estimation of this elasticity when economic agents select into treatment on both idiosyncratic treatment effects and perceived costs of adopting treatment. Existing approaches to estimating this elasticity require assuming selection on observables (as noted by Heckman (1979)), imposing strong parametric assumptions (Heckman, 1979), or access to sufficiently high powered instruments to estimate a control variable nonparametrically (Ahn & Powell, 1993; Das et al., 2003; Eisenhauer et al., 2015). This contrasts starkly with estimating treatment effects, where linear instrumental variables estimates a local average treatment effect in the presence of selection on unobservables and without imposing any parametric assumptions (Imbens & Angrist, 1994).

In this paper, I focus on the elasticity of irrigation adoption to its gross returns in India. Irrigation is of first order importance in Indian agriculture. From 1960 to 2010, during India's Green Revolution, the irrigated share of agricultural land grew from 18% to 54%; over 60% of this growth came from the expansion of tubewells for groundwater extraction. This extraction is not benign; Rodell et al. (2009) find extraction caused water tables in northwest India to fall 3.3m from 2000-2010, or 0.21 standard deviations of depth to water table across districts. Falling water tables, by increasing the costs of groundwater irrigation, have been shown to increase poverty (Sekhri, 2014), decrease land values (Jacoby, 2017), and cause outmigration and decrease area under irrigation (Fishman et al., 2017). This has important implications for economic efficiency: groundwater extraction is a classic example of "tragedy of the commons", as farmers do not internalize the increase in pumping costs their extraction causes for neighboring farmers through declining water tables (Jacoby, 2017).

Despite potentially large externalities from groundwater extraction, formulating optimal policy responses to declining water tables in India is difficult for two reasons. First, the elasticity of irrigation to many counterfactual policies is unknown. Second, empirical estimates of the impacts per unit decline in water tables on agricultural profits are not available, as agricultural profits in developing countries are notoriously difficult to measure reliably.[1] An estimate of the elasticity of groundwater irrigation for agriculture to its gross returns would solve both of these challenges. For the first, responses of irrigation to a policy are proportional to the elasticity of irrigation to its gross returns times the effect of the policy on relative profits under irrigation. For the second, effects of declining water tables on adoption of irrigation are proportional to their effects on farmer profits times the elasticity of irrigation to its gross returns.

I estimate an elasticity of irrigation adoption to its gross returns. To do so, I first build a generalized Roy model where farmers adopt irrigation if their gross returns to irrigation are greater than their costs of irrigation; this allows for selection into irrigation on unobservable heterogeneity in gross returns, and I make no parametric assumptions about the joint distribution of gross returns and costs. Under this model, I show that a linear instrumental variable estimator using an instrument for gross revenue under irrigation estimates the sum of weighted averages of gross returns to irrigation (a "local average treatment effect") and inverse semielasticities of demand for irrigation (a "local average surplus effect"). This builds on formulas for instrumental

[1] Challenges in the measurement of agricultural profits in developing countries are discussed at length in Foster & Rosenzweig (2010) and Karlan et al. (2014a), among others. To list two, first, absent administrative data, long household surveys are required to capture the full set of inputs used in smallholder agriculture. Second, smallholder agriculture intensively uses non-marketed inputs (primarily household labor) which are difficult to value.

variables bias from Angrist et al. (1996); here, the "bias" is the estimand of interest, a local average surplus effect. Existing results imply the local average treatment effect, and therefore the local average surplus effect, is identified with a continuous instrument for irrigation (Heckman & Vytlacil, 2005) or bounded with a discrete instrument for irrigation (Mogstad et al., 2017).[2] Under stronger assumptions, which still allow sorting on unobserved heterogeneity in both gross returns and costs, I show weighted linear instrumental variables with an instrument for irrigation is a consistent estimator of this local average treatment effect; the weights adjust the compliers to the instrument for irrigation to match the compliers to the instrument for gross revenue under irrigation on observables.[3]

The generalized Roy model I use to study selection into irrigation on its gross returns builds on a long literature surveyed in Heckman & Vytlacil (2007a,b); these models have been used to study sectoral choice and wage premia (Roy (1951)), education and skill premia (Willis & Rosen, 1979), and, closest to my setting, hybrid maize seed and its gross returns (Suri, 2011). I build most closely on Eisenhauer et al. (2015), who establish nonparametric identification of agents' willingness to pay for treatment (irrigation) from an instrument for treatment and an instrument for treatment effects (gross returns to irrigation). I instead assume the existence of an instrument for potential outcome under treatment (gross revenue under irrigation), and establish nonparametric identification of the inverse semielasticity of adoption of treatment to the treatment effect under weaker conditions. These weaker conditions are the union of the assumptions of the standard local average treatment effect framework (Imbens & Angrist, 1994) and the assumptions needed for point identification of economic surplus from a change in costs when potential outcomes are independent of treatment conditional on observables (Willig, 1978; Small & Rosen, 1981).

I estimate that a 1% increase in the gross returns to irrigation causes a 0.7% increase in the irrigated share of agricultural land. I estimate this elasticity using climate and soil characteristics as an instrument for gross revenue under irrigation, and using hydrogeology as an instrument for irrigation. I use this elasticity to infer changes in profits from changes in adoption of irrigation caused by shocks to profitability of irrigation. Fishman et al. (2017) estimate the effect of declining water tables on adoption of irri-

[2]An extra monotonicity assumption is needed; the marginal farmers induced to adopt irrigation by the instrument for irrigation and the instrument for gross revenue under irrigation are assumed to be the same conditional on the propensity score and observables; I discuss this assumption in Section **??**. For point identification, conditions on the conditional support of the propensity score are needed as well.

[3]This approach generalizes assumptions made in Angrist & Fernandez-Val (2010) under which linear instrumental variable estimators can be reweighted on observables to recover the same local average treatment effect. In doing so it contributes to a number of recent papers which enables comparison of compliers to different instruments under monotonicity by estimating marginal treatment effects (Kowalski, 2016; Arnold et al., 2018; Mountjoy, 2018)) or bounding local average treatment effects (Mogstad et al., 2017).

gation; with their estimate, my estimate of this elasticity implies that that the 3.3m decline in depth to groundwater observed in northwest India from 2000-2010 decreased economic surplus by 1.2% of gross revenue per hectare. These losses are large; for comparison, Government of India (2018) anticipate losses in India due to climate change of 1.8%/decade over the next century. I compare my estimate to a simple physics based back-of-the-envelope that considers losses only from farmers' increased electricity costs; my estimate is six times as large as that back of the envelope, consistent with farmers' cost share of electricity in irrigation.

I incorporate my estimate of the economic costs of declining water tables into a model of optimal taxation of electricity for groundwater irrigation, following Allcott et al. (2014). A social planner chooses subsidies for electricity use, trading off the value of subsidies as a transfer to farmers with their deadweight loss and the negative externalities generated from induced marginal extraction. These externalities vary across districts, as water tables fall more rapidly in thinner aquifers, and these falls are experienced by more farmers when a larger share of land is irrigated. I calibrate the model using data on groundwater extraction and aquifer characteristics for districts in Rajasthan, the state in northwest India most known for falling water tables. I find the observed electricity subsidy in Rajasthan is responsible for a 1.5 meter fall in water tables, 46% of the observed rate of decline in northwest India. However, this subsidy increases farmer surplus by 5.9% of gross agricultural revenue, and on the margin implies the social planner is paying 1.56 Rs for 1.00 Rs in surplus transferred to farmers, not far from a similar shadow cost in the US from Hendren (2016). Externalities are important: of the 1.56 Rs, 0.31 Rs are lost to deadweight loss, while 0.25 Rs are lost to negative externalities from induced marginal extraction.

I consider a counterfactual where the social planner optimally varies subsidies across districts, relatively decreasing subsidies in high externality districts, while holding fixed total subsidy payments. This alternate policy reduces the effect of subsidies on water table declines by 16%, but it increases farmer surplus by only 0.07% of gross agricultural revenue. This increase in surplus is small in magnitude relative to the reallocation of surplus from subsidies from farmers in districts with high externalities to farmers in districts with low externalities, consistent with political economy motives for electricity subsidies (Dubash, 2007). However, the magnitude of surplus gains, and more generally the magnitude of externalities, are much larger under smaller calibrations of the discount rate: while transfers and deadweight loss are static, falls in the water table are permanent in the districts I consider, implying the social planner must trade off transfers to farmers today with lost profits for farmers in future decades.

In providing these estimates, I build on a deep literature on the economics of irrigation. Most directly, I contribute to existing results of the impacts of surface water irrigation (Duflo & Pande, 2007) and declining water tables (Sekhri, 2014; Fishman et al., 2017) on welfare proxies in India, and hedonic estimates of the value of access to groundwater in India (Jacoby, 2017) and in the US (Schlenker et al., 2007). In contrast, I estimate sufficient parameters for many optimal policy calculations: the economic losses from a 1 meter decline in the water table, and the elasticity of demand for irrigation to its gross returns. In this sense, I build on estimates of the elasticity of groundwater extraction to electricity subsidies (Badiani & Jessoe, 2017) and output subsidies for water intensive crops (Chatterjee et al., 2017). I use this estimated elasticity to build on the optimal control literature, summarized in Koundouri (2004a), and applied in India by Sayre & Taraz (2018); a large body of work has used complicated, calibrated dynamic models of management of aquifers to characterize optimal policy.[4] Contributing to this literature, I take a sufficient statistics approach, building a simple public economic model following Chetty (2009) and Allcott et al. (2014): empirical estimates of elasticities are used where possible, and calibrated parameters enter transparently into counterfactuals.

The rest of the paper is organized as follows. Section 1.2 describes the data used and the context. Section 1.3 presents the model, including results on identification and estimation. Section 1.4 describes the empirical strategy I use. Section 1.5 presents the main results, including the impacts of groundwater depletion on rural surplus, and Section 1.6 discusses their robustness. Section 1.7 considers optimal subsidies for electricity for groundwater irrigation, building on results from Section 1.5. Section 1.8 concludes.

1.2 Data and context

1.2.1 Context

India's Green Revolution, starting in the 1960's, was a time of rapid growth in agricultural productivity, driven by increased adoption of new high yielding varieties of seeds, fertilizers, pesticides, and irrigation (Evenson & Gollin, 2003). Irrigation was a particularly important component: large investments were made in the expansion of surface water irrigation, with over 2,400 large dams constructed from 1971-1999 (Duflo & Pande, 2007), but the majority of growth of irrigation was ground water irrigation (Gandhi & Bhamoriya, 2011). The irrigated share of agricultural land in India expanded from 18% to 54% from 1960 to 2008, while the share of agricultural land irrigated using tubewells grew from 0% to 22%, accounting for 63% of the overall growth in irrigation. Reduced form evidence suggests that access to groundwater has large impacts on social welfare (Sekhri, 2014; Fishman et al., 2017; Jacoby, 2017) and is an important driver of

[4]Results from these models can be sensitive to the calibration: Gisser & Sanchez (1980) famously find small gains from optimal policy relative to laissez faire, but Koundouri (2004b) argue their findings are driven by their steep calibrated marginal benefit curves, while Brozović et al. (2010) argue they are driven by the characteristics of the aquifer they study.

adoption of modern agricultural technologies (Sekhri, 2014). This evidence suggests a large share of agricultural productivity growth during the Green Revolution may have been caused by access to groundwater.

Groundwater is stored in underground aquifers, which are underground layers of permeable rock or other materials that hold water. The meters below ground level at which groundwater is available is often referred to as the depth to water table, and varies both across aquifers and within aquifer. For agriculture in India, as in much of the world, this groundwater is typically extracted using tubewells. In a tubewell, a narrow pipe, typically PVC or stainless steel, is bored into the ground, fitted with a strainer cap, and installed with a pump used to pump the water to the surface. Drilling tubewells is costly: according to the 2007 Minor Irrigation Census, the fixed cost of infrastructure for groundwater irrigation in the average district was 26,600 Rs/ha, just over 1 year of agricultural revenue per hectare. This cost varies substantially across districts, with a coefficient of variation of 0.55. This variation is partially driven by the accessibility of groundwater. At greater depths to water table, wells must be drilled deeper, which is more costly (Jacoby, 2017). Additionally, at these lower depths, more expensive and more powerful pumps are required (Sekhri, 2014). Moreover, different types of soils can store different quantities of water, and vary in their permeability. These hydrogeological characteristics affect the rate at which groundwater resources can be extracted that balances natural rates of recharge ("potential aquifer yield" or "safe yield"), the rate at which the water table falls per unit of water extracted ("specific yield"), and the number of wells required per unit of water extracted (Fishman et al., 2017).

Although some of this variation in accessibility of groundwater is driven by exogenous hydrogeological characteristics of the districts, human activity can impact this accessibility. In many districts, ancient groundwater resources are trapped in confined aquifers; these resources are exhaustible. Rodell et al. (2009) use satellite data to show declining water tables in northwestern India, while Suhag (2016) show that the Indian Central Groundwater Board's calculations based on hydrology models imply overexploitation of groundwater resources in the same region. Appendix Figure A.1 shows that this overexploitation (high withdrawals of groundwater as a percentage of natural rates of recharge) is most prevalent in states that experienced the largest increases in agricultural productivity during the Green Revolution, highlighting the link between agricultural productivity and groundwater extraction. In many places, declining water tables are believed to have significantly increased costs of groundwater extraction (Fishman et al., 2017; Jacoby, 2017). On the other hand, rainwater capture and surface water irrigation have the potential to replenish groundwater reserves and reduce dependency on groundwater (Sekhri, 2013).

This decline has been accelerated by implicit subsidies for groundwater irrigation. Most significantly, most states in India do not have volumetric pricing of electricity, but instead charge pump capacity fees. These fees partially substitute for volumetric pricing, since many farmers pump groundwater whenever electricity is available during the growing seasons. However, the levels of fees correspond to large subsidies for electricity, ranging from 52% to 100% subsidies (Fishman et al., 2016). Badiani & Jessoe (2017) use panel variation in these subsidies to estimate an elasticity of water use to the price of electricity of -0.18, suggesting these subsidies contribute meaningfully to declining water tables. However, they point out that this inelastic demand for electricity suggests limited deadweight loss from the subsidies. Since a commonly stated motivation for subsidies is as a transfer to farmers (Dubash, 2007), a social planner who places a high value on marginal consumption by farmers, potentially due to a lack of availability of other policy instruments for making such transfers, might find it optimal to trade off a small deadweight loss to increase transfers to farmers. Moreover, subsidies may correct for the presence of market power in water markets, which might cause socially suboptimal rates of groundwater extraction (Gine & Jacoby, 2016).

In addition to traditional concerns of inefficiency due to subsidies or other wedges, rates of groundwater extraction may be higher than is socially optimal due to negative externalities in pumping groundwater. As farmers extract groundwater, water is drawn from nearby parts of the aquifer, decreasing the water table for neighboring farmers (Theis, 1935), and increasing their costs of extracting groundwater. In the presence of such externalities, farmers will not internalize the increased costs their pumping causes to other farmers. Jacoby (2017) suggests externalities may be particularly important in confined aquifers in India; wells are frequently tightly clustered, and interference between wells is a concern, especially during the dry season.

An estimate of the magnitude of this externality is necessary to determine an optimal tax, or subsidy, for groundwater irrigation. To calculate this externality, one can decompose it into two terms. First, increased pumping of groundwater causes a decline in the water table. The impact of increased pumping on the water table varies significantly across aquifers: pumping one cubic meter of water causes the water table to decline by as much as 20,000 cubic meters in thin, confined aquifers, and by as little as 5 cubic meters in thick, unconfined aquifers (Gisser & Sanchez, 1980; Brozović et al., 2010).

Second, these declines in the water table cause decreases in the profitability of irrigated agriculture, as the cost of groundwater extraction increases. These increases in costs are an externality: they are almost completely experienced by farmers other than the farmer extracting the unit of water. Estimating this increase in costs is hard: costs are notoriously hard to observe in agricultural data (Foster & Rosenzweig, 2010; Karlan et al., 2014a), and as a result empirical estimates of the economic costs of declining water tables are unavailable. In India, past work has estimated impacts of declining water tables on welfare proxies, including poverty headcount (Sekhri, 2014) and outmigration (Fishman et al., 2017). However, calculating the externality requires an estimate of the economic damages from a unit decline in water tables. Existing approaches to estimating this have focused largely on the United States, and have typically used hedonic regressions (see Koundouri (2004a) for a review); these approaches may not be feasible in developing country settings such as India, where the assumption of frictionless land markets and full information is less likely to hold.[5]

1.2.2 Data

I merge data from multiple sources on agriculture in India. Since district boundaries in India have changed multiple times over the past century, all analysis is done using 1961 state and district boundaries. Descriptive statistics for all variables used in analysis are presented in Table 1.1.

Primary agricultural outcomes come from two sources. First, I merge together the World Bank India Agriculture and Climate Data Set, which contains data from 1956-1987, with the ICRISAT Village Dynamics in South Asia Macro-Meso Database, which contains data from 1966-2011. I refer to this merged dataset as “Ag ’56-’11”.[6] The merged dataset contains annual district level data on crop specific land allocations (rainfed and irrigated), prices, and yields. I use this to construct an imbalanced panel of 222 districts in 11 states from 1956-2011 of agricultural revenue per hectare and irrigated share of agricultural land. While more districts are observed in this data set, I restrict to districts which appear in all primary data sets used for analysis to maintain comparability across specifications.[7] For much of the analysis, I restrict to the most recent 5 year cross section in this data set.

I supplement this with the 2012 Agricultural National Sample Survey, which included questions on household level land allocations and agricultural production by crop, crucially both on irrigated and rainfed land; I refer to these observations as plots. The data also contain household level expenditures on agricultural inputs by category. I refer to this dataset as “NSS ’12”. 35,200 households were surveyed, and the survey is intended to be representative at the district level. The sampling of villages from which

[5]Many studies have also used contingent valuation approaches, which can be severely biased. A noteable alternative approach is taken by Hagerty (2018), who studies water markets in the United States. However, they estimate the willingness to pay for one unit of water, which is different from the economic costs of a one unit decline in water tables. One notable exception is Jacoby (2017), who applies a hedonic regression in India to estimate the economic value of having a borewell using exogenous drilling failures as an instrument. Since the presence of a functioning borewell is easily observable, the assumptions underpinning a hedonic regression are likely to hold. However, this estimate cannot be converted into an estimate of the economic costs of a one unit decline in water tables without strong assumptions.

[6]The former dataset has been used by many papers analyzing agriculture in India, including Duflo & Pande (2007) and Sekhri (2014) studying irrigation, while the latter dataset has been used by Allen & Atkin (2015) among others.

[7]Most notably, this restriction drops Chhattisgarh, Jharkhand, and West Bengal.

surveyed households were selected was stratified on share of village land irrigated; because this stratification is correlated with treatment (irrigation), I use survey weights in all analysis with this data. Moreover, to maintain comparability with Ag '56-'11, I weight plots by area, I restrict to crops observed in Ag '56-'11, and I reweight districts so each district receives the same weight. Both revenue per hectare and input expenditures per hectare are noisily measured at high quantiles; I winsorize them at 100,000 Rs/ha (95th percentile for revenue per hectare, 99th percentile for input expenditures per hectare).

For data on irrigation technologies, I use the 2007 Minor Irrigation Census. This survey censuses minor irrigation schemes (culturable command area less than 2000 hectares), which account for 65% of irrigated area and almost all groundwater irrigation. I refer to this dataset as "Irr '07". In this, I observe district level counts of minor irrigation schemes by type (dugwell, shallow tubewell, deep tubewell, surface flow scheme, surface lift scheme), hectares of potential created and used for surface water and ground water schemes, and counts of ground and surface water schemes by cost.[8]

I use potential aquifer yield as my instrument for costs of irrigation, a measure of the sustainable rate of extraction of groundwater from a typical tubewell. I constructed this measure by georeferencing a hydrogeological map of India from the Central Ground Water Board (CGWB) which categorizes all land by potential aquifer yield and aquifer type. The measure ranges from 0 L/s to 40 L/s.[9] In all analysis I divide by 40 to normalize this measure to range from 0 to 1, and I plot variation in this measure across districts in Panel (a) of Figure 1.1.

I use a measure of log relative potential irrigated crop yield as my instrument for potential gross revenue under irrigation. For data on potential crop yield, I use the FAO GAEZ database; this source is discussed at length in Costinot et al. (2016). Among other products, it includes constructed measures of potential yields under 5 input scenarios (low rainfed, intermediate rainfed/irrigated, high rainfed/irrigated) based on climate and soil characteristics. I construct potential rainfed crop yield as the weighted average of potential crop yields under the intermediate rainfed scenario. I construct relative potential irrigated crop yield as the ratio of the weighted average of potential

[8] I observe 5 categories, corresponding to [0 Rs., 10,000 Rs.), [10,000 Rs., 50,000 Rs.), [50,000 Rs., 100,000 Rs.), [100,000 Rs., 1,000,000 Rs.), [1,000,000 Rs., ∞). I code each of these as 10,000 Rs., 50,000 Rs., 100,000 Rs., 300,000 Rs., and 1,000,000 Rs. Alternative codings do not affect significance of any results nor magnitudes of any results in logs, but magnitudes in levels are sensitive to the coding of the [100,000 Rs., 1,000,000 Rs.) category. Estimates of the pseudo treatment effect elasticity of demand are unaffected.

[9] All land is cateogorized as unconsolidated formations (>40 L/s, 25-40 L/s, 10-25 L/s, <10 L/s), consolidated/semi-consolidated formations (1-25 L/s, 1-10 L/s, 1-5 L/s), and hilly areas (1 L/s), which I code as 40, 25, 10, 1, 25, 10, 1, and 1 L/s, respectively. This measure is strongly correlated with the measure of aquifer depth used by Sekhri (2014), and the measure of whether groundwater formations are unconsolidated or consolidated used by D'Agostino (2017).

crop yields under the intermediate irrigated scenario to potential rainfed crop yield.[10] I plot variation in log relative potential irrigated crop yield across districts in Panel (b) of Figure 1.1.[11] This measure is likely to be correlated with gross revenue under rainfed agriculture; I therefore control for log potential rainfed crop yield in all primary specifications. I discuss the construction of relative potential irrigated crop yield and potential rainfed crop yield in more detail in Appendix A.1.

I make use of some supplementary datasets. I use data from the Indian Central Groundwater Board's network of monitoring tubewells on seasonal depth to water table from 1995 to 2017; I refer to this dataset as "Well '95-'17". Data on the groundwater share of irrigation by district in 2001 is from the FAO Global Map of Irrigation Areas. Data sources of all calibrated parameters for counterfactual exercises in Section 1.5.4 and Section 1.7 are cited in Table 1.7.

1.3 Model

I consider a model of profit maximizing farmers deciding whether to irrigate their land. Following Suri (2011), I use a generalized Roy model to model the selection decision: although only farmers' gross revenue conditional on their adoption decision is observed, farmers decide to irrigate if their gross revenue under irrigation minus gross revenue under rainfed agriculture (gross returns to irrigation) is greater than their relative costs of irrigating. Past work has established nonparametric identification of parameters of these models from panel data (Suri, 2011), instruments for costs (Heckman & Vytlacil, 2005), instruments for treatment effects (Adão (2016); in this context, treatment effects are the gross returns to irrigation), and instruments for both costs and treatment effects (Das et al., 2003; Eisenhauer et al., 2015).

In Section 1.3.1, I consider a simple econometric model to motivate the more general framework. In Section 1.3.2, I setup a generalized Roy model building on the work cited above. I assume the presence of a conventional cost instrument, but I also impose a novel exclusion restriction on an outcome instrument: I assume the outcome instru-

[10]The weights used are state-by-state shares of land allocated to different crops. To identify effects from variation in potential crop yield, and not variation in weights, I control for state fixed effects in all analysis. Other work has used the difference in yields under different scenarios as an instrument for returns to technology adoption (Bustos et al., 2016).

[11]The measure is almost identical if I use the high input scenarios; in India, for almost all crops, potential yields under the high input scenario are closely approximated by a crop specific multiple of potential yields under the intermediate and low input scenarios. Regressing potential yields from the rainfed high input scenario on the rainfed intermediate input scenario yields R^2 ranging from 0.87 to 1, while regressing potential yields from the irrigated intermediate input scenario on the rainfed intermediate input scenario yields R^2 ranging from 0.04 and 0.06 on the low end (for water intensive sugarcane and rice) to 0.90 and 1 on the high end (for drought resilient sorghum and pearl millet).

ment does not affect gross revenue under rainfed agriculture (potential outcome under control). In Section 1.3.3, I define the marginal treatment effect (following Heckman & Vytlacil (2005)), and two novel parameters, the marginal surplus effect and the treatment effect elasticity of demand. The marginal surplus effect builds on Willig (1978) and Small & Rosen (1981): it is the inverse semielasticity of demand for irrigation, which equals the effect on profits caused by shifts to profitability of irrigation, as inferred by changes in adoption of irrigation. The treatment effect elasticity of demand captures the percentage increase in adoption of irrigation caused by a 1% increase in treatment on the treated (the effect of irrigation on gross revenue for inframarginal irrigators); it is inversely proportional to the marginal surplus effect and unitless, which facilitates interpretation and comparison across studies. In Section 1.3.4, I establish nonparametric identification of the marginal surplus effect. I show that the treatment effect elasticity of demand is not nonparametrically identified without strong assumptions on the instruments, but a pseudo treatment effect elasticity of demand, that serves as a reasonable approximation in many contexts, is. In Section 1.3.5, I discuss estimation of the marginal surplus effect. I show that linear instrumental variables using the outcome instrument estimates the sum of a local average treatment effect (a weighted average of marginal treatment effects) and a local average surplus effect (a weighted average of marginal surplus effects), and that these weights are nonparametrically identified. I compare the linear instrumental variables approach to a control function approach, and show that with the novel exclusion restriction the control function approach is overidentified.

1.3.1 A simplified econometric model

Consider the following econometric model

$$Y_i = \beta_0 + \beta_1 D_i + \beta_2 D_i W_i + \epsilon_i$$
$$D_i = \gamma_0 + \gamma_1 Z_i + \gamma_2 W_i + \eta_i$$

where Y_i is an observed outcome for agent i and D_i is the agent's endogenous adoption of a binary treatment. I make the independence assumption that $(Z_i, W_i) \perp (\epsilon_i, \eta_i)$. Z_i shifts agents decisions to adopt treatment. W_i shifts agents decisions to adopt treatment through its effect on treatment effects; $\beta_1 + \beta_2 w$ is the treatment effect for agents with $W_i = w$. The estimand of interest is $\frac{\gamma_2}{\beta_2}$, or the effect of a unit increase in treatment effects on adoption of treatment. An implicit exclusion restriction has been made here, that W_i does not affect outcomes for agents who do not adopt treatment.

I consider estimation of β_2 by linear instrumental variables, using Z_i and W_i as instruments for D_i and $D_i W_i$. This yields the following IV estimand for β_2, the effect of an increase in W_i on treatment effects.

$$\hat{\beta}_2 = \frac{\frac{\text{Cov}(Y_i, W_i)}{\text{Cov}(D_i, W_i)} - \frac{\text{Cov}(Y_i, Z_i)}{\text{Cov}(D_i, Z_i)}}{\frac{\text{Cov}(D_i W_i, W_i)}{\text{Cov}(D_i, W_i)} - \frac{\text{Cov}(D_i W_i, Z_i)}{\text{Cov}(D_i, Z_i)}}$$

This estimator is the ratio of two terms. The denominator is nonzero when there is a first stage for the IV estimator (W_i and Z_i are correlated with $D_i W_i$ relative to D_i differentially). The numerator is the difference between two linear IV estimators. The first of these estimators, but not the second, violates the exclusion restriction for instrumental variables in the more general correlated random coefficients model $Y_i = \beta_0 + \beta_{1i} D_i + \epsilon_i$.[12]

What is this difference between IV estimators under this model? For expositional purposes, I assume that Z_i and W_i are each binary, that they are independent, and that they are each 0 (1) with probability $\frac{1}{2}$ ($\frac{1}{2}$).[13]

$$\frac{\text{Cov}(Y_i, W_i)}{\text{Cov}(D_i, W_i)} - \frac{\text{Cov}(Y_i, Z_i)}{\text{Cov}(D_i, Z_i)} = \frac{\beta_2 \mathbf{E}[D_i]}{\gamma_2}$$

The difference between the two linear IV estimators is β_2, the change in treatment effects, times $\mathbf{E}[D_i]$, average adoption, divided by γ_2, the change in adoption. This is an inverse semielasticity of adoption to the treatment effect. The first IV estimator, $\frac{\text{Cov}(Y_i, W_i)}{\text{Cov}(D_i, W_i)}$, is the sum of two terms: $\beta_1 + \beta_2 \mathbf{E}[W_i]$, the local average treatment effect for agents induced to adopt treatment by W_i or Z_i, and an inverse semielasticity $\frac{\beta_2 \mathbf{E}[D_i]}{\gamma_2}$, the direct effect of W_i on outcomes per unit change in adoption of treatment.

The result is that the difference between two linear IV estimators, the first using an "instrument" for potential outcome under treatment, and the second using an instrument for treatment, estimates an inverse semielasticity of adoption of treatment to the treatment effect when the distribution of the "instrument" for potential outcome under treatment has no skew. However, it is not clear what this approach estimates when non-linearities or more flexible patterns of selection are permitted. With this motivation, I now ask if a similar approach can be used to estimate an inverse semielasticity of adoption in a generalized Roy model, where agents select into treatment on heterogeneous

[12] Note that this estimator I propose of β_2 is different from the natural estimator in the interacted model $Y_i = \beta_0 + \beta_1 D_i + \beta_2 D_i W_i + \beta_3 W_i + \epsilon_i$, using Z_i and $Z_i W_i$ as instruments for D_i and $D_i W_i$. When W_i is binary, one can show this $\hat{\beta}_2 = \frac{\text{Cov}(Y_i, Z_i | W_i = 1)}{\text{Cov}(D_i, Z_i | W_i = 1)} - \frac{\text{Cov}(Y_i, Z_i | W_i = 0)}{\text{Cov}(D_i, Z_i | W_i = 0)}$ (Hull, 2018). Under the more general econometric model presented in Section 1.3.2, even local versions of this estimator, and the one I propose under the exclusion restriction $\beta_3 = 0$, estimate different parameters. Loosely speaking, the estimator in the interacted model estimates the effect of W_i on the local average treatment effect, while the estimator in the model I propose estimates the effect of W_i on treatment on the treated.

[13] The latter two are without loss of generality as long as $(Z_i, W_i) = (z, w)$ with positive probability for all $(z, w) \in \{0, 1\}^2$, as this can be achieved by reweighting. A more general model that relaxes many of these assumptions is developed in Section **??**.

treatment effects and costs of adoption.[14]

1.3.2 Environment

Farmers ("agents") decide whether to adopt irrigation ("treatment") to maximize their profits ("surplus"), which is their gross revenue ("outcome") net of any costs, broadly defined. Let Y_{1i} be the gross revenue farmer i receives when they irrigate ("potential outcome under treatment"), and Y_{0i} be the gross revenue farmer i receives when they engage in rainfed agriculture ("potential outcome under control"). Let C_{1i} be farmer i's relative costs of adopting irrigation ("costs of adoption"). Let D_i be an indicator for farmer i's decision to irrigate ("treatment indicator"). Farmers maximize profits, $\pi_i = D_i(Y_{1i} - C_{1i}) + (1 - D_i)Y_{0i}$ ("surplus"). I assume the researcher observes $Y_i = D_iY_{1i} + (1 - D_i)Y_{0i}$, farmer i's gross revenue ("outcome"), and D_i, farmer i's decision to irrigate ("adoption decision"), but does not observe profits, costs, or counterfactual revenue.

The surplus maximization assumption implies

$$D_i = \mathbf{1}\{Y_{1i} - C_{1i} - Y_{0i} > 0\} \tag{1.1}$$

Equation 1.1 is equivalent to the generalized Roy modeling framework discussed in Heckman & Vytlacil (2007a,b). Agents adopt treatment if their treatment effect ($Y_{1i} - Y_{0i}$) is greater than their costs of adoption (C_{1i}).

Next, I assume the presence of instruments z and w. z is a conventional instrument, in that it shifts agents' costs of adoption, C_{1i}, without affecting their potential outcomes, Y_{1i} and Y_{0i}. I refer to it as the "cost instrument". However, w is a nonstandard instrument: it shifts agents' potential outcome under treatment, Y_{1i}, without shifting their costs of adoption, C_{1i}, or their potential outcome under control, Y_{0i}. I refer to it as the "outcome instrument". Additional assumptions are explained below.

Assumption 1.

$$\begin{aligned} Y_{1i}(w) &= V_{\gamma i}\gamma_W(w) + V_{1i} \\ C_{1i}(z) &= V_{\gamma i}\gamma_Z(z) + V_{Ci} \\ Y_{0i} &= V_{0i} \end{aligned}$$

[14] Wooldridge (2015) proposes control function approaches that allow for selection on unobservable treatment effect heterogeneity and that allow for multiple endogenous regressors. However, what linear estimators with multiple endogenous regressors estimate when the structural model is misspecified may not be useful (Kirkeboen et al., 2016; Hull, 2018; Mountjoy, 2018), while linear instrumental variables with a single endogenous regressor retains a LATE interpretation without any assumptions on functional forms (Heckman & Vytlacil, 2005). I ask if this robustness can be extended to linear instrumental variables with W_i.

Assumption 2. *γ_W and γ_Z are each monotonic in their arguments, and $V_{\gamma i} > 0$ $\forall i$. The distribution of $V_i \equiv \frac{-V_{1i}+V_{Ci}+V_{0i}}{V_{\gamma i}}$ is continuous and has a strictly increasing cumulative distribution function F_V and smooth density f_V.*

Assumption 1 implicitly makes a number of assumptions. First, w and z each satisfy exclusion restrictions. Only Y_{1i} is structurally a function of w, and only C_{1i} is structurally a function of z. These exclusion restrictions are strong assumptions, and I discuss possible violations in my empirical context in Section 1.6. That only Y_{1i} is structurally a function of w is a novel exclusion restriction in generalized Roy models.[15] It is most similar to Eisenhauer et al. (2015), who assume there is a regressor excluded from just C_{1i}, while I assume w is excluded from C_{1i} and Y_{0i}. That z is excluded from Y_{1i} and Y_{0i} is the standard exclusion restriction made to estimate a local average treatment effect.

Second, (z, w) are weakly separable from unobserved heterogeneity, through the index $(\gamma_W(w) - \gamma_Z(z))$. Combined with Assumption **??**, this implies monotonicity in an index of (z, w). It also implies the more general weak separability assumption made in Willig (1978), Small & Rosen (1981), and Bhattacharya (2017), who assume weak separability of price and product quality to estimate welfare impacts of changes to product quality on consumers. Crucially, this assumption guarantees that z and w enter choices and surplus symmetrically, so impacts on choices are strictly increasing in impacts on potential surplus under treatment. However, although weak separability only requires that (z, w) enter jointly through a flexible index, the more restrictive functional form I use is the most general that satisfies weak separability, the exclusion restrictions, the monotonicity assumptions, and the additive generalized Roy structure.[16] Despite the restrictiveness of these assumptions, variability in $V_{\gamma i}$ flexibly captures, for example, that more productive farmers might be more responsive to shifts in the instruments, something that similar work does not allow.[17]

Assumption 2 makes all remaining technical assumptions. The assumptions on monotonicity of γ_Z and γ_W are standard for instrumental variables, and reasonable in my context.[18] That the distribution of V_i is continuous and strictly increasing is a standard technical assumption.

[15] It is not novel in one-sided selection models, such as studying labor market participation, which two-sided models nest with a normalization of $Y_{0i} = 0$ (no earnings for non-participants). In these models, w is a wage shifter, D_i is the labor market participation decision, and z is an instrument for participation.

[16] The proof is in Appendix A.2.1.

[17] Specifically, Eisenhauer et al. (2015) and Adão (2016) require their instruments (z, w) are additively separable from unobserved heterogeneity, which implies that their instrument w has a homogeneous effect across agents conditional on observables. However, approaches in Das et al. (2003) and Eisenhauer et al. (2015) are straightforward to generalize to this environment.

[18] Specifically, in my context, I assume that potential revenue under irrigation is strictly increasing in potential irrigated crop yields, and that costs of irrigating are strictly decreasing in potential aquifer yield.

Additionally, define

$$U_i = F_V(V_i)$$

U_i is distributed Uniform[0,1], and orders agents from highest to lowest propensity to adopt treatment. Note that Equation 1.1, combined with Assumption **??** and the definition of V_i in Assumption 2, can now be rewritten as $D_i = \mathbf{1}\{U_i < F_V(\gamma_W(w) - \gamma_Z(z))\}$. Therefore, the share of agents who adopt treatment $\mathbf{E}[D_i(z,w)] = F_V(\gamma_W(w) - \gamma_Z(z))$.

Lastly, let Z_i and W_i be agent i's realized value of the instruments z and w. I make an independence assumption that will be sufficient for identification.

Assumption 3.

$$(Z_i, W_i) \perp (V_{0i}, V_{Ci}, V_{1i}, V_{\gamma i})$$

1.3.3 Marginal surplus effects and marginal treatment effects

Within this structure, it is now possible to define the marginal treatment effect and the marginal surplus effect.

$$\text{MTE}(u; w) = \mathbf{E}[Y_{1i}(w) - Y_{0i} | U_i = u] \tag{1.2}$$

$$\text{MSE}(u) = \frac{u}{f_V(F_V^{-1}(u))} \mathbf{E}[V_{\gamma i} | U_i < u] \tag{1.3}$$

The definition of the marginal treatment effect in Equation 1.2 is standard and follows Heckman & Vytlacil (2005). The definition of the marginal surplus effect in Equation 1.3 is novel. To interpret this, note that the ratio $\frac{u}{f_V(F_V^{-1}(u))}$ is just a Mills ratio for the random variable V_i, evaluated at $v = F_V^{-1}(u)$. The numerator, u, is the share of agents adopting treatment. The denominator, $f_V(F_V^{-1}(u))$, is the density of agents on the margin, which is similar to an elasticity: when the density of marginal agents is large, small increases in potential surplus under treatment cause large movements of agents into treatment. The third term reflects the extent to which inframarginal adopters of treatment are relatively more affected by shifts to z and w than compliers.

Following this intuition, we can arrive at a key result.

$$\frac{d\mathbf{E}[Y_i(z,w)]/dz}{d\mathbf{E}[D_i(z,w)]/dz} = \text{MTE}(\mathbf{E}[D_i(z,w)]; w) \tag{1.4}$$

$$\frac{d\mathbf{E}[\pi_i(z,w)]/dz}{d\mathbf{E}[D_i(z,w)]/dz} = \frac{d\mathbf{E}[\pi_i(z,w)]/dw}{d\mathbf{E}[D_i(z,w)]/dw} = \text{MSE}(\mathbf{E}[D_i(z,w)]) \tag{1.5}$$

Equation 1.4 gives the standard result on marginal treatment effects: the marginal treatment effect is the change in average outcomes per unit change in adoption of treatment caused by a shift to z. Equation 1.5 gives a new result on the marginal surplus effect: the marginal surplus effect is the change in average surplus per unit change in adoption of treatment caused by a shift to z or w.[19]

Additionally, following Heckman & Vytlacil (2007a,b), it follows from Equation 1.4 that one can define impacts on outcomes of policies that shift z in terms of MTE and $\mathbf{E}[D_i]$ alone. Similarly, it follows from Equation 1.5 that one can define impacts on surplus of policies that shift z or w in terms of MSE and $\mathbf{E}[D_i]$ alone.

$$\frac{\mathbf{E}[Y_i(z',w)] - \mathbf{E}[Y_i(z,w)]}{\mathbf{E}[D_i(z',w)] - \mathbf{E}[D_i(z,w)]} = \underbrace{\frac{\int_{\mathbf{E}[D_i(z,w)]}^{\mathbf{E}[D_i(z',w)]} \text{MTE}(u;w)du}{\mathbf{E}[D_i(z',w)] - \mathbf{E}[D_i(z,w)]}}_{\text{policy relevant treatment effect}} \tag{1.6}$$

$$\frac{\mathbf{E}[\pi_i(z',w')] - \mathbf{E}[\pi_i(z,w)]}{\mathbf{E}[D_i(z',w')] - \mathbf{E}[D_i(z,w)]} = \underbrace{\frac{\int_{\mathbf{E}[D_i(z,w)]}^{\mathbf{E}[D_i(z',w')]} \text{MSE}(u)du}{\mathbf{E}[D_i(z',w')] - \mathbf{E}[D_i(z,w)]}}_{\text{policy relevant surplus effect}} \tag{1.7}$$

Equation 1.6 is the standard result from Heckman & Vytlacil (2007a,b) that the impact of a broad class of policies on average outcomes is equal to the product of a policy relevant treatment effect and the impact of the policy on adoption of treatment, where the policy relevant treatment effect is a weighted average of marginal treatment effects. Equation 1.7 is a new result that shows that the impact of a broad class of policies on average surplus is equal to the product of a policy relevant surplus effect and the impact of the policy on adoption of treatment, where the policy relevant surplus effect is a weighted average of marginal surplus effects.

Lastly, to interpret Equation 1.5, it is helpful to draw a comparison to consumer theory. There, a classic result is that the marginal surplus effect is price divided by the price elasticity of demand (Willig, 1978; Small & Rosen, 1981). Alternatively, one could phrase this as the price elasticity of demand is equal to the price divided by the marginal surplus effect. An equivalent result holds here. I define

$$\text{TOT}(u;w) = \mathbf{E}[Y_{1i}(w) - Y_{0i}|U_i < u] \tag{1.8}$$

$$\epsilon^*(u;w) = \frac{\text{TOT}(u;w)}{\text{MSE}(u)} \tag{1.9}$$

Equation 1.8 gives the standard definition of treatment on the treated. Note that it has the standard interpretation, that $\text{TOT}(\mathbf{E}[D_i(z,w)];w) = \mathbf{E}[Y_{1i}(w) - Y_{0i}|D_i(z,w) = 1]$. Given the analogy in consumer theory, one might hope that $\epsilon^*(u;w)$, as defined in Equation 1.9, is the treatment effect elasticity of demand. Equation 1.10 shows this

[19] The derivations of Equation 1.4 and Equation 1.5 is in Appendix A.2.1.

result below.

$$\frac{\text{TOT}(\mathbf{E}[D_i(z,w)];w)}{\mathbf{E}[D_i(z,w)]}\frac{d\mathbf{E}[D_i(z,w)]/dw}{\partial\text{TOT}(\mathbf{E}[D_i(z,w)];w)/\partial w} = \epsilon^*(\mathbf{E}[D_i(z,w)];w) \tag{1.10}$$

Equation 1.10, combined with Equation 1.9, shows that the marginal surplus effect can be interpreted as the ratio of treatment on the treated to the treatment effect elasticity of demand for treatment.[20]

1.3.4 Identification

The identification of marginal surplus effects and marginal treatment effects follows from classic results on local instrumental variables from Heckman & Vytlacil (1999, 2005). I now assume that (Z_i, W_i) have a smooth density that is strictly positive at (z, w). Independence of the instruments and standard results on nonparametric identification imply the expectations $\mathbf{E}[Y_i(z,w)]$ and $\mathbf{E}[D_i(z,w)]$ and their derivatives with respect to z and w are identified (Matzkin, 2007).[21] As in Heckman & Vytlacil (2005), Equation 1.4 therefore establishes identification of marginal treatment effects from local instrumental variables using the cost instrument.

For identification of marginal surplus effects, the key result is what local instrumental variables using the outcome instrument estimates.

$$\frac{d\mathbf{E}[Y_i(z,w)]/dw}{d\mathbf{E}[D_i(z,w)]/dw} = \text{MTE}(\mathbf{E}[D_i(z,w)];w) + \text{MSE}(\mathbf{E}[D_i(z,w)]) \tag{1.11}$$

Local instrumental variables using the outcome instrument estimates the marginal treatment effect plus the marginal surplus effect.[22] This is the local version of the result for the linear model in Section 1.3.1.

Identification of marginal surplus effects follows simply from subtracting Equation 1.4 from Equation 1.11.

$$\text{MSE}(\mathbf{E}[D_i(z,w)]) = \frac{d\mathbf{E}[Y_i(z,w)]/dw}{d\mathbf{E}[D_i(z,w)]/dw} - \frac{d\mathbf{E}[Y_i(z,w)]/dz}{d\mathbf{E}[D_i(z,w)]/dz} \tag{1.12}$$

The intuition for this result is visible in Figure 1.2. Both the cost instrument z and the outcome instrument w affect agent adoption decisions and surplus through a common index, because of the weak separability assumption. Whether surplus under treatment

[20] The derivation of Equation 1.10 is in Appendix A.2.1.

[21] Formally, $\mathbf{E}[Y_i(z,w)] = \mathbf{E}[Y_i|Z_i = z, W_i = w]$ and $\mathbf{E}[D_i(z,w)] = \mathbf{E}[D_i|Z_i = z, W_i = w]$.

[22] The derivation of Equation 1.11 is in Appendix A.2.1.

increases from $Y_{1i} - C_{1i}$ to $Y_{1i}^{*} - C_{1i}$ (shock to w, as in Panel (a)) or to $Y_{1i} - C_{1i}^{*}$ (shock to z, as in Panel (b)), the effect on choices is a sufficient statistic for the effect on surplus; the marginal surplus effect is well defined. However, their effects on outcomes differ. In Panel (b), we can see that the cost instrument increases outcomes proportional to the marginal treatment effect: potential outcomes are unaffected by the cost instrument, but the induced increase in adoption $\mathbf{E}[D_i]$ causes agents' outcomes to increase by their treatment effect. However, in Panel (a), we can see that the outcome instrument has two effects on outcomes. The first effect is proportional to the marginal treatment effect: adoption $\mathbf{E}[D_i]$ increases because surplus under treatment increases, and this increase in adoption $\mathbf{E}[D_i]$ causes agents' outcomes to increase by their treatment effect. However, the second effect is proportional to the marginal surplus effect. This is the direct effect on outcomes caused by the increase in Y_{1i}; the increase in Y_{1i} and the increase in $Y_{1i} - C_{1i}$ are the same (because of the exclusion restriction), so this increase is exactly the same as the effect of the outcome instrument on surplus.

Note, however, that unlike marginal surplus effects and marginal treatment effects, treatment on the treated and the treatment effect elasticity of demand are not identified without either parametric assumptions or an identification at infinity argument. This contrasts with the standard consumer theory setting, where typically a price elasticity of demand is estimated, and marginal surplus effects can be calculated using that price elasticity. To allow comparison of results with price elasticities, I instead define the pseudo treatment effect elasticity of demand to be

$$\epsilon(u;w) = \frac{\text{MTE}(u;w)}{\text{MSE}(u)} \tag{1.13}$$

which, following the results above, is also identified. It is biased relative to the treatment effect elasticity of demand: $\epsilon^{*}(u;w) = \frac{\text{TOT}(u;w)}{\text{MTE}(u;w)}\epsilon(u;w)$, so the pseudo treatment effect elasticity of demand, which requires less restrictive assumptions for identification, will be too small (large) when treatment on the treated is large (small) relative to the marginal treatment effect.[23]

[23]Despite this, the pseudo treatment effect elasticity of demand is still useful. In some cases, instead of observing the outcome Y_i, the researcher might observe the outcome Y_i times an unknown constant (in agriculture, this could be yields measured using satellite data, as in Burke & Lobell (2017)) or costs $D_i C_{1i}$ times an unknown constant (in my context, this is fixed infrastructure costs for irrigation). In both cases, the pseudo treatment effect elasticity of demand can still be consistently estimated. In my context, this permits an overidentification test. In other cases, one might estimate the pseudo treatment effect elasticity of demand in one context, and extrapolate to another where the marginal treatment effect is known but an instrument to estimate the marginal surplus effect is unobserved.

1.3.5 Estimation

For estimation, I now assume that a set of observable characteristics of each agent, X_i, are also observed. All assumptions above are now made conditional on $X_i = x$, and all results above now hold conditional on $X_i = x$. No additional assumptions are made except where explicitly stated.

Instrumental variables

The nonparametric identification results suggest the application of local instrumental variable estimators. In practice, as discussed in Carneiro et al. (2011) and Eisenhauer et al. (2015), local instrumental variable estimators are difficult to implement in practice while conditioning on (Z_i, W_i, X_i) jointly. Frequently, their implementation relies on strong restrictions on how (W_i, X_i) can enter outcome equations. However, as Imbens & Angrist (1994) and Heckman & Vytlacil (2005) show, linear instrumental variables using a conventional instrument, such as Z_i, makes no such assumptions: instead, it only requires the researcher to estimate the expectation of Z_i conditional on all variables which are not excluded from outcome equations (in this case, (W_i, X_i)). Then, linear instrumental variables estimates a local average treatment effect, or a weighted average of marginal treatment effects. Flexibly controlling for observables in linear instrumental variables is well understood (for example, see Chernozhukov et al. (2016)), and does not require any assumptions on how non-excluded observables enter outcome equations, in contrast to how local instrumental variable methods are often implemented (Carneiro et al., 2011).

Just as linear instrumental variables with Z_i estimates a local average treatment effect, linear instrumental variables with W_i estimates the sum of a local average treatment effect and a local average surplus effect, where a local average surplus effect is a weighted average of marginal surplus effects. Formally,

$$\beta_Z^{IV} \equiv \frac{\text{Cov}(Y_i, Z_i - \mathbf{E}[Z_i|W_i, X_i])}{\text{Cov}(D_i, Z_i - \mathbf{E}[Z_i|W_i, X_i])} = \text{LATE}_Z \tag{1.14}$$

$$\text{LATE}_Z = \int \text{MTE}(u; w, x)\omega_Z(u; w, x)dudwdx \tag{1.15}$$

$$\beta_W^{IV} \equiv \frac{\text{Cov}(Y_i, W_i - \mathbf{E}[W_i|Z_i, X_i])}{\text{Cov}(D_i, W_i - \mathbf{E}[W_i|Z_i, X_i])} = \text{LATE}_W + \text{LASE}_W \tag{1.16}$$

$$\text{LATE}_W = \int \text{MTE}(u; w, x)\omega_W(u; w, x)dudwdx \tag{1.17}$$

$$\text{LASE}_W = \int \text{MSE}(u; x)\omega_W(u; w, x)dudwdx \tag{1.18}$$

Equation 1.14 and Equation 1.15 are the result from Heckman & Vytlacil (2005): linear instrumental variables using the cost instrument estimates a local average treatment effect, which is a weighted average of marginal treatment effects. As Heckman & Vytlacil (2005) show, these weights ω_Z are nonparametrically identified, positive, and integrate to 1. The new result is Equation 1.16: linear instrumental variables using the outcome instrument estimates a local average treatment effect plus a local average surplus effect. The local average surplus effect is a weighted average of marginal surplus effects. I show in Appendix A.2.2 that the LATE_W and LASE_W weights, ω_W, are nonparametrically identified, positive, and integrate to 1. This extends the result on the linear model from Section 1.3.1 to a generalized Roy model with nonlinearities and selection on heterogeneous treatment effects.

InterpoLATE-ing

There are multiple approaches in the literature to estimation of LATE_W. First, nonparametric bounds on LATE_W using LATE_Z are derived in Mogstad et al. (2017), by considering the largest and smallest possible values of LATE_W consistent with marginal treatment effects that would result in estimating LATE_Z. Second, if variation in treatment effects is explained by observables, Angrist & Fernandez-Val (2010) show weighted linear instrumental variables with the cost instrument can estimate LATE_W. Third, one could instead estimate marginal treatment effects directly using the cost instrument, and recover an estimate of LATE_W from the marginal treatment effects and an estimate of the LATE_W weights. Alternatively, Brinch et al. (2017) propose an approach to recovering marginal treatment effects from estimates of local average treatment effects, by imposing restrictions on outcome equations and flexibly modeling the distribution of unobservable heterogeneity.

I build on Angrist & Fernandez-Val (2010), and assume that variation in local average treatment effects is explained by observables. Specifically, I partition $X_i = (\tilde{X}_i, S_i)$, and assume that local average treatment effects conditional on S_i are homogeneous estimated using W_i or Z_i. Formally, define

$$\text{LATE}_{(\cdot)|s} = \frac{\int \text{MTE}(u; w, (\tilde{x}, s))\omega_{(\cdot)}(u; w, (\tilde{x}, s))du dw d\tilde{x}}{\int \omega_{(\cdot)}(u; w, (\tilde{x}, s))du dw d\tilde{x}}$$

to be the conditional local average treatment effect.[24] I assume

Assumption 5a. $LATE_{Z|s} = LATE_{W|s}\ \forall s \in Supp(S_i)$

Although this is a strong assumption, I show in Appendix A.2.1 that this still poten-

[24] In my empirical context, S_i are a vector of state dummies; across state geographic heterogeneity and policies are likely to explain a significant share of treatment effect heterogeneity.

tially allows for arbitrary linear marginal treatment effects. This allows for "essential heterogeneity" (Heckman et al., 2006), substantially weakening the assumption made in Angrist & Fernandez-Val (2010), who assume these conditional local average treatment effects are also equal to a conditional average treatment effect. The difference is while their goal is to estimate the average treatment effect and other population moments, my goal is to estimate LATE_W, which requires a much weaker assumption.[25]

From the definition of the conditional local average treatment effect, it is clear that LATE_W is a weighted average of $\text{LATE}_{W|s}$, and therefore $\text{LATE}_{Z|s}$. It therefore follows that LATE_W can be estimated by weighted linear instrumental variables using Z_i. Letting $\overline{\omega}_{(\cdot)}(s) \equiv \int \omega_{(\cdot)}(u; w, (\tilde{x}, s)) du dw d\tilde{x}$, this yields the following estimator of LASE_W.[26]

$$\beta_W^{IV} - \beta_Z^{WIV} = \text{LASE}_W \tag{1.19}$$

$$\beta_Z^{WIV} \equiv \frac{\text{Cov}\left((\overline{\omega}_W(S_i)/\overline{\omega}_Z(S_i))Y_i, Z_i - \mathbf{E}[Z_i|W_i, X_i]\right)}{\text{Cov}\left((\overline{\omega}_W(S_i)/\overline{\omega}_Z(S_i))D_i, Z_i - \mathbf{E}[Z_i|W_i, X_i]\right)} \tag{1.20}$$

The difference between weighted linear instrumental variable estimators is a consistent estimator of LASE_W. Intuitively, the weights make the z compliers resemble the w compliers on the observable S_i.

Additionally, the ratio of the local average treatment effect to the local average surplus effect estimated using weighted instrumental variables estimates a weighted average of pseudo treatment effect elasticities of demand.

$$\frac{\beta_Z^{WIV}}{\beta_W^{IV} - \beta_Z^{WIV}} = \int \epsilon(u; w, x) \left(\frac{\omega_W(u; w, x)\text{MSE}(u; x)}{\int \omega_W(u; w, x)\text{MSE}(u; x) du dw dx} \right) du dw dx \tag{1.21}$$

This result follows straightforwardly from $\beta_Z^{WIV} = \text{LATE}_W$, and substituting the definition $\epsilon(u; w, x) = \frac{\text{MTE}(u;w,x)}{\text{MSE}(u;x)}$. The weights $\frac{\omega_W(u;w,x)\text{MSE}(u;x)}{\int \omega_W(u;w,x)\text{MSE}(u;x)du dw dx}$ are nonparametrically identified, positive, and integrate to 1.

This estimator of a local average surplus effect may be underpowered, if there are many w compliers but very few z compliers for some S_i, but there is balance for other S_i. In Appendix A.2.2, I propose feasible reweighted instrumental variable estimators using both z and w to minimize the variance of the resulting estimator of a local average surplus effect; I refer to these estimators as β_W^{WIV} and β_Z^{WIV}. Additionally, estimating $\overline{\omega}_{(\cdot)}$, even under Assumption 5a, requires estimating the effect of w and z on adoption

[25] Note that this is still much stronger than assumptions made by Brinch et al. (2017) and Mogstad et al. (2017). However, the estimator I propose is much simpler to implement. Additionally, in Section 1.5.2, I estimate a parametric version of the model from Section 1.3.2 that does not impose this assumption, and estimates of this model suggest bias from violations of this assumption is small in my context.

[26] The proof of Equation 1.19 is in Appendix A.2.1.

conditional on $S_i = s$, something I am underpowered for in my setting. Given this constraint, I calculate these weights under the assumption that the first stages for w and z (the derivatives of the propensity score conditional on $S_i = s$ with respect to w and z) are constant across $S_i = s$. However, the estimator is still consistent (although no longer efficient) if the first stage for w is a constant multiple of the first stage for z across $S_i = s$.

ExtrapoLASE-ing

Just as with a local average treatment effect, a single estimate of a local average surplus effect need not be policy relevant. I propose an approach similar to Brinch et al. (2017), who use estimates of outcomes for always takers, compliers, and never takers to recover the marginal treatment effect with a discrete instrument under parametric assumptions. Instead, I recover the marginal surplus effect from estimates of local average surplus effects. Recall that the local average surplus effect is a weighted average of marginal surplus effects, and the weights ω_W are identified. Furthermore, recall that $\text{MSE}(u;x) = \frac{u}{f_V(F_V^{-1}(u;x);x)}\mathbf{E}[V_{\gamma i}|U_i < u, X_i = x]$. Given this, with parametric restrictions on $\text{MSE}(u;x)$, implied by restrictions on the joint distribution of (V_{γ_i}, V_i) conditional on $X_i = x$, one can identify $\text{MSE}(u;x)$ from local average surplus effects and the weights they place on different marginal surplus effects.

In particular, I assume the marginal surplus effect is linear. Unlike a marginal treatment effect, for many distributions a marginal surplus effect will have a 0 intercept, and therefore a single parameter (the slope) is sufficient to characterize a linear marginal surplus effect.[27,28] A linear marginal surplus effect is therefore identified from a single estimate of a local average surplus effect, and the weights ω_W. Formally, I assume

Assumption 5b. $MSE(u) = ku$

Note that this assumption is neither necessary nor sufficient for linear marginal treatment effects conditional on $X_i = x$, and allows for flexible nonlinearities in the effects of the cost and outcome instruments on costs and potential outcome under treatment, respectively, conditional on $X_i = x$. Under this assumption, estimation of the marginal surplus effect from an estimate of the local average surplus effect is straightforward.

$$k = \frac{\text{LASE}_W}{\int u\omega_W(u;w,x)dudwdx} \tag{1.22}$$

[27] Specifically, bounded $V_{\gamma i}$ and the distribution of V_i not having fat tails are sufficient for the marginal surplus effect to have a 0 intercept; this is a standard property of a Mills ratio.

[28] One parametrization that yields a linear marginal surplus effect is $V_i \sim \text{Uniform}[a, a+k]|X_i = x$, and $V_{\gamma i} = 1\ \forall i$.

In general, estimation of $\omega_W(u; w, x)$ can be hard, even though it is nonparametrically identified. I simplify the problem by estimating $\omega_W(u; w, x)$ under the assumption that $\mathbf{E}[D_i(z, w; x)]$ is linear.

Parametric control function ("Heckit")

Past work has developed control function approaches that could be used to estimate a marginal surplus effect, including parametric (Heckman, 1979), semiparametric (Ahn & Powell, 1993), and nonparametric approaches (Das et al., 2003). In fact, the natural estimator of the marginal surplus effect building on the estimator of Das et al. (2003) is asymptotically equivalent to a local instrumental variables estimator suggested by Equation 1.12. However, the control function estimator is overidentified; this is because it requires observations of $\mathbf{E}[Y_i|D_i, W_i, Z_i, X_i]$ and $\mathbf{E}[D_i|W_i, Z_i, X_i]$, while the instrumental variable approach I propose only requires observations of $\mathbf{E}[Y_i|W_i, Z_i, X_i]$ and $\mathbf{E}[D_i|W_i, Z_i, X_i]$. Specifically, the exclusion restriction that Y_{0i} is not a function of w is more easily testable with more disaggregated data.

As an alternative to the instrumental variable approach to estimating a marginal surplus effect presented previously, I consider a two step parametric control function approach using a standard Heckman selection correction. As in Björklund & Moffitt (1987), I assume idiosyncratic variation in (Y_{1i}, C_{1i}, Y_{0i}) is jointly normally distributed. Although the normality assumption appears restrictive, Kline & Walters (2017) show that in many cases, parametric control function approaches exactly or closely match the same moments as linear IV estimators, and thus produce identical or similar estimates of local average treatment effects.

Assumption 5c.

$$\begin{pmatrix} Y_{1i} \\ C_{1i} \\ Y_{0i} \end{pmatrix} \sim N\left(\begin{pmatrix} (g_W + c_0)W_i + X_i'\mu_1 \\ g_Z Z_i + X_i'\mu_C \\ c_0 W_i + X_i'\mu_0 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{1c} & \Sigma_{10} \\ \Sigma_{1c} & \Sigma_{cc} & \Sigma_{c0} \\ \Sigma_{10} & \Sigma_{c0} & \Sigma_{cc} \end{pmatrix}\right)$$

Details of the estimation are in Appendix Section A.2.3. From the estimated model, it is straightforward to calculate the marginal surplus effect; this calculation under normality is similar to the expression for the treatment effect elasticity of demand under normality in French & Taber (2011).

$$\text{MSE}(u; x) = \frac{\sigma_V u}{\phi(\Phi^{-1}(u))} \tag{1.23}$$

where $\sigma_V = \text{Var}(V_i)$, ϕ is the normal density function, and Φ is the normal cumulative distribution function.

This parametric control function approach is a useful benchmark for the instrumental variable approach I propose. It also allows enables two additional tests of the instrumental variable approach. First, it allows me to test the exclusion restriction that Y_{0i} is not a function of w. Second, it allows me to test the performance of the weighted instrumental variable estimator. Specifically, I follow Andrews et al. (2018) and calculate the informativeness of the weighted (and unweighted) instrumental variable estimators of LASE_W and LATE_Z for the control function estimators of LASE_W and LATE_Z, respectively.

1.4 Empirical strategy

1.4.1 Notation and context specific concerns

Following Section 1.3.5 and the end of Section 1.3.5, but adapting to my empirical context, I consider observations of $(Y_{ins}, D_{ins}, Z_{ns}, W_{ns}, (X_{ns}, S_s))$ for each plot i, located in district n in state s. Y_{ins} is plot i's realized gross revenue. D_{ins} is an indicator for whether plot i is irrigated. Z_{ns} is plot i's value of the cost instrument, its potential aquifer yield. W_{ns} is plot i's value of the outcome instrument, its log relative potential irrigated crop yield. X_{ns} is a vector of controls for plot i, which in my main specifications is log potential rainfed crop yield. S_s is a vector of state dummies.

The instruments, (Z_{ns}, W_{ns}), and controls (X_{ns}, S_s), are constant within district. All analysis reports robust standard errors clustered at the district level.

In regressions using district level data, I observe area weighted average outcomes for the district. I use Y_{ns} for average gross revenue per hectare, and D_{ns} for share of land irrigated at the district level. That Y_{ns} and D_{ns} might vary across districts with the same values of the instruments, even though we can treat Y_{ns} and D_{ns} as population averages within district, is consistent with the distribution of unobservables varying across districts. The independence assumption therefore implies that instruments are assigned across districts independent of this distribution.

In analysis using data from NSS '12, I observe plot level data.[29] Y_{ins} is now gross revenue per hectare for plot i, and D_{ins} is a dummy for irrigated. The sampling in the Agricultural NSS was stratified on village level irrigation status, which is endogenous; as a result, I use survey weights to recover unbiased estimates. To maintain comparability with regressions using district level data, I also weight by plot size, and normalize

[29] To be more precise, observations are at the level of household-by-crop-by-irrigation adoption, which one can think of as aggregated across plots, proportional to area, on which households grow the same crop and make the same irrigation adoption decision.

weights such that the sum of weights in each district is 1.

In analysis using Irr '07, I use the negative of average fixed costs of irrigation infrastructure per agricultural hectare as an outcome. This provides a useful check on results from other datasets, as I discuss in Section 1.4.2.

1.4.2 Instrumental variables

My objective is to construct 2SLS estimators of the form in Equation 1.14 and 1.16. With a large number of clusters, one could estimate the conditional expectations of Z_{ns} and W_{ns} nonparametrically. With the 222 districts I observe, I instead take a parametric approach and assume $\mathbf{E}[Z_{ns}|W_{ns}, X_{ns}, S_s]$ and $\mathbf{E}[W_{ns}|Z_{ns}, X_{ns}, S_s]$ are linear conditional on S_s. With this, I estimate by OLS

$$Y_{ins} = \beta_Z^{RF} Z_{ns} + \delta_{1s} W_{ns} + \delta_{2s} X_{ns} + \alpha_{1s} + \epsilon_{1,ins} \tag{1.24}$$

$$D_{ins} = \beta_Z^{FS} Z_{ns} + \delta_{3s} W_{ns} + \delta_{4s} X_{ns} + \alpha_{2s} + \epsilon_{2,ins} \tag{1.25}$$

$$Y_{ins} = \beta_W^{RF} W_{ns} + \delta_{5s} Z_{ns} + \delta_{6s} X_{ns} + \alpha_{3s} + \epsilon_{3,ins} \tag{1.26}$$

$$D_{ins} = \beta_W^{FS} W_{ns} + \delta_{7s} Z_{ns} + \delta_{8s} X_{ns} + \alpha_{4s} + \epsilon_{4,ins} \tag{1.27}$$

Note that coefficients on controls are allowed to vary by state s in all specifications. Let $\beta_Z^{IV} = \beta_Z^{RF}/\beta_Z^{FS}$, and $\beta_W^{IV} = \beta_W^{RF}/\beta_W^{FS}$. I use $\beta_W^{IV} - \beta_Z^{IV}$ as an estimate of a local average surplus effect, and $\beta_Z^{IV}/(\beta_W^{IV} - \beta_Z^{IV})$ as an estimate of a pseudo treatment effect elasticity of demand.

These estimators may be inconsistent if $\text{LATE}_W \neq \text{LATE}_Z$. I therefore also implement the weighted instrumental variable estimator constructed in 1.3.5; this estimator will be consistent for a local average surplus effect and a pseudo treatment effect elasticity of demand under Assumption 5a.

To validate the approach, I also use the negative of average fixed costs of irrigation infrastructure per agricultural hectare as an outcome. This is consistent with the modeling framework; as Björklund & Moffitt (1987) and Eisenhauer et al. (2015) note, there is a duality between costs and benefits in the generalized Roy model; the difference is only which is treated as observable. To expand briefly, we are using $-qD_iC_{1i}$ as the outcome instead of Y_i, and $Y_{1i} - (1-q)C_{1i} - Y_{0i}$ as costs instead of C_{1i}, where q is the share of fixed costs in costs of irrigation times the discount rate (to convert infrastructure costs, which is a stock, into a flow); I assume q is constant. Instruments are now switched: W_i becomes the cost instrument, and Z_i becomes the outcome instrument. Estimated marginal treatment effects are $-q$ times marginal treatment effects, since $Y_{1i} - Y_{0i} = C_{1i}$ for marginal agents. Estimated marginal surplus effects are q times marginal surplus effects, since responses to increased surplus from decreased costs of irrigation and increased surplus from increased gross revenue under irrigation are the same. Therefore, the estimated pseudo treatment effect elasticity of demand (the ratio of the local average treatment effect to the local average surplus effect) when using negative fixed costs as an outcome should be the negative of the estimate using gross revenue as an outcome.[30]

1.4.3 Control function

To estimate the control function approach, I use NSS '12, in which I observe plot level data. This is crucial because this approach relies on observing average outcomes conditional both on the values of the instruments and on adoption of treatment, something the instrumental variables approach does not need. To separate differences in results coming from different methods and different data sets, I first estimate a local average surplus effect using linear instrumental variables in NSS '12. I follow Section 1.3.5 in estimating the control function approach. Controls include state fixed effects and their interaction with log potential rainfed crop yield, but the cost instrument z and outcome instrument w are not interacted with state fixed effects. Additional details of the approach are in Appendix A.2.3.

1.5 Results

1.5.1 Instrumental variables

Table 1.2 presents unweighted instrumental variable regressions in Ag '07-'11. Columns 1 and 2 show a strong first stage with the cost instrument and the outcome instrument, with t-statistics of 5.0 and 4.2, respectively. The instrumental variable coefficient in Column 6, which uses the cost instrument, is a local average treatment effect. Marginal irrigators increase their agricultural revenue by 22,600 Rs/ha when they adopt irrigation. For ease of interpretation, the same specification with log revenue per hectare as the outcome gives a coefficient of 0.95. This is similar to Duflo & Pande (2007), who estimate an elasticity of production with respect to dam induced irrigation of 0.61,

[30] Note that imposing $C_{0i} = 0$ is no longer a normalization in order for these interpretations of results using fixed costs as an outcome to be valid. This creates two problems. First, it creates the potential for exclusion restriction violations due to Z_i affecting costs of rainfed agriculture. This is not a concern in my context, since Z_i affects the costs of extracting groundwater. Second, it affects the interpretation of q. Assumptions that would imply q is constant are very strong, and likely require all costs of irrigation to involve drilling for and pumping groundwater, ruling out irrigation reducing growing season labor costs for rice cultivation, for example. I therefore interpret these results as suggestive robustness.

which they note is in the lower range of existing estimates. The instrumental variable coefficient in Column 7, which uses the outcome instrument, is the sum of a local average treatment effect and a local average surplus effect.

Table 1.3 presents instrumental variable and weighted instrumental variable estimates used to recover a local average surplus effect and pseudo treatment effect elasticity of demand; for compactness, each cell corresponds to a single regression. Columns correspond to a single set of estimates, while rows correspond to estimators. Column 1 presents the same results as are in Table 1.2. Row 5 of Column 1 is the difference between the IV estimator using the outcome instrument and the IV estimator using the cost instrument, which estimates a local average surplus effect if the two local average treatment effects (for cost instrument compliers and outcome instrument compliers) are the same. The estimated local average surplus effect is 31,700 Rs/ha. To facilitate interpretation, an estimate of the pseudo treatment effect elasticity of demand is presented in row 6: the resulting point estimate is 0.72, although it is imprecisely estimated.

Column 2 presents results with the weighted instrumental variable estimator, which corrects for potential bias from differences in shares of cost instrument and outcome instrument compliers in different states. The estimated local average surplus effect with this estimator, 49,800 Rs/ha, is larger (although not statistically significantly so), and the estimated pseudo treatment effect elasticity of demand is similar.

Columns 3 and 4 present results with negative infrastructure costs as the outcome using unweighted and weighted instrumental variables, respectively; as described in Section 1.4.2, the roles of the instruments are now switched. The local average treatment effect estimates imply marginal irrigation infrastructure costs of 59,100-86,900 Rs/ha. Unlike estimates with agricultural productivity as an outcome, these instrumental variable estimates are economically significantly different from OLS estimates, consistent with unobservable heterogeneity in costs of irrigation driving selection.[31] Although interpreting the local average surplus effect estimates is difficult, following the reasoning in Section 1.4.2, pseudo treatment effect elasticity of demand estimates should be the negative of estimates using agricultural productivity as an outcome. Estimates of this elasticity using infrastructure costs are statistically and economically indistinguishable from estimates using agricultural productivity, but are much more precisely estimated. The estimates imply a 1% increase in the gross returns to irrigation causes a 0.7% increase in adoption of irrigation, times a bias term equal to the ratio of gross returns for average irrigators to gross returns for marginal irrigators.

[31] The difference is not statistically significant (for the Hausman test, $p = 0.12$ for unweighted IV and $p = 0.13$ for weighted IV), so I interpret this difference as potentially suggestive of selection on unobservable heterogeneity in costs of irrigation.

1.5.2 Control function

Before estimating key model parameters using a two step control function approach, I first compare instrumental variable estimates of the local average surplus effect in NSS '12, on which the control function approach is implemented, to the estimates from Ag '07-'11. The estimate of the local average surplus effect in Column 1 on Table 1.4 is similar, but noisier; I interpret this to mean direct comparisons of control function estimates using NSS '12 to instrumental variable estimates using Ag '07-'11 are reasonable, although they should still be made with caution.

I present the estimated coefficients from the control function approach in Table 1.5. A few things to note. First, the estimated effect of the outcome instrument on potential revenue under rainfed agriculture, c_0, is not significantly different from 0, so the overidentification test fails to reject. Second, the estimated standard deviation of idiosyncratic profitability of irrigation of 25,800 Rs/ha, σ_V, is large: as reference, the observed standard deviation of agricultural revenue per hectare is 26,100 Rs/ha, although these two measures need not be similar. Third, the selection terms are imprecisely estimated, although there is potentially suggestive evidence that there is selection on costs, consistent with the differences between instrumental variables and OLS estimators with fixed costs as the outcome in Section 1.5.1.

To compare the control function approach to the instrumental variable approach, Column 3 of Table 1.4 shows estimates of LATE_Z, LATE_W, and LASE_W from the control function approach, along with bias from violations of the exclusion restriction.[32] The local average surplus effect, 54,300 Rs/ha, is larger than estimates from either instrumental variable method and is more precisely estimated. The estimated bias from differences between local average treatment effects is small, at -3,300 Rs/ha. The estimated bias from violations of the exclusion restriction is also small, at 4,600 Rs/ha. These biases happen to offset, and the total bias in the instrumental variable estimator of the local average surplus effect is just 1,200 Rs/ha.

However, just because the control function estimates imply the linear IV estimator has a small bias in this case does not mean it is a good estimator of a local average surplus effect. To judge this, I follow Andrews et al. (2018) and calculate the informativeness of the IV and WIV estimators of LATE_Z and LASE_W for the equivalent control function estimates. This does not capture bias, which is small in this context but need not be in others, but does capture the extent to which structural estimates of LATE_Z and LASE_W are explained by IV estimators. Kline & Walters (2017) note that in many cases, IV and structural estimates of LATE_Z are numerically equivalent, which would yield an informativeness of 1; I therefore use the informativeness of IV estimates of LATE_Z for structural estimates of LATE_Z as a benchmark. Table **??** shows these

[32] I discuss the construction of these in Section A.2.3.

measures. The IV and weighted IV estimators of LATE_Z both have high informativeness of structural estimates (0.51 and 0.46, respectively). The IV estimator of LASE_W has a low informativeness of the structural estimator (0.12). However, the WIV estimator of LASE_W has an informativeness of the structural estimate that is similar to that of IV estimates of LATE_Z for structural estimates of LATE_Z (0.50). I interpret this as evidence that the instrumental variable approach is, at the least, a useful complement to traditional structural approaches one could use to estimate marginal surplus effects, as the two approaches should yield similar results.

1.5.3 MSE

Estimated marginal surplus effects and local average surplus effects for the instrumental variable estimator (in Ag '07-'11), the weighted instrumental variable estimator (in Ag '07-'11), and the control function estimator (in NSS '12) are presented in Figure 1.3. The instrumental variable estimates of marginal surplus effects are constructed from the local average surplus effect estimates as described in Section 1.3.5. The control function estimate of the local average surplus effect is constructed from the marginal surplus effect estimate as described in Section 1.3.5. First, note that although the weighted IV local average surplus effect is 57% larger than the IV estimate, the weighted IV marginal surplus effect is only 30% larger. This is because the weighted IV local average surplus effect places more weight on larger margins of adoption, where marginal surplus effects will typically be larger (and are by assumption with the functional forms I use). Second, the control function estimate of the marginal surplus effect is larger than the IV estimate, but it is close to the WIV estimate over empirically relevant margins of adoption. As a result, for counterfactual exercises, I pick the "median" of the three estimates and use the WIV estimate of the marginal surplus effect. Third, note that distributional assumptions can have a large impact on estimates of the marginal surplus effect when extrapolating outside of frequently observed margins of adoption.

1.5.4 Groundwater depletion and rural surplus

With an estimate of the marginal surplus effect, we can calculate the effects of declining water tables on surplus. To do so, with the marginal surplus effect it is sufficient to have an estimate of the impact of declining water tables on adoption of irrigation. Let b be the depth to water table in meters. I calibrate $d\mathbf{E}[D_i]/db = -.0024/\text{m}$ based on estimates from Fishman et al. (2017), which I assume to be constant.[33] This yields

$$\frac{d\mathbf{E}[\pi_i]}{db} = \text{MSE}(\mathbf{E}[D_i])\frac{d\mathbf{E}[D_i]}{db}$$

[33] This, and all other calibrated parameters used in counterfactual exercises, are in Table 1.7.

I use this approach to calculate the impact of declining water tables on economic surplus, and report estimates in Table 1.8. Column 1 reports estimates of the impact of a 1m decline in water tables on economic surplus in Rs/ha. The WIV marginal surplus effect implies a 1m decline in water tables reduces surplus per irrigated hectare by 172 Rs, or 0.7% of agricultural productivity per hectare in India in 2009. Across monitoring wells in India, one standard deviation of depth to water table is 15.4m, implying a one standard deviation increase in depth to water table would cause a loss of surplus per irrigated hectare equal to 10.8% of 2009 Indian agricultural productivity per hectare.

To assess the plausibility of this estimate, I do an alternative calculation. Instead, I ask how much farmers' private electricity costs of pumping groundwater would increase if depth to water table fell by 1m; an appeal to the envelope theorem suggests this is a direct loss of surplus for farmers. I then scale this up by the inverse share of electricity costs in costs of declining water tables; I consider values of 3 and 6 for this.[34] The IV and weighted IV estimates of the marginal surplus effect are 4.3 and 5.5 times larger than the increase in farmers' private electricity costs of pumping groundwater from a 1m decline in water tables, respectively. I interpret this as validation of that these estimates are reasonable to use for the remaining counterfactuals.

Next, I use the estimated marginal surplus effects, or local average surplus effects, to calculate the lost surplus from declining water tables in Haryana, Punjab, and Rajasthan, from 2000-2010, as estimated by Rodell et al. (2009). My preferred estimate, using the WIV marginal surplus effect, finds lost surplus of 365 Rs/ha, or 1.16% of agricultural productivity per hectare in northwest India. Other estimates range from 251 to 430 Rs/ha, while back of the envelope calculations scaling increased electricity costs are 197 and 395 Rs/ha.

1.6 Robustness

I present an analysis of robustness of the estimated local average surplus effect here. Sections 1.6.1, 1.6.2, and 1.6.3 discuss the exclusion restrictions that the outcome in-

[34] I calculate this share in two ways. For the first approach, I begin by noting that, on the margin, costs of adopting irrigation should equal benefits. I therefore use the IV LATE for the cost instrument on agricultural productivity as a measure of the costs of adopting irrigation. Next, I assume that the share of electricity costs in costs of declining water tables equals one minus the share of irrigation infrastructure in costs of adopting irrigation. Lastly, I use the IV LATE on fixed costs as a measure of fixed costs of adopting irrigation. To convert this to a flow, I multiply by 0.2, a common interest rate on credit in India (Hussam et al., 2017). This calculation yields an electricity cost share of 0.5. Alternatively, I assume that only fixed costs and electricity costs increase when water tables decline, and I assume they do so in proportion to their aggregate shares. I calculate the share of fixed costs using the approach above, and I calibrate electricity expenditures per irrigated hectare at 1,470 Rs/ha. This calculation yields an electricity cost share of 0.12. These yield a range of 2 to 8.

strument does not affect costs, that the outcome instrument does not affect potential revenue under rainfed agriculture, and that the cost instrument does not affect potential revenue, respectively. Section 1.6.4 discusses potential violations of the weak separability assumption. Section 1.6.5 discusses endogenous attrition, or that the instruments may increase gross cultivated area.

1.6.1 $W_n \not\Rightarrow C_{1i}$

The outcome instrument W_n might affect costs of agriculture if farmers reoptimize in response to increases in potential revenue under irrigation, and increase expenditures on inputs conditional on irrigating. If this is the case, direct effects on potential revenue driven by W_n may be the sum of increases in surplus and increases in costs; any such increases in costs are an exclusion restriction violation. To test this, in Column 4 of Table 1.4, I use household level data on agricultural input expenditures from NSS '12 as the outcome, and I compare instrumental variable estimates using the cost instrument Z_n and the outcome instrument W_n of the effect of irrigation D_n. Additionally, the cost instrument Z_n should have a direct effect on input expenditures related to pumping groundwater, so I exclude these.[35] This is a standard overidentification test: both Z_n and W_n should be valid instruments for the effect of irrigation on agricultural inputs excluding direct expenditures on irrigation if farmers do not reoptimize. Row 5 shows I fail to reject this overidentification test, and the estimate is a precise 0.

Alternatively, the outcome instrument may affect direct costs of irrigating through falling water tables. The outcome instrument should cause increases in extraction of groundwater, which would cause water tables to fall, which in turn will increase costs of irrigation. I test for this in Table 1.9. In Columns 6 and 7 of the first subtable, I fail to reject the null of no depletion caused by increases in irrigation caused by W_n. However, the coefficients are not small: they suggest a fully irrigated district has water tables that are 18m deeper than a district with no irrigation (1.2 standard deviations of depth to water table across monitoring wells), and depletion is 2m/year faster. However, this will not meaningfully bias my estimates: multiplying 18m by the 172 Rs/ha cost increase caused by a 1m fall in water tables, this implies that costs increased by 3,110 Rs/ha, which is less than 10% of my estimates of the local average surplus effect.

[35] Specifically, I drop the categories "Diesel", "Electricity", and "Irrigation". While one might be tempted to use these categories to construct a measure of agricultural profits, they crucially do not include depreciation of irrigation infrastructure.

1.6.2 $W_n \not\Rightarrow Y_{0i}$

The outcome instrument W_n might affect potential revenue under rainfed agriculture; it is constructed using FAO GAEZ data on predicted relative yields under irrigated agriculture. This is negatively correlated with predicted yields under rainfed agriculture, as places with high returns to irrigation typically have low yields under rainfed agriculture. I address this in two ways. First, I consider including more or less flexible controls for FAO GAEZ potential rainfed crop yield. All primary specifications include controls for state fixed effects interacted with potential rainfed crop yield, I compare this baseline specification to specifications with alternative controls in Table 1.10. First, Column 2 shows a specification with no controls. The estimated local average surplus effect is biased downward, as relative potential irrigated yields are negatively correlated with rainfed yields. Columns 3, 4, 5, and 6 include progressively more flexible controls, with controls in my preferred specification (in Column 1) falling between Column 4 and Column 5. Estimates of the local average surplus effect range from 39,600 Rs/ha to 56,900 Rs/ha, compared to 31,700 Rs/ha with unweighted instrumental variables, although the precision begins to decrease as more controls are added.

Alternatively, the effect of the outcome instrument on rainfed yields is identified. Flexible models which allow for this in Ag '07-'11 are underpowered, but the control function approach I implement in NSS '12 is sufficiently powered to test this under more parametric restrictions. I implement this overidentification test in Row 2 of Table 1.5; I fail to reject the outcome instrument has no effect on rainfed yields, and the 0 is small and precise. I assess the magnitude of bias from exclusion restriction violations in Table 1.4, Column 3: the bias in instrumental variables from violations of the exclusion restriction is estimated to be 4,600 Rs/ha, less than 10% of the control function estimate of the local average surplus effect.

1.6.3 $Z_n \not\Rightarrow (Y_{0i}, Y_{1i})$

The cost instrument Z_n decreases costs of groundwater irrigation by enabling lower cost tubewell irrigation. In India, prior to the Green Revolution, almost no agricultural land was irrigated using tubewells, so the cost instrument should have no effect on irrigation or agricultural revenue before the start of the Green Revolution. I estimate a difference in difference specification in Table 1.11, comparing coefficients on the cost instrument Z_n, the outcome instrument W_n, and the rainfed yield control $\log \text{RF yield}_n$, along with their interactions with a post Green Revolution start dummy.[36] To facilitate comparison across years, I use log agricultural productivity instead of its level. The cost instrument has no significant effects on irrigation or agricultural productivity before the Green Revolution, when tubewells are not available as a technology. In contrast, the

[36] I follow Sekhri (2014) and define 1966 to be the start of the Green Revolution.

outcome instrument increases revenue even before the Green Revolution, as other forms of irrigation were already available as a technology. However, the outcome instrument has limited effects on adoption of irrigation: increases in the returns to irrigation have a small effect on adoption of irrigation when there is large variation in the costs of irrigation, as was the case before the expansion of tubewell irrigation.

Alternatively, the cost instrument Z_n might affect potential revenue directly if farmers reoptimize in response to decreases in costs of irrigation, and increase expenditures on inputs conditional on irrigating. To some extent, Column 4 of Table 1.4 should alleviate those concerns, as effects of the cost instrument on input expenditures are small. However, I explicitly excluded any expenditures specific to irrigation, as the cost instrument should have direct negative effects on these. Additionally, that the magnitudes of the LATE estimates in Columns 1 and 2 of Table 1.3 are reasonable should alleviate concerns of large bias, but given the limited precision with which they are estimated, this is also insufficient. To construct a test for reoptimization, I argue that if falling costs of irrigation cause farmers to reoptimize, we should see shifting of crop choice under irrigation towards water intensive crops; this appears as a violation of monotonicity, where the instrument decreases area irrigated under crops with low water intensity. I test for this in Table 1.12. Because I test for effects on every crop in the data, I adjust inference for multiple hypothesis testing; after this adjustment, no monotonicity violations are detected. Decreases in costs of irrigation cause shifts away from rainfed rice, maize, and wheat, and into irrigated rice.

1.6.4 Weak separability

In general, monotonicity with multiple instruments is a much stronger assumption than monotonicity with a single instrument. This is equally true here: through the lens of the model, it requires farmers can only differ in their responsiveness to the instruments through $V_{\gamma i}$. This is violated if some farmers' surplus under irrigation is relatively more responsive to the cost instrument. I consider likely violations of this in this section.

The clearest violation of monotonicity is the presence of surface water. Farmers with access to surface water will not have their costs of irrigation respond to the cost instrument, since they will irrigate using surface water even if their costs of pumping groundwater fall. However, these farmers will still respond to the outcome instrument, since their revenue under irrigation will still shift up. Let Surface_i be a dummy for access to surface water. To see how this violates monotonicity, one can write this modified model as

$$\begin{aligned} Y_{1i}(w) &= V_{\gamma i}\gamma_W(w) + V_{1i} \\ C_{1i}(z) &= (1 - \text{Surface}_i)V_{\gamma i}\gamma_Z(z) + V_{Ci} \\ Y_{0i} &= V_{0i} \end{aligned}$$

I take two approaches to handling this. First, I drop states where more than one third of irrigation is surface water, and present results in Column 2 of Table 1.13. States with large shares of surface water may bias up estimation of a local average surplus effect, if the outcome instrument increases revenues in those states but does not affect adoption of irrigation. The estimated local average surplus effect restricted to states with low shares of surface water is in fact slightly larger, suggesting such bias is not large in this context.

Second, I take a more model driven approach. I make the additional assumption that $\text{Surface}_i \perp (W_i, Z_i, V_{1i}, V_{Ci}, V_{0i}, V_{\gamma i})|X_i$, or that access to surface water for irrigation is exogenous conditional on the controls X_i. Additionally, I assume that everyone with access to surface water irrigates. This latter assumption I test: I show in Columns 1 and 2 of Table 1.9 that the outcome instrument (in the first subtable) and the cost instrument (in the second subtable) cause significant increases in groundwater irrigation, but not surface water irrigation. Under these assumptions, all results on estimation still hold, but when conducting counterfactuals using the local average surplus effect that affect only groundwater, it must be scaled down by the share of groundwater in irrigation. When I applied the local average surplus effect to estimation of the welfare losses from falling water tables in Section 1.5.4, the estimate of the effect of falling water tables on groundwater irrigation I use was from communities without access to surface water irrigation. On the other hand, when I use the local average surplus effect to recover an estimate of the elasticity of irrigation to the price of electricity, I must account for having estimated the local average surplus effect nationally, where the groundwater share of irrigation is 0.66.

1.6.5 Attrition

An addition concern is attrition: when costs of irrigation fall, some farmers will shift from rainfed agriculture to irrigated agriculture, but land that was fallow will also become irrigated, and farmers may begin to multiple crop. This constitutes endogenous selection into the sample. To account for this, I allow land to shift from either rainfed agriculture or fallow into irrigated agriculture in response to the instruments. Instead of looking at the share of agricultural land that is irrigated, I look at the share of district land that is irrigated. However, I do not observe the reservation rent on fallow land, or the gross revenue under rainfed agriculture that land would need to yield in order to be cultivated. However, an extended model implies that selection out of fallow should be the same in response to the outcome instrument and the cost instrument, so I test robustness of the results to imputation of a range of reservation rents; I use both 0 Rs/ha and 20,000 Rs/ha (just under the average revenue per hectare on rainfed plots in NSS '12). The results of this exercise are in Columns 3 and 4 of Table 1.13. The estimated local average surplus effect is smaller, but not significantly different, and does not depend on the choice of reservation rent.

1.7 Optimal policy

In Section 1.5.4, I calculated the lost surplus per hectare from a one meter decline in the water table. I now apply this estimate to optimal policy for groundwater subsidies. As discussed in Section 1.2.1, irrigation is implicitly subsidized in India through subsidies for electricity for pumping groundwater. Although there is not volumetric electricity pricing, pump capacity fees implicitly price electricity at an average of one third of marginal cost (Fishman et al., 2016; Badiani & Jessoe, 2017). Following Allcott et al. (2014), I consider a policy maker maximizing social surplus in choosing how to set pump capacity fees. Despite deadweight loss, subsidies may be optimal because the policy maker has a preference for redistribution, and is willing to spend $\lambda > 1$ Rs to transfer 1 Rs to farmers, a stated motive behind electricity subsidies (Dubash, 2007).[37] However, the impacts of marginal pumping induced by the subsidy on depth to water table of other farmers are not internalized by farmers increasing their pumping. This negative externality, and the deadweight loss from the subsidies, must be traded off by the social planner against the value of the subsidies as a transfer.

In Section 1.7.1, I model the planner's problem, and in Section 1.7.2, I discuss calibration of key parameters, including the marginal surplus effect. In Section 1.7.3, I use the model to calculate the gains from decentralizing the setting of pump capacity fees in Rajasthan. Rajasthan is in northwestern India, where I estimated the lost surplus from declining water tables in Section 1.5.4, and relative to other states in the region has greater heterogeneity of aquifer characteristics, and therefore in the magnitude of the negative externality. I quantify potential gains from reducing relative subsidies in districts with large negative pumping externalities.

[37] Whether the policy maker is justified in acting as if $\lambda > 1$ is a question beyond the scope of this paper, but for electricity subsidies $\lambda > 1$ may be efficient if other transfers to farmers create greater deadweight loss (Hendren, 2014) or have high leakage (Niehaus & Sukhtankar, 2013).

1.7.1 Planner's problem

I model groundwater irrigation closely following Shah et al. (1995). In period t, farmers have access to an available stock of groundwater, S_t, from which they can pump groundwater for irrigation. If farmer i irrigates ($D_{it} = 1$), they receive revenue $Y_{1i}(a_{it})$ and incur costs $C_{1i}(a_{it}; S_t)$, where a_{it} is quantity of water farmer i would extract to maximize surplus conditional on irrigating in period t. If farmer i does not irrigate, they receive revenue Y_{0i}. Costs $C_{1i}(a_{it}; S_t)$ include fixed costs $k_i(S_t)$, linear electricity costs $m_i(S_t)p_t a_{it}$, where p_t is the price per kWh in period t, and other linear variable costs $c_i(S_t)a_{it}$. Farmers are atomistic, in that farmers do not internalize any impact their extraction a_{it} has on the available stock of groundwater S_t. Farmers maximize surplus π_i by solving

$$\pi_i = \int_0^T e^{-rt} \max_{a_{it}, D_{it}} \left[D_{it} Y_{1i}(a_{it}) - D_{it} \underbrace{\left((c_i(S_t) + m_i(S_t)p_t)a_{it} + k_i(S_t)\right)}_{C_{1i}(a_{it}; S_t)} + (1 - D_{it})Y_{0i} \right] dt \tag{1.28}$$

I make a few additional realistic assumptions on electricity use and groundwater extraction. I model the evolution of the stock of groundwater simply; it falls by one unit per unit of extraction, so $\dot{S}_t = A_t \equiv \int D_{it} a_{it} di$. To extract a unit of water, the electricity required $m_i(S_t) = (h_i + b(S_t))m$, where $h_i + b(S_t)$ is the depth to groundwater for farmer i. The electricity requirement per unit of water per meter of depth to groundwater, m, is simply the energy required to lift one unit of water by one meter divided by the pump efficiency. The global component of depth to groundwater, $b(S_t) = S_t/\alpha\overline{L}$, where α is the specific yield of the aquifer (the fall in the water table per unit of groundwater extracted), and $\overline{L}$ is the area of the aquifer in hectares; as a result, when one meter hectare of groundwater is extracted, farmers experience an increase in depth to groundwater of $1/\alpha\overline{L}$ meters.

The social planner chooses p_t, the price of electricity charged to farmers, to maximize social surplus. Social surplus is total farmer surplus times λ plus profits from the electricity sector. Total agricultural electricity use in period t is $M_t \equiv \int D_{it} m_i(S_t) a_{it} di$, and the cost of producing a unit of electricity is c_t. The social planner solves

$$\max_p V(p) \equiv \lambda \int \pi_i di + \int e^{-rt}(p_t - c_t) M_t dt \tag{1.29}$$

I make three additional simplifications. First, I ignore rebound effects, where increases in the price of electricity today, by reducing extraction of groundwater, increase the available stock of groundwater, which reduces future costs of extraction and in turn increases future extraction. I further assume that current extraction is a good approximation of future extraction. In fact, extraction is growing (Rodell et al., 2018). These two simplifications have offsetting effects: rebound implies externalities are smaller than I estimate, while growing extraction implies externalities are larger than I estimate. I anticipate that these biases are small, as my calibrated elasticity is low (which reduces the bias from ignoring rebound) and my calibrated discount rate is high (which reduces the bias from ignoring rebound and growth in extraction). Third, I assume that current costs of electricity generation and electricity subsidies are a good approximation of future costs and subsidies. This is difficult to know, but I consider it a natural starting point for analysis.

I consider the social planner's first order condition for social surplus maximization with respect to the period 0 price of electricity. When writing the social planner's first order condition, I normalize by total electricity use M_0, and multiply by -1; this normalized first order condition can be interpreted as changes in social welfare per rupee of surplus transferred to farmers. I follow the public economics literature and express this first order condition in terms of reduced form sufficient statistics (Chetty, 2009). I define $\epsilon_{M,p}$ to be the elasticity of electricity use to the price of electricity, and $\epsilon_{A,p}$ to be the elasticity of groundwater extraction to the price of electricity.

$$
\begin{aligned}
-\frac{1}{M_0}\frac{dV(p)}{dp_0} = & \underbrace{\lambda - 1}_{\text{Transfer value}} \underbrace{- \epsilon_{M,p}\frac{p_0 - c_0}{p_0}}_{\text{DWL}} \\
& \underbrace{-\frac{\lambda}{r}\epsilon_{A,p}\frac{(L/\alpha\overline{L})\overbrace{(\partial\mathbf{E}[D_{i0}]/\partial b_0)\text{MSE}(\mathbf{E}[D_{i0}])}^{\text{Farmer cost of 1m fall in water table/ha}}}{p_0M_0/A_0}}_{\text{Pumping externality (farmer)}} \underbrace{-\frac{1}{r}\epsilon_{A,p}\frac{(L/\alpha\overline{L})\overbrace{(p_0 - c_0)(mA_0/L)}^{\text{Utility cost of 1m fall in water table/ha}}}{p_0M_0/A_0}}_{\text{Pumping externality (utility)}}
\end{aligned}
\tag{1.30}
$$

I consider each term in Equation 1.30. The first term, $\lambda - 1$, is the value the social planner places on shifting one rupee from public funds to farmers. The second term, $-\epsilon_{M,p}\frac{p_0-c_0}{p_0}$, is the standard deadweight loss term. It is the elasticity of electricity use to the price of electricity times a term that captures the distortion from subsidies.

The third term, $\frac{\lambda}{r}\epsilon_{A,p}\frac{(L/\alpha\overline{L})(\partial\mathbf{E}[D_{i0}]/\partial b_0)\text{MSE}(\mathbf{E}[D_{i0}])}{p_0M_0/A_0}$, is the pumping externality experienced by farmers per Rs of transfer. It is scaled by λ, because changes in farmer surplus, whether from transfers or increased pumping costs from externalities, are valued the same by the social planner. It is scaled by $\frac{1}{r}$, because while transfers are experienced immediately, and deadweight loss is based on the farmer's static optimization, the externality from a unit fall in the water table is experienced indefinitely by all farmers. It is scaled by $\epsilon_{A,p}$ because the externality caused per rupee of transfer is proportional to the extraction caused per rupee of transfer. The remainder $\frac{(L/\alpha\overline{L})\frac{\partial \mathbf{E}[D_{i0}]}{\partial b_0}\text{MSE}(\mathbf{E}[D_{i0}])}{p_0 M_0/A_0}$ captures the distortion. The numerator is the externality per unit of water extracted, and equals the fall in water table experienced by farmers per unit of water extracted $L/\alpha\overline{L}$ times the lost farmer surplus per unit fall in the water table $\frac{\partial \mathbf{E}[D_{i0}]}{\partial b_0}\text{MSE}(\mathbf{E}[D_{i0}])$. The denominator is the electricity cost per unit of water extracted, p_0M_0/A_0. The full term $\frac{1}{r}\epsilon_{A,p}\frac{(L/\alpha\overline{L})\frac{\partial \mathbf{E}[D_{i0}]}{\partial b_0}\text{MSE}(\mathbf{E}[D_{i0}])}{p_0 M_0/A_0}$, is the externality ratio, or the Rs of externality created per Rs of surplus transferred to farmers.

The fourth term, $\frac{1}{r}\epsilon_{A,p}\frac{(L/\alpha\overline{L})(p_0-c_0)(mA_0/L)}{p_0M_0/A_0}$, is the pumping externality experienced by the utility per Rs of transfer. The utility experiences the externality because of the wedge between the price farmers pay for electricity and the marginal cost of generation. It is scaled by $\frac{1}{r}$, $\epsilon_{A,p}$, and inversely proportional to p_0M_0/A_0 for the same reasons the pumping externality experienced by farmers is. The numerator, $(L/\alpha\overline{L})(p_0-c_0)(mA_0/L)$, is lost profits experienced by the utility per unit of water extracted caused by the increase in electricity required to pump groundwater caused by falls in the water table. The wedge $p_0 - c_0$ is the future difference between the price of electricity and the marginal cost of generation, as the increased electricity use caused by the externality occurs indefinitely.

1.7.2 Calibration

I discuss a few key aspects of the calibration. Note that all parameters used in the calibration are in Table 1.7.

First, I take two approaches to calibrating $\epsilon_{A,p}$ and $\epsilon_{M,p}$. In both cases, I assume electricity use for extracting groundwater is a constant proportion of extraction, so $\epsilon_{A,p} = \epsilon_{M,p}$. This need not hold in the model above, in the presence of heterogeneity in responsiveness to the price of electricity that is correlated with idiosyncratic depth to groundwater h_i. For the first approach, I use an estimate from Badiani & Jessoe (2017), $\epsilon_{A,p} = -0.18$. For the second approach, I use my preferred estimate of a local average surplus effect to calculate this elasticity; the inverse of a local average surplus effect is a semielasticity of irrigation to its gross returns. I calculate $\epsilon_{A,p} = -0.045$.[38] This estimate is likely to be biased downwards, since it ignores intensive margin responses of extraction to changes in the subsidy. I therefore interpret it as a lower bound, and I show estimates using both $\epsilon_{A,p} = -0.18$ and $\epsilon_{A,p} = -0.045$.

Second, the numerator of the externality ratio, $(L/\alpha\overline{L})(\partial\mathbf{E}[D_{i0}]/\partial b_0)\text{MSE}(\mathbf{E}[D_{i0}])$ can

[38]Specifically, I approximate $\epsilon_{A,p} \approx \frac{p_0M_0/\mathbf{E}[D_{i0}]L}{0.66\text{LASE}}$, where 0.66 is the groundwater share of irrigated land. I use LASE = 49,800 Rs/ha, and electricity expenditures per irrigated hectare by farmers of $p_0M_0/\mathbf{E}[D_{i0}]L = 1{,}470$ Rs/ha.

be decomposed into the product of three terms. The first, $1/\alpha$, is the inverse specific yield of the aquifer, or the total fall in the water table per unit of water extracted. The second, $L\mathbf{E}[D_{i0}]/\overline{L}$ is the share of the aquifer that is irrigated; this captures the fraction of a fall in the water table experienced by farmers. These first two terms will vary across aquifers, which may fall within district or cross district boundaries. For this exercise I assume each district is a single, contiguous aquifer; however, with more granular data, this exercise is straightforward at the aquifer level. The third, $(\partial\mathbf{E}[D_{i0}]/\partial b_0)(\text{MSE}(\mathbf{E}[D_{i0}])/\mathbf{E}[D_{i0}])$, is the lost surplus per irrigated hectare per unit fall in the depth to groundwater. My preferred estimate of this is 172 Rs/ha/m in Table **??**, which I use for this exercise.

Third, calculating m, the electricity needed to pump one unit of groundwater one meter, is a simple physics problem which depends only on the depth to water table and the efficiency of extraction. Shah (2009) suggests 40% is a reasonable efficiency in the Indian context. Further, I assume that $M_0 = A_0 b_0 m$, or that electricity use for irrigation is groundwater extraction times depth to groundwater times the electricity needed to pump one unit of groundwater one meter.[39] This calculation yields total agricultural electricity use that is 36% of reported electricity use. I assume this difference is driven by depth to water table in farmers' wells being significantly deeper than the depths to water table in India's monitoring wells. I scale up my estimates of electricity use M_0 by a constant proportion across districts to match this total.

Fourth, a key decision is which parameters I allow to vary across districts. In this exercise, I focus on heterogeneity in optimal subsidies that stems from variation in the magnitude of the pumping externality. I therefore allow the key parameters which determine the pumping externality to vary: the average specific yield, the depth to water table, and the irrigated share of land. The externality ratio is inversely proportional, inversely proportional, and proportional to each of these parameters, respectively. I do a variance decomposition of the log externality ratio across districts: 11% of the variation is attributed to specific yield, 52% is attributed to irrigated share of land, and 37% is attributed to depth to water table.[40]

Fifth, for counterfactuals, a necessary decision is to determine which parameters are permitted to respond endogenously to changes in the policy, and which are not. The only parameters I allow to vary in response to changes in p are A_0, the total extraction of groundwater in the current period, and $\mathbf{E}[D_i]$, the irrigated share of the aquifer. For both, I use $\epsilon_{A,p}$ as the relevant elasticity. As mentioned previously, I ignore rebound; equivalently stated, I do not allow farmers to respond to changes in depth to water

[39] Depth to groundwater is measured using the median depth to groundwater by district across monitoring tubewells in Well '95-'17

[40] The externality experienced by the utility varies with the extraction of groundwater per irrigated hectare by district, which I also allow to vary. Setting this to the average extraction across districts does not meaningfully change any results, so I do not emphasize it.

table b_t, but I do calculate the changes in rates of depletion implied by the changes in A_0. Additionally, I undertake the analysis as if the policy change were permanent; future decreases in $\mathbf{E}[D_i]$ caused by increases in electricity prices reduce negative externalities, and future increases in $p_t - c_t$ caused by increases in electricity prices reduce the negative externality on the utility. Both of these effects reduce the magnitude of optimal variation in subsidies relative to ignoring these responses. In sum, this represents a compromise between a full numerical simulation of the model, as would be standard in the optimal control literature, and the simpler sufficient statistics approach I undertake, and I leave I comparison of my approach to a full numerical simulation to future work.

Sixth, for aggregating across districts, it is necessary to know district specific levels of extraction A_0 at baseline subsidy levels; I collect this data from district groundwater brochures from the Central Ground Water Board, which estimate groundwater withdrawals in each district in an idiosyncratic year ranging from 2004 to 2011, with a modal year of 2008.

Seventh, I make two sample restrictions for districts for the counterfactual exercise. First, I only use districts for which depth to water table, district irrigated land share, and average aquifer specific yield are available; this brings me from 24 districts in the main analysis to 22. Second, I drop districts where more than 7% of irrigation uses surface water. In districts with high levels of surface water irrigation, optimal policy requires a different set of considerations: surface water irrigation has positive externalities, as it causes recharge of groundwater, and surface water and groundwater irrigation may be substitutes. This reduces the set of districts from 22 to 14.

1.7.3 Results

Figure 1.4 presents the optimal district specific electricity taxes in Rajasthan. To calculate optimal taxes, I first calibrate the social planners willingness to pay to increase farmer surplus by 1 unit, λ, under the assumption current policy is optimal subject to the constraint that there is a single subsidy at the state level, which yields $\lambda = 1.56$. Note that this λ is just the inverse marginal value of public funds; as a reference, this is similar to the inverse marginal value of public funds for SNAP, a public assistance program in the United States, as calculated in Hendren (2016).

Panel (a) presents the optimal tax by district. The optimal tax is relatively low in districts in northwestern Rajasthan, which tend to have lower land shares of irrigation, cultivating bajra instead of more water intensive wheat and maize, lower depths to water table, and higher specific yields, and therefore relatively small pumping externalities. Panel (b) presents the externality ratio and deadweight loss in each district as a function of the electricity tax. First, note that negative externalities are almost triple deadweight loss in the highest externality district, but close to 0 in other districts. Second, current subsidy levels reduce farmer surplus on the margin in the district with the largest pumping externalities, as the marginal pumping induced by current levels of subsidies in that district reduces farmer surplus by more than their value as a transfer.

Table 1.14 presents results for total subsidies, deadweight loss, farmer surplus, and groundwater depletion, all relative to a no subsidy policy, under three scenarios. Column 1 presents the status quo. Total subsidies equal 6.6% of agricultural production, but deadweight loss from the subsidies is 0.65% of agricultural production, despite the high subsidy level. This follows from the low estimate of the price elasticity of electricity demand in agriculture I use from Badiani & Jessoe (2017). Externalities experienced by the utility are small relative to externalities experienced by farmers, as despite the high subsidies, electricity for pumping groundwater is a low share of costs of falling water tables. Negative pumping externalities induced by subsidies are meaningful, at 0.45% of agricultural production, but smaller than deadweight loss; however, this masks substantial heterogeneity. Additionally, subsidies were responsible for declines in water tables of 1.51m from 2000-2010, 46% of the observed decline in northwestern India.

Column 2 of Table 1.14 presents a scenario where the social planner chooses district specific subsidies to maximize social welfare under the same λ that implies the policy in Column 1 is the optimal state level policy, while holding total subsidies fixed. This policy involves increasing subsidies in districts with small pumping externalities, while decreasing subsidies in districts with large pumping externalities. First, note that this policy increases deadweight loss: this follows from the constant elasticity assumption, which implies a constant subsidy across locations minimizes deadweight loss holding fixed total subsidy payments. However, the increased deadweight loss is smaller than the decrease in negative pumping externalities. Negative externalities relative to the no subsidy policy fall by 25%, the total distortion relative to no subsidy falls by 7%, and the effect of subsidies on depth to groundwater decreases by 16%. However, total farmer surplus increases by only 0.07% of agricultural production.

Columns 4 and 6 present equivalent exercises, but using a lower calibrated elasticity (0.045) and a lower calibrated discount rate (0.08), respectively. Focusing on Column 4, the lower elasticity implies the inefficiency from subsidies is small: the λ which implies current policy is the optimal state level policy is 1.12. As a result, potential gains from spatially explicit policy are small. This highlights the importance of having a more precise estimate of this elasticity. Focusing on Column 6, the lower discount rate magnifies externalities, which in turn increases the potential gains from spatially explicit policy from 0.07% of agricultural production to 0.29% of agricultural production. It also implies that subsidies are very inefficient as transfers due to large negative externalities.

In this exercise, although "optimal" district specific subsidies increase total surplus, for high calibrations of the discount rate they do reduce farmer surplus in high externality districts, as relatively inefficient subsidies are reduced in those districts. As a result, this "optimal" policy may not be politically feasible. However, alternative more feasible policies can replicate the proposed optimal electricity tariff, while generating potentially larger gains. First, Badiani & Jessoe (2017) and Fishman et al. (2017) find that responses to changes in the cost of groundwater extraction tend to be on the extensive margin (in reduced area under irrigation) and not intensive margin (through reduced pumping). As a result, impacts of changing electricity tariffs can be replicated through other policies that change incentives to irrigate.[41] Additionally, Chatterjee et al. (2017) document that output subsidies for water intensive crops create incentives to increase groundwater extraction. Therefore, policies which reduce input subsidies complementary to irrigation or output subsidies for water intensive crops while increasing subsidies for inputs complementary to rainfed agriculture could increase the efficiency of farmer subsidies, especially in districts with large pumping externalities.

1.8 Conclusion

This analysis suggests that groundwater depletion in India from 2000-2010 permanently reduced economic surplus by 1.2% of gross agricultural revenue. This is similar to anticipated losses in India due to climate change of 1.8%/decade under the 4°C warming scenario (Government of India (2018)), and is especially concerning given accelerating rates of depletion (Jacoby (2017)). Policy solutions without economic tradeoffs may not be easy to come by: without reducing total electricity subsidies, the spatially explicit subsides I study can only increase surplus by a magnitude equal to losses from less than 1 year of groundwater depletion. Moreover, this policy reduces farmer surplus in districts with large externalities, and therefore may be politically infeasible. However, understanding the magnitudes of these externalities and the losses from depletion enables quantifying the potential efficiency gains from investments in surface water irrigation, or subsidies for inputs complementary to rainfed agriculture.

To undertake this analysis, I have expanded on tools from the program evaluation literature and microeconomic theory to define the marginal surplus effect. While marginal treatment effects capture the impact of policies or shocks which increase adoption of some treatment (such as college attendance) on observable outcomes, marginal surplus effects capture the direct impact of these policies or shocks on the economic surplus of inframarginal adopters. This is an important metric for policy across a range of contexts, such as health and safety regulations for workers, environmental regulations for firms, or, in this study, groundwater depletion in agriculture.

[41] Note that implementation of volumetric pricing could have a very different set of impacts on electricity use, especially with respect to efficiency, than the changes to electricity pricing as implemented through pump capacity fees that I consider.

Figure 1.1: Cost and benefit shifters

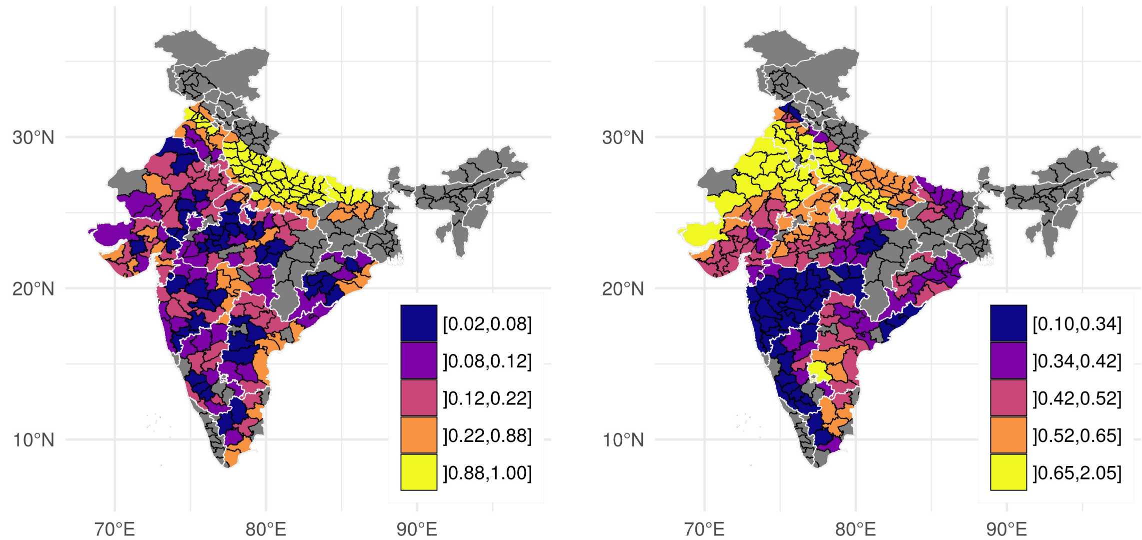

(a) Potential aquifer yield Z_n

(b) log rel. potential irrigated crop yield W_n

Notes: Variation in the cost instrument Z_n (potential aquifer yield, Panel (a)) and the outcome instrument W_n (log relative potential irrigated crop yield, Panel (b)) across districts in India is presented here. Colors correspond to quintiles of their respective distributions. District boundaries are in black, and state boundaries are in white.

Figure 1.2: Model comparative statics

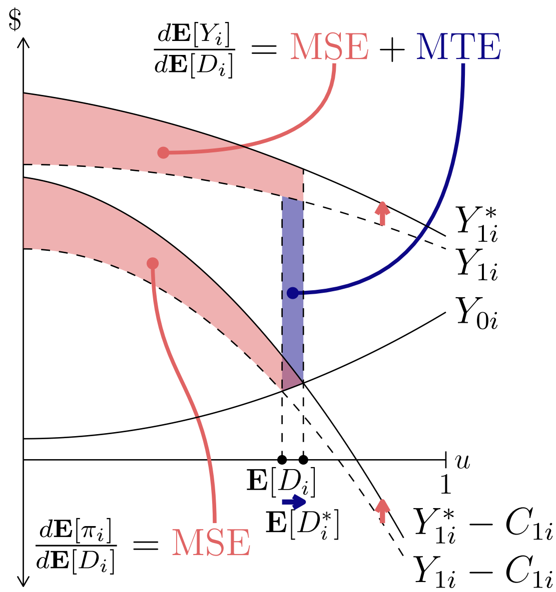

(a) Increased outcome under treatment

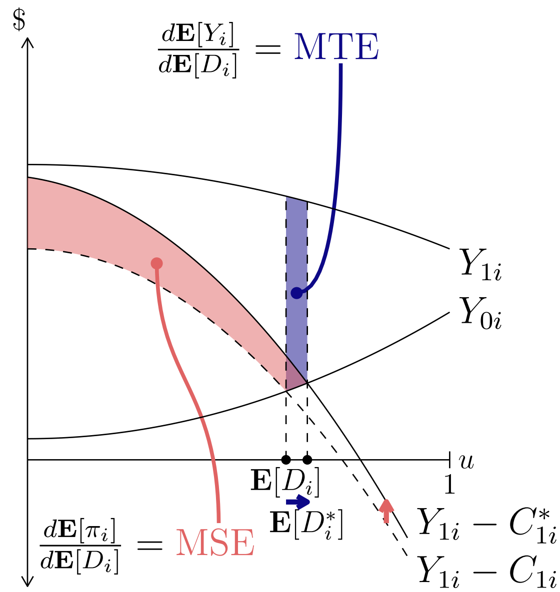

(b) Decreased costs of treatment

Notes: Panel (a) shows the effects of shifting w, the instrument for potential outcome under treatment (which shifts potential outcome under treatment Y_{1i} to Y_{1i}^*), while Panel (b) shows the effects of shifting z, the instrument for costs of adopting treatment (which shifts costs C_{1i} to C_{1i}^*). Changes in the share of agents adopting treatment, from $\mathbf{E}[D_i]$ to $\mathbf{E}[D_i^*]$, are displayed. Changes in average surplus $\mathbf{E}[\pi_i]$ or changes in average outcomes $\mathbf{E}[Y_i]$ are shaded. Marginal treatment effects are in purple, and are equal to the change in average outcomes per unit change in adoption of treatment caused by shifts to z. Marginal surplus effects are in pink, and are equal to the change in average surplus per unit change in adoption of treatment caused by shifts to either z or w. The change in average surplus caused by both z and w is proportional to the marginal surplus effect. However, the change in average outcomes caused by z is proportional to the marginal treatment effect, while the change in average outcomes caused by w is proportional to the marginal surplus effect plus the marginal treatment effect.

Figure 1.3: Marginal surplus effect estimates

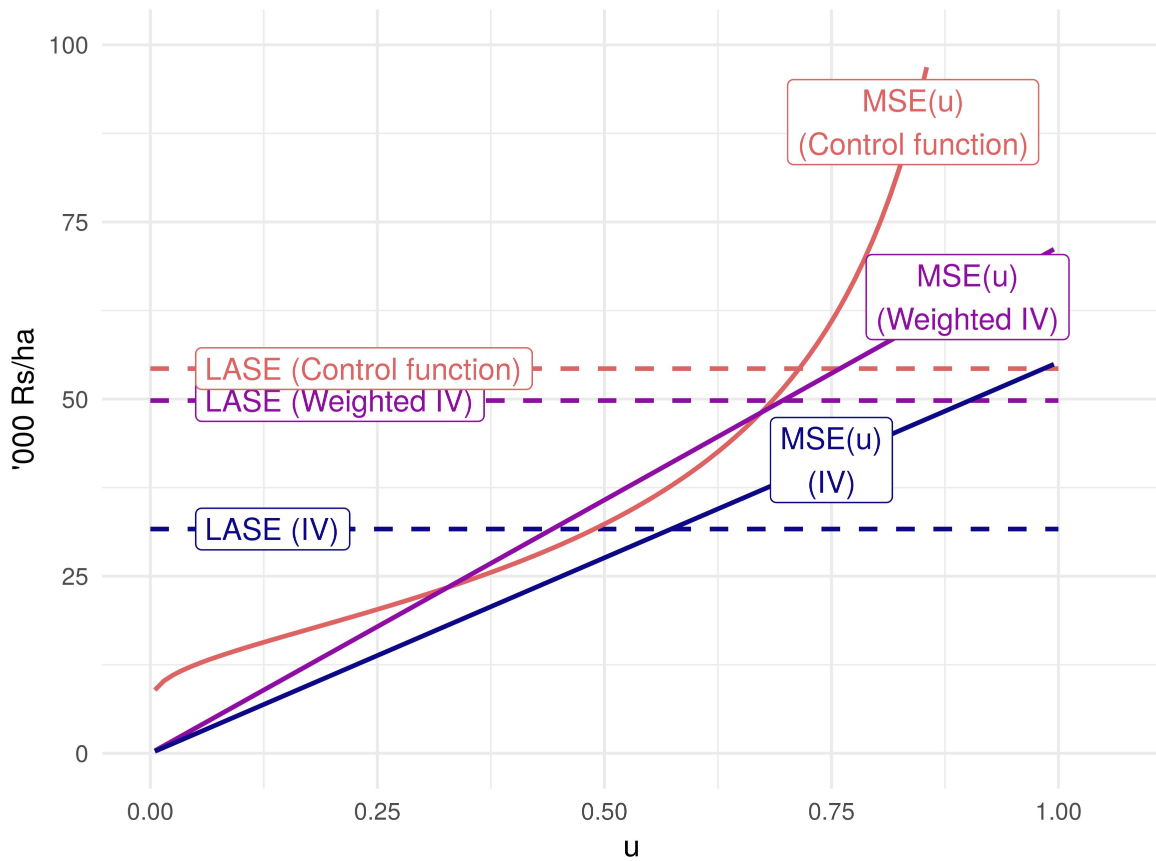

Notes: Solid lines present estimates of marginal surplus effects (the change in average surplus per unit change in adoption caused by shifts to either costs or outcomes under treatment), while dashed lines present estimates of local average surplus effects (a weighted average of marginal surplus effects). Dashed lines for IV and Weighted IV estimators are the estimates of local average surplus effects used to construct marginal surplus effects, following Section 1.3.5. The control function estimate of the local average surplus effect is constructed by replacing outcomes and treatment in the IV regression using w with control function estimates of predicted changes in surplus and changes in propensity scores.

Figure 1.4: Optimal electricity taxes in Rajasthan

(a) Optimal tax, by district

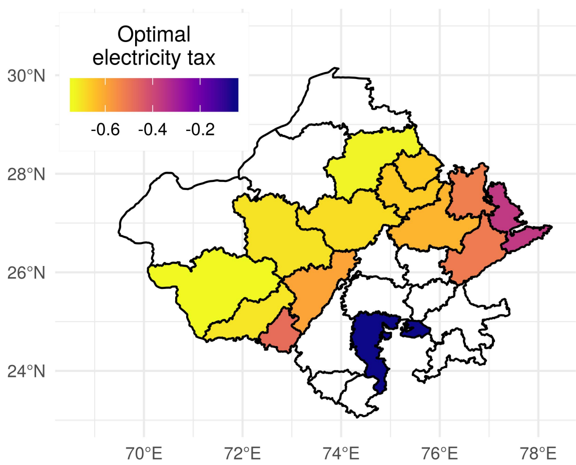

(b) Externality ratios and DWL, by district

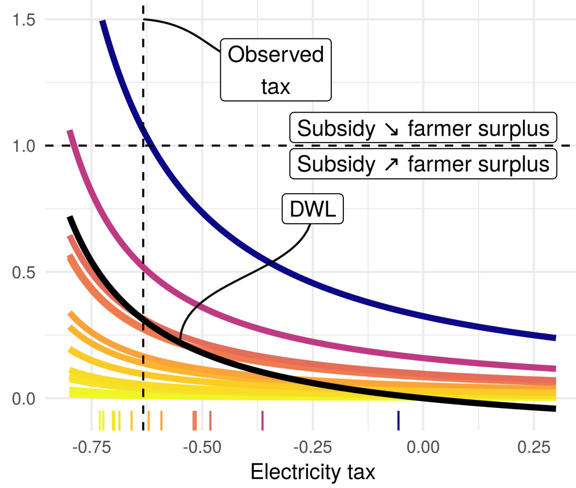

Notes: This figure presents the results of the optimal policy exercise. In Panel (a), I plot the optimal electricity tax by district in Rajasthan, dropping districts with missing data or high levels of surface water irrigation. In Panel (b) I plot farmer externality ratios (the negative externality on farmers created by induced marginal groundwater extraction per unit of transfer to farmers, which varies across districts) and deadweight loss (DWL) as a function of the electricity tax. The optimal electricity tax solves

$$\lambda - 1 = (\text{DWL}) + \lambda(\text{Farmer externality ratio}) + (\text{Utility externality ratio})$$

λ is the willingness to pay of the social planner to increase farmer surplus by 1 unit. I use $\lambda = 1.56$ for values reported in this figure, which implies current subsidies are optimal if the planner is constrained to a single state level subsidy. I assume a constant elasticity of demand for electricity and water to the price of electricity. Both deadweight loss and externality ratios vary with the tax as electricity use and groundwater extraction respond. Farmer externality ratios by district are plotted in Panel (b). These externality ratios drive variation across districts in the optimal tax, and are the product of the inverse specific yield, inverse depth to water table, and the share of aquifer irrigated. The vertical dotted line in Panel (b) is the observed tax in Rajasthan (Fishman et al., 2016), while the horizontal dotted line is at 1: as discussed in Section 1.7.1, when the farmer externality ratio is above 1, any subsidy decreases farmer surplus, while when the farmer externality ratio is below 1, any subsidy increases farmer surplus (although subsidies are still costly to the social planner, due to increased net fiscal outlays, deadweight loss, and negative externalities on utilities). A tick is added to the bottom of the graph for the optimal tax in each district.

Table 1.1: Descriptive statistics

	Mean	SD	Min	Max	# of obs.	# of clu.
Ag '07-'11						
Y_n Agricultural productivity ('000 Rs/ha)	24.9	15.1	1.3	125.0	884	222
D_n Share irrigated	0.550	0.273	0.017	1.000	884	222
Z_n Potential aquifer yield (40 L/s)	0.336	0.349	0.025	1.000	884	222
W_n log relative potential irrigated crop yield	0.533	0.254	0.098	2.050	884	222
X_n log potential rainfed crop yield (log t/ha)	0.690	0.503	-2.234	1.285	884	222
Share rice	0.268	0.265	0.000	0.977	884	222
Share wheat	0.211	0.190	0.000	0.631	884	222
NSS '12						
Y_i Agricultural productivity ('000 Rs/ha)	36.6	26.1	0.0	100.0	33,778	222
$Y_i \| D_i = 1$ Irrigated plots	44.9	26.3	0.0	100.0	23,957	220
$Y_i \| D_i = 0$ Rainfed plots	22.0	18.3	0.0	100.0	9,821	189
Area (ha)	1.778	2.540	0.001	40.823	33,778	222
D_i Irrigated	0.637		0.000	1.000	33,778	222
Agricultural inputs net irrigation ('000 Rs/ha)	15.4	15.5	0.0	100.0	26,280	222
Any bank loan	0.310		0.000	1.000	26,280	222
Irr '07						
Infrastructure costs/irrigated ha ('000 Rs/ha)	26.6	14.5	3.4	85.1	222	222
Groundwater share of irrigation	0.658	0.257	0.022	1.000	222	222
Deep tubewells/irrigated ha	0.025	0.057	0.000	0.616	222	222
Shallow tubewells/irrigated ha	0.130	0.213	0.000	1.821	222	222
Dugwells/irrigated ha	0.251	0.401	0.000	2.961	222	222
Well '95-'17						
Depth to water table (mbgl)	14.3	15.4	-1.1	534.0	123,199	203

Notes: Descriptive statistics on the primary datasets are presented here. Units are in parentheses, and standard deviations are omitted for binary variables. Observations in Ag '07-'11 are district-year, observations in NSS '12 are household-plot (for agricultural productivity, area, and irrigated) or household (for agricultural inputs and any bank loan), observations in Irr '07 are district, and observations in Well '95-'17 are well-season. Clusters are districts. To maintain comparability to Ag '07-'11 and Irr '07, statistics for the NSS '12 are calculated weighting using sampling weights times plot area, with weights scaled so each district receives identical weight. Similarly, statistics for Well '95-'17 are weighted so each district-year receives identical weight. All subsequent analysis maintains these weights.

Table 1.2: Instrumental variables estimates

	Share irrigated		Agricultural productivity ('000 Rs/ha)				
	First stage $\left(\beta^{FS}_{(\cdot)}\right)$		Reduced form $\left(\beta^{RF}_{(\cdot)}\right)$		OLS	IV $\left(\beta^{IV}_{(\cdot)} = \frac{\beta^{RF}_{(\cdot)}}{\beta^{FS}_{(\cdot)}}\right)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Z_n (cost instrument)	0.278***		6.3				
	(0.056)		(4.1)				
W_n (outcome instrument)		0.791***		42.9***			
		(0.188)		(10.2)			
D_n (share irrigated)					23.9***	22.6*	54.3***
					(2.8)	(13.1)	(14.5)
Instrument (IV only)	-	-	-	-	-	Z_n	W_n
State FE	X	X	X	X	X	X	X
State FE × X_n	X	X	X	X	X	X	X
State FE × Z_n	-	X	-	X	-	-	X
State FE × W_n	X	-	X	-	-	X	-
# of observations	884	884	884	884	884	884	884
# of clusters	222	222	222	222	222	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses. Regression table contains instrumental variable estimates from Ag '07-'11 using potential aquifer flow Z_n and log relative potential irrigated crop yield W_n as instruments. In each case, the effect of share irrigated on agricultural productivity per hectare is instrumented for. Controls in all specifications include state fixed effects and state fixed effects interacted with log potential rainfed crop yield X_n. The estimated local average surplus effect is the coefficient on share irrigated in Column 7 minus the coefficient on share irrigated in Column 6; estimates of local average surplus effects and pseudo treatment effect elasticities of demand are presented in Table 1.3.

Table 1.3: Local average surplus effect estimates

	Agricultural productivity		(−) Infrastructure costs	
	Ag '07-'11		Irr '07	
	IV	WIV	IV	WIV
	(1)	(2)	(3)	(4)
Z_n				
β_Z^{FS} (first stage)	0.278***	0.245***	0.574***	0.575***
	(0.056)	(0.073)	(0.217)	(0.221)
$\beta_Z^{IV} = \frac{\beta_Z^{RF}}{\beta_Z^{FS}} = \text{LATE}_Z$	22.6*	32.9**	-59.1**	-86.9*
	(13.1)	(15.7)	(25.5)	(47.4)
State FE × W_n	X	X	X	X
W_n				
β_W^{FS} (first stage)	0.791***	0.654***	0.275***	0.258***
	(0.188)	(0.216)	(0.068)	(0.095)
$\beta_W^{IV} = \frac{\beta_W^{RF}}{\beta_W^{FS}} = \text{LASE}_W + \text{LATE}_W$	54.3***	82.7***	28.3	32.6
	(14.5)	(28.5)	(18.2)	(24.3)
State FE × Z_n	X	X	X	X
Surplus effects				
$\beta_W^{IV} - \beta_Z^{IV} \approx \text{LASE}_W$	31.7*	49.8	87.4***	119.6**
	(17.9)	(30.8)	(33.7)	(55.9)
$\frac{\beta_Z^{IV}}{\beta_W^{IV} - \beta_Z^{IV}} \approx$ Treatment effect elasticity of demand	0.715	0.660	-0.676***	-0.727***
	(0.733)	(0.607)	(0.156)	(0.172)
State FE	X	X	X	X
State FE × X_n	X	X	X	X
LASE: p-value [pairs bootstrap-c p-value]	0.077 [0.136]	0.106 [0.144]	0.010 [0.052]	0.033 [0.072]
$(Z_n, W_n) = (\text{Aquifer yield}_n, \text{Irr. crop yield}_n)$	X	X	-	-
$(Z_n, W_n) = (\text{Irr. crop yield}_n, \text{Aquifer yield}_n)$	-	-	X	X
# of observations	884	884	222	222
# of clusters	222	222	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses, and each cell reports a coefficient from a separate regression. Estimates from Columns 1 and 2 are directly comparable, while the relative interpretation of estimates from Columns 3 and 4 is discussed in Section 1.4.2 and 1.5.1. Rows 1 and 3 report first stage coefficients with irrigated share of agricultural land D_n as the dependent variable. Rows 2 and 4 report instrumental variable estimates with gross revenue (for Columns 1 and 2) or negative fixed costs of irrigation infrastructure (for Columns 3 and 4) as the dependent variable ('000 Rs/ha). Row 5 reports estimates of the local average surplus effect, and Row 6 reports estimates of a pseudo treatment effect elasticity of demand. Estimators in Columns 2 and 4 are weighted to balance the share of compliers in each state across β_Z^{IV} and β_W^{IV} as discussed in Section 1.3.5. All specifications include as controls state fixed effects and state fixed effects interacted with log potential rainfed crop yield X_n. The instrument Z_n is potential aquifer yield in Columns 1 and 2 and log relative potential irrigated crop yield in Columns 3 and 4, and the instrument W_n is log relative potential irrigated crop yield in Columns 1 and 2 and potential aquifer yield in Column 3 and 4. Pairs bootstrap-c p-values for estimates of local average surplus effects are calculated following Young (2018).

Table 1.4: LASE robustness, NSS

	Agricultural productivity			Agricultural inputs
	Ag '07-'11	NSS '12		
	IV	IV	IV (CF predictions)	IV
	(1)	(2)	(3)	(4)
Z_n				
β_Z^{FS}	0.278***	0.289***	0.257	0.375***
	(0.056)	(0.073)	(0.056)	(0.078)
β_Z^{IV}	22.6*	37.5***	13.4	12.0*
	(13.1)	(13.9)	(13.5)	(6.6)
State FE × W_n	X	X	X	X
W_n				
β_W^{FS}	0.791***	0.834***	0.881	0.852***
	(0.188)	(0.226)	(0.218)	(0.214)
β_W^{IV}	54.3***	67.9***	10.1 + 54.3 + 4.6	8.6
	(14.5)	(23.9)	$\underbrace{(10.7)}_{\text{LATE}_W} + \underbrace{(20.0)}_{\text{LASE}_W} + \underbrace{(19.3)}_{\text{bias}_W}$	(7.3)
State FE × Z_n	X	X	X	X
Surplus effects				
$\beta_W^{IV} - \beta_Z^{IV}$	31.7*	30.4	55.5	-3.3
	(17.9)	(26.3)	(19.7)	(9.7)
State FE	X	X	X	X
State FE × X_n	X	X	X	X
# of observations	884	33,778	33,778	26,280
# of clusters	222	222	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses, and each cell reports a coefficient from a separate regression. Rows 1 and 3 report first stage coefficients with irrigated share of agricultural land D_n as the dependent variable. Rows 2 and 4 report instrumental variable estimates with gross revenue (for columns 1, 2, and 3) or expenditures on agricultural inputs net of irrigation (for column 4) as the dependent variable ('000 Rs/ha). Row 5 reports estimates of the local average surplus effect. Estimators in Columns 2, 3, and 4 are weighted using sample weights times plot area, with weights scaled so each district receives identical weight. Column 3 uses control function predicted outcomes and propensity scores as outcomes in the reduced form and first stage, respectively. This allows decomposition of β_W^{IV} into a LATE, a LASE, and bias from violations of the exclusion restriction $W_n \not\perp Y_{0i}$, which is identified using the control function approach. All specifications include as controls state fixed effects and state fixed effects interacted with log potential rainfed crop yield X_n. The instrument Z_n is potential aquifer yield, and the instrument W_n is log relative potential irrigated crop yield.

Table 1.5: Control function estimates

g_C	-34.1 (13.6)**
c_0	4.0 (17.0)
g_Y	78.6 (27.5)***
σ_V	25.8 (9.5)***
$\frac{\text{Cov}(-V_{1i}, V_i - \mathbf{E}[V_i \mid X_i])}{\sigma_V^2}$	0.21 (0.44)
$\frac{\text{Cov}(V_{0i}, V_i - \mathbf{E}[V_i \mid X_i])}{\sigma_V^2}$	0.11 (0.24)
$\frac{\text{Cov}(V_{Ci}, V_i - \mathbf{E}[V_i \mid X_i])}{\sigma_V^2}$	0.68 (0.46)
# of observations	33778
# of clusters	222

Notes: Robust standard errors clustered at the district level are used to construct 95% confidence intervals in square brackets. Parameters are estimated by a two step control function approach as detailed in Section 1.3.5 and A.2.3, and standard errors are adjusted for the two step procedure. g_C is the effect of the cost instrument Z_{Cn} (potential aquifer yield) on cost per hectare of irrigation, g_Y and c_0 are the effects of the outcome instrument Z_{Yn} (log relative potential irrigated crop yield) on relative revenue per hectare from irrigation and revenue per hectare from rainfed agriculture, respectively. σ_V is the standard deviation of idiosyncratic relative profitability of irrigated agriculture. The three covariance terms decompose the variance of idiosyncratic relative profitability of irrigated agriculture into components from idiosyncratic revenue from irrigated agriculture, idiosyncratic revenue from rainfed agriculture, and idiosyncratic costs of irrigated agriculture, respectively.

Table 1.6: Informativeness of IV estimators for CF predicted LATE and LASE

Descriptive Statistic (IV estimator)	Estimate of interest (CF prediction)	Informativeness
β_Z^{IV}	LATE_Z	0.506
$\beta_W^{IV} - \beta_Z^{IV}$	LASE_W	0.118
β_Z^{WIV}	LATE_Z^{WIV}	0.455
$\beta_W^{WIV} - \beta_Z^{WIV}$	LASE_W^{WIV}	0.504

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The informativeness of 4 IV estimators for their target parameters estimated using a control function approach are presented here. Informativeness is calculated following Andrews et al. (2018), who note that it can be interpreted as the R^2 from the population regression of the target parameter on the corresponding IV estimator in their joint asymptotic distribution. IV estimators β_Z^{IV} and β_W^{IV} use Z (potential aquifer flow) and W (log relative potential irrigated crop yield) as instruments, respectively, for the effect of D (irrigation) on Y (gross revenue per hectare). CF predictions replace Y and D with their predictions using a two step control function approach following Kline & Walters (2017). LATE comparisons control for state FE, W (Z), and state FE interacted with X for $\beta_Z^{(\cdot)}$ ($\beta_W^{(\cdot)}$), and LASE comparisons control for state FE, state FE interacted with W (Z), and state FE interacted with X for $\beta_Z^{(\cdot)}$ ($\beta_W^{(\cdot)}$). WIV estimators use weights to balance compliers on state FE, with weights constructed as described in Section 1.3.5. Cluster robust variance covariance matrices are estimated clustered at the district level.

Table 1.7: Calibrated parameters

	Value [Low, High]	Source
Calibrated parameters		
$\epsilon_{A,p}$, $\epsilon_{M,p}$	-0.18	Badiani & Jessoe (2017)
r, upper bound (rural credit interest rate)	0.20	Hussam et al. (2017)
r, lower bound (India 30 year bond yield)	0.08	
m (energy/m^3 of water/m)	6.8 Wh/m^3/m	Shah (2009)
$d\mathbf{E}[D_i]/db$	-.0024/m	Fishman et al. (2017)
Calibrated parameters (Rajasthan)		
p	1.21 Rs/kWh	Fishman et al. (2016)
c	3.30 Rs/kWh	Fishman et al. (2016)
b (depth to water table)	[5m, 66m]	Well '95-'17
α (specific yield)	[0.015, 0.068]	Narain et al. (2006)
$\mathbf{E}[D_i]L/\overline{L}$ (aquifer share irrigated)	[0.015, 0.492]	Ag '07-'11
$A/\mathbf{E}[D_i]L$ (groundwater use/irrigated ha)	[0.065, 0.650] m ha/ha	Ag '07-'11
Rajasthan 2008 agricultural electricity use	9,791 GWh	Rajasthan DES (2011)
India statistics		
$\mathbf{E}[D_i]L$ (irrigated ha)	60 million ha	Ag '07-'11
$A/\mathbf{E}[D_i]L$ (avg. groundwater use/irrigated ha)	0.43 m ha/ha	Shah (2009), Ag '07-'11
p	1.05 Rs/kWh	Fishman et al. (2016)
$pM/\mathbf{E}[D_i]L$ (avg. elec. exp./irrigated ha)	1,470 Rs/ha	Fishman et al. (2016), Ag '07-'11
Estimates		
MSE(u)	71,500u Rs/ha	Section 1.5.3
$\epsilon_{A,p}$, $\epsilon_{M,p}$ (lower bound)	-0.045	Section 1.7.2

Notes: This table contains the calibrated parameters for the counterfactual exercises in Section 1.5.4 and Section 1.7. Values are provided as points when a single estimate is used, and as a range when the value used is allowed to vary across districts. Ranges for specific yield, depth to water table, aquifer share irrigated, and groundwater use/irrigated ha are specific to Rajasthan. Depth to water table is estimated as the median post monsoon Kharif reading from the network of monitoring tube wells, bottom winsorized at 5m.

Table 1.8: Lost surplus from groundwater depletion

	1m decline Rs/irrigated ha (1)	3.3m decline, NW India Rs/ha [% of productivity/ha] (2)
IV		
LASE		251 [0.80%]
MSE	132	282 [0.90%]
Weighted IV		
LASE		394 [1.26%]
MSE	172	365 [1.16%]
Control Function		
LASE		430 [1.37%]
Back of envelope		
3x Electricity costs	93	197 [0.63%]
6x Electricity costs	186	395 [1.26%]

Notes: This table presents estimates of the lost surplus from groundwater depletion using estimates of local average surplus effects and marginal surplus effects from Section 1.5.1 and 1.5.3, and calibrated parameters from Table 1.7. Column 1 presents the impact of a 1m decline in the water table on costs per irrigated hectare. Column 1 IV and WIV estimates are calculated using the estimated marginal surplus effect, and the calibrated effect of a 1m decline in water tables on adoption of irrigation. Column 1 back of the envelope approaches calculate the increased electricity costs farmers would have to pay to pump groundwater one additional meter, exclusively using calibrated parameters from Table 1.7. Column 2 presents the impact of a 3.3m decline in water tables in Northwestern India (Haryana, Punjab, and Rajasthan), the estimate of 2000's water table declines from Rodell et al. (2009).

Table 1.9: Irrigation technology

	Irr '07					Well '95-'17	Well '07-'11
	Groundwater ha/ha	Surface water ha/ha	Deep tubwell/ha	Shallow tubewell/ha	Dugwell/ha	Depletion (mbgl/year)	Depth to water table (mbgl)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
D_n (share irrigated)	0.718***	0.282	0.091*	0.067	-0.019	2.14	18.1
	(0.222)	(0.222)	(0.053)	(0.078)	(0.127)	(1.39)	(19.1)
Instrument	W_n	W_n	W_n	W_n	W_n	W_n	W_n
State FE	X	X	X	X	X	X	X
State FE $\times X_n$	X	X	X	X	X	X	X
State FE $\times Z_n$	X	X	X	X	X	X	X
# of observations	222	222	222	222	222	85,804	28,169
# of clusters	222	222	222	222	222	198	176

	Irr '07					Well '95-'17	Well '07-'11
	Groundwater ha/ha	Surface water ha/ha	Deep tubwell/ha	Shallow tubewell/ha	Dugwell/ha	Depletion (mbgl/year)	Depth to water table (mbgl)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
D_n (share irrigated)	1.000***	0.000	-0.002	0.325***	-0.395**	3.75	67.5
	(0.207)	(0.207)	(0.031)	(0.096)	(0.155)	(2.59)	(41.4)
Instrument	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n
State FE	X	X	X	X	X	X	X
State FE $\times X_n$	X	X	X	X	X	X	X
State FE $\times W_n$	X	X	X	X	X	X	X
# of observations	222	222	222	222	222	85,804	28,169
# of clusters	222	222	222	222	222	198	176

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses. In the first subtable, coefficients on share irrigated are estimated using W_n, log relative potential irrigated crop yield, as an instrument. In the second subtable, coefficients on share irrigated are estimated using Z_n, potential aquifer flow, as an instrument. Controls X_n are log potential rainfed crop yield.

Table 1.10: LASE robustness, controls

	Agricultural productivity (Y_n)					
	Ag '07-'11					
	IV	IV	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Z_n						
β_Z^{FS}	0.278***	0.465***	0.229***	0.239***	0.310***	0.394***
	(0.056)	(0.035)	(0.053)	(0.058)	(0.060)	(0.072)
β_Z^{IV}	22.6*	22.3***	26.3**	17.7	34.8***	36.1***
	(13.1)	(4.6)	(12.4)	(15.1)	(11.3)	(10.0)
W_n	X	-	X	X	X	X
State FE × W_n	X	-	-	-	X	X
State FE × X_nW_n	-	-	-	-	-	X
State FE × W_n^2	-	-	-	-	-	X
W_n						
β_W^{FS}	0.791***	0.302***	0.522***	0.756***	0.502**	0.400*
	(0.188)	(0.101)	(0.187)	(0.187)	(0.220)	(0.222)
β_W^{IV}	54.3***	26.3***	83.2***	57.4***	76.6**	84.3*
	(14.5)	(8.6)	(27.1)	(15.5)	(33.7)	(45.4)
Z_n	X	-	X	X	X	X
State FE × Z_n	X	-	-	-	X	X
State FE × X_nZ_n	-	-	-	-	-	X
State FE × Z_n^2	-	-	-	-	-	X
Surplus effects						
$\beta_W^{IV} - \beta_Z^{IV}$	31.7*	4.0	56.9**	39.6*	41.9	48.2
	(17.9)	(9.6)	(28.3)	(20.6)	(35.1)	(47.0)
X_n	X	-	X	X	X	X
State FE	X	-	X	X	X	X
State FE × X_n	X	-	-	X	X	X
State FE × X_n^2	-	-	-	-	X	X
# of observations	884	884	884	884	884	884
# of clusters	222	222	222	222	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses, and each cell reports a coefficient from a separate regression. Rows 1 and 3 report first stage coefficients with irrigated share of agricultural land D_n as the dependent variable. Rows 2 and 4 report instrumental variable estimates with gross revenue as the dependent variable ('000 Rs/ha). The control X_n is log potential rainfed crop yield, the instrument Z_n is potential aquifer yield, and the instrument W_n is log relative potential irrigated crop yield.

Table 1.11: Placebo before Green Revolution

	D_{nt} (share irrigated)	$\log Y_{nt}$ (log agricultural productivity)
	(1)	(2)
Z_n	0.050	0.028
	(0.063)	(0.103)
$1\{t > 1966\}Z_n$	0.116***	0.120*
	(0.042)	(0.067)
W_n	0.182	0.755*
	(0.144)	(0.411)
$1\{t > 1966\}W_n$	0.335***	0.781***
	(0.114)	(0.240)
$\log \text{RF yield}_n$	0.144*	1.150***
	(0.085)	(0.260)
$1\{t > 1966\} \log \text{RF yield}_n$	0.193***	0.200
	(0.074)	(0.154)
State-by-year FE	X	X
# of observations	11,799	11,799
# of clusters	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses. Z_n is potential aquifer yield, W_n is log relative potential irrigated crop yield, and RF yield$_n$ is log potential rainfed crop yield. Outcomes are from Ag '56-'11.

Table 1.12: Irrigation and crop choice

		$D_i 1\{\text{Crop}_i = (.)\}$												
	D_i (irrigated)	Wheat	Rice	Cotton	Soya	Bajra	Gram	Maize	Jowar	Sugar	RM	Tur	Groundnut	Potato
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Z_n	0.289***													
	(0.073)													
D_i (irrigated)		0.314	0.983***	-0.160	0.094	0.001	-0.249**	0.039	0.012	-0.060	0.028	0.009	-0.028	0.016
		(0.219)	(0.297)	(0.108)	(0.083)	(0.143)	(0.119)	(0.116)	(0.030)	(0.107)	(0.070)	(0.012)	(0.024)	(0.045)
Instrument (IV only)	-	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n
BH q-value	-	0.495	0.012	0.495	0.558	0.996	0.236	0.796	0.796	0.796	0.796	0.796	0.558	0.796
State FE	X	X	X	X	X	X	X	X	X	X	X	X	X	X
State FE $\times X_n$	X	X	X	X	X	X	X	X	X	X	X	X	X	X
State FE $\times W_n$	X	X	X	X	X	X	X	X	X	X	X	X	X	X
# of observations	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778
# of clusters	222	222	222	222	222	222	222	222	222	222	222	222	222	222

		$(1 - D_i)1\{\text{Crop}_i = (.)\}$												
	D_i (irrigated)	Wheat	Rice	Cotton	Soya	Bajra	Gram	Maize	Jowar	Sugar	RM	Tur	Groundnut	Potato
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Z_n	0.289***													
	(0.073)													
D_i (irrigated)		-0.087**	-0.513***	-0.053	-0.220**	0.158	-0.081	-0.187***	-0.027	0.000	-0.021	0.058	-0.027	0.000
		(0.038)	(0.148)	(0.096)	(0.107)	(0.097)	(0.057)	(0.072)	(0.045)	(0.000)	(0.027)	(0.083)	(0.054)	(0.000)
Instrument (IV only)	-	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n	Z_n
BH q-value	-	0.096	0.007	0.674	0.127	0.269	0.347	0.060	0.674	0.871	0.674	0.674	0.674	0.419
State FE	X	X	X	X	X	X	X	X	X	X	X	X	X	X
State FE $\times X_n$	X	X	X	X	X	X	X	X	X	X	X	X	X	X
State FE $\times W_n$	X	X	X	X	X	X	X	X	X	X	X	X	X	X
# of observations	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778	33,778
# of clusters	222	222	222	222	222	222	222	222	222	222	222	222	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses. Coefficients on share irrigated are estimated using Z_n, potential aquifer yield, as an instrument. D_i is an irrigation indicator for plot i, and Crop_i is the crop cultivated on plot i. Controls X_n are log potential rainfed crop yield, and W_n is log relative potential irrigated crop yield. D_i and Crop_i are from NSS '12. Following Benjamini & Hochberg (1995) and Anderson (2008), BH p-value are multiple inference adjusted p-values (adjusted within table).

Table 1.13: LASE robustness, surface water and endogenous cultivation

	Agricultural productivity (Y_n)		$Y_n L_n/\overline{L}_n$	$(Y_n L_n + 20(\overline{L}_n - L_n))/\overline{L}_n$
	Ag '07-'11			
	IV	IV	IV	IV
	(1)	(2)	(3)	(4)
Z_n				
β_Z^{FS}	0.278***	0.279***	0.456***	0.456***
	(0.056)	(0.054)	(0.075)	(0.075)
β_Z^{IV}	22.6*	15.1	34.5***	17.2***
	(13.1)	(12.9)	(5.6)	(5.7)
State FE × W_n	X	X	X	X
W_n				
β_W^{FS}	0.791***	0.777***	0.559**	0.559**
	(0.188)	(0.227)	(0.241)	(0.241)
β_W^{IV}	54.3***	64.3***	55.9***	39.5**
	(14.5)	(19.3)	(17.5)	(18.4)
State FE × Z_n	X	X	X	X
Surplus effects				
$\beta_W^{IV} - \beta_Z^{IV}$	31.7*	49.2**	21.4	22.4
	(17.9)	(20.7)	(17.4)	(18.0)
State FE	X	X	X	X
State FE × X_n	X	X	X	X
GJ, HA+PJ, MH, RJ, UP	-	X	-	-
Endog. $D_n L_n/\overline{L}_n$	-	-	X	X
# of observations	884	447	884	884
# of clusters	222	133	222	222

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses, and each cell reports a coefficient from a separate regression. Rows 1 and 3 report first stage coefficients with irrigated share of agricultural land D_n as the dependent variable. Rows 2 and 4 report instrumental variable estimates with gross revenue as the dependent variable ('000 Rs/ha). The control X_n is log potential rainfed crop yield, the instrument Z_n is potential aquifer yield, and the instrument W_n is log relative potential irrigated crop yield. Column 2 restricts observations to districts in the five 1961 states with the smallest shares of surface water irrigation. Columns 3 and 4 use share of district land irrigated, instead of share of district agricultural land irrigated, as treatment D_n. Columns 3 and 4 use agricultural production plus a reservation rent for uncultivated land (0 in Column 7 and 20,000 Rs/ha in Column 8) per hectare of district land as the outcome, instead of agricultural revenue per cultivated hectare.

Table 1.14: Optimal electricity taxes in Rajasthan

	$\epsilon = 0.18$, $r = 0.2$		$\epsilon = 0.045$, $r = 0.2$		$\epsilon = 0.18$, $r = 0.08$	
	Status quo (1)	Optimal (2)	Status quo (3)	Optimal (4)	Status quo (5)	Optimal (6)
λ (implied by status quo)	1.56	1.56	1.12	1.12	2.13	2.13
	Billion Rs [% of agricultural production]					
Total subsidy	10.32 [6.59%]	10.32 [6.59%]	10.32 [6.59%]	10.32 [6.59%]	10.32 [6.59%]	10.32 [6.59%]
Deadweight loss	1.02 [0.65%]	1.09 [0.69%]	0.27 [0.17%]	0.29 [0.18%]	1.02 [0.65%]	1.20 [0.76%]
Externality (utility)	0.10 [0.06%]	0.07 [0.05%]	0.03 [0.02%]	0.02 [0.01%]	0.25 [0.16%]	0.17 [0.11%]
Farmer surplus						
Subsidy	9.31 [5.94%]	9.24 [5.90%]	10.06 [6.42%]	10.04 [6.41%]	9.31 [5.94%]	9.13 [5.83%]
Externality (farmer)	0.70 [0.45%]	0.52 [0.33%]	0.20 [0.13%]	0.16 [0.10%]	1.75 [1.12%]	1.12 [0.71%]
Total	8.61 [5.50%]	8.72 [5.57%]	9.86 [6.30%]	9.88 [6.31%]	7.56 [4.83%]	8.01 [5.12%]
	m/decade [% of 2000-2010 decline]					
Water table decline	1.51 [45.7%]	1.26 [38.3%]	0.40 [12.2%]	0.34 [10.4%]	1.51 [45.7%]	1.19 [36.1%]

Notes: This table presents the results of the optimal policy exercise. Columns 1, 3, and 5 present results from maintaining the status quo ($p = 1.21$ Rs/ha in all districts, with marginal cost $c = 3.30$ Rs/ha). Columns 2, 4, and 6 present results from optimal subsidies holding fixed total subsidies. ϵ is the calibrated elasticity of groundwater extraction/electricity use to the price of electricity, and r is the calibrated discount rate. λ is the inverse marginal value of public funds for a marginal change to state level subsidies under the status quo. All cells report impacts of the policy relative to no subsidies.

Chapter 2

Irrigation in Rwanda: Farmers' Responses to a Massive Expansion of the Production Possibility Frontier

2.1 Introduction

Agricultural productivity growth in sub-Saharan Africa has lagged severely compared to the rest of the world. Diagnostically, a key difference between farmers in sub-Saharan Africa and elsewhere in the world is a low use of modern inputs (World Bank, 2007). One explanation for relatively low growth in productivity is that farmers' productive decisions are constrained, preventing farmers from realizing the production possibility frontier (Udry, 1996a). A large and recent literature has used field experiments to examine the effects of relaxing some of these constraints on technology adoption, with particular focuses on credit (Giné & Yang, 2009; Carter et al., 2013; Beaman et al., 2014; Karlan et al., 2014b; Crépon et al., 2015; Tarozzi et al., 2015), risk (Cai et al., 2015; Cole et al., 2013; Emerick et al., 2016; Karlan et al., 2014b), and information (BenYishay & Mobarak, 2014; Beaman et al., 2018; Conley & Udry, 2010; Kondylis et al., 2017; Cole & Fernando, 2016). These studies have produced robust evidence that experimentally generating slack in either credit, risk, or information constraints leads to higher levels of technology adoption and often improved yields. Yet, these gains typically accrue to a minority of farmers, which suggests that none of these constraints can unilaterally explain the agricultural productivity gap. Two competing hypotheses may explain these results. First, the candidate new technologies available to farmers in these studies may have fundamentally low returns. For example, it has been proposed that the returns to modern input use may be lower in the African context due to, among other reasons, counterfeit inputs (Bold et al., 2017) and soil depletion (Barrett & Bevis, 2015). If slack is generated experimentally in various constraints in environments where the potential returns from technology adoption are heterogeneous and/or modest on average, it may explain these adoption results. Alternatively, farmers may be subject to a variety of heterogeneous and competing constraints, including failures in land and labor markets where the recent experimental evidence base has made less progress.[1] In such a world, generating slack in one constraint only advances productive potential up until the next constraint binds. Thus, while the evidence base on constraints to technology adoption has matured substantively in recent years, there remains much we do not understand about low agricultural productivity and the low use of modern inputs in Africa.

An additional difference between agricultural productivity growth in sub-Saharan Africa and the rest of the world is the presence of irrigation. As of 2015, only 3.3% of arable land in Africa was irrigated, compared to 50.5% and 63.2% in South and East Asia, respectively (FAO, 2019). Irrigation fundamentally changes the production possibility frontier in several ways: it adds additional agricultural seasons; reduces the risk of poor rainfall realizations, and allows cultivation of high value crops which require water use and water control. From a technology adoption perspective, we can thus view irrigation as not only a directly productive technology with its own adoption decision, but also one which is complementary to a variety of other technological adoption decisions both due to technical complementarities and to a reduction in risk. The empirical evidence base, largely from India, suggests that irrigation greatly increases agricultural productivity and farmer welfare (Duflo & Pande, 2007; Dillon, 2011; Sekhri, 2014; **?**). Irrigation expanded rapidly in South Asia over the past half century, which together with these positive results suggests that it is likely to be an important driver of the productivity growth experienced by that region.

At the same time, we know much less about how the gains to irrigation are mediated by farmer constraints, which may meaningfully influence whether we should expect irrigation to be similarly productive in SSA. This gap is due to two factors. First, cross-sectional comparisons are challenged by an extraordinarily limited evidence base on irrigation in SSA itself.[2] Second, the aggregate nature of plausibly exogeneous variation

[1] It has been well documented that weak property rights constrain productive investments, both theoretically and empirically (Besley, 1995; Goldstein & Udry, 2008; Besley & Ghatak, 2010). Consistent with this, a growing body of work has found that programs to establish secure property rights, through land titling or demarcation, can encourage productive investments (Deininger & Feder, 2009; Goldstein et al., 2018), including in Rwanda (Ali et al., 2014). However, quasi-experimental work in the United States has found evidence that agricultural land market may fail even in settings with well established property rights (Bleakley & Ferrie, 2014).

[2] This is likely attributable to the shortage of irrigation schemes in SSA. A literature review conducted in March 2019 identified only one study using quantitative methods to examine the effects of irrigation in SSA (Dillon, 2011); that study uses PSM techniques to examine the returns to irrigation in Mali

in irrigation access (for example, slope characteristics of river basins (Duflo & Pande, 2007) or aquifer characteristics (Sekhri, 2014)) has limited our ability to separate the direct effects of irrigation access from general equilibrium effects induced by this access. A consequence of this distinction is that we have learned little about how constraints to farmer production influence the returns to irrigation. If constraints like credit, risk, information, or factor market failures limit adoption of productive technologies in SSA, then they may also both limit irrigation use and the returns to irrigation which could be realized with large-scale adoption of complementary technologies.

In this paper, we study the impacts of irrigation and constraints to irrigation adoption in Rwanda. Our study takes two steps. First, we establish that irrigation is profitable. We leverage the technical characteristics of several new gravity-fed hillside irrigation schemes to estimate the reduced form impacts of irrigation on production decisions and profits. Specifically, these schemes share some features in common: a main canal was constructed along a contour of the hillside according to engineering specifications, and plots below this canal receive water access while plots above this canal do not. We collected 4 years of data on 3,000 plots in very close proximity to this canal to allow estimates of the impacts of irrigation on plot and farm-level outcomes using a very granular spatial regression discontinuity design.

Once we have identified the productive potential of irrigation in this context, we exploit our granular spatial discontinuity in access to irrigation to test for whether factor market failures constrain adoption of irrigation in Rwanda. We extend the model in Benjamin (1992) to show that absent market failures, access to irrigation on one plot should not change input decisions on other plots. In practice, we examine how farmers who have plots that receive access to irrigation change input decisions on their *other* plots, compared to farmers with plots close to the boundary who do not receive access to irrigation. The plot-specific nature of the irrigation shock allows us to examine how an increase in input demand on one plot affects productivity on other plots.

We find the following. Two to four years after the construction of irrigation, our treatment on the treated (TOT) estimates suggest that irrigated plots are 70% more productive than plots within a few meters which do not receive access. This increase in production is entirely attributable to a change in cultivation practices during the dry season: farmers substitute away from bananas, a low productivity perennial, and towards cultivating horticulture. While rainy season inputs and yields from different cropping choices are comparable in magnitude, dry season horticulture represents a large increase in input demands and a large increase in output relative to banana cultivation or to leaving the plot fallow in the dry season. Given that the dry season is only 3 months long, we find that irrigation allows a 70% increase in yields and cash profits, achieved in only 1/4 of the year. The effect on economic profits - cash profits net of household labor costs - depends on the shadow wage faced by household workers, who provide a large majority of the labor to these farms. Household labor is notoriously difficult to value in rural settings. In the two extremes, if we value this labor at zero we can therefore interpret the cash profit estimate as a 70% increase in economic profits; conversely if we value it at the market wage the net effect on economic profits is close to zero because irrigation causes a substantial increase in labor demand.

Despite this large change in productive potential, adoption is only partial. Two to four years after the system was brought online, dry season horticultural cultivation has been roughly constant at about 25% of available plots. At this level of adoption, the sustainability of the system is in doubt: even the large gains in revenues to adopters are unable to generate enough surplus to pay for routine maintenance costs with only 25% of farmers benefiting. To understand why adoption may be so low, we turn to the literature on separation in agricultural households which identifies the presence of frictions in rural factor markets (Singh et al., 1986; Benjamin, 1992; Udry, 1996b; LaFave & Thomas, 2016; Dillon & Barrett, 2017). In particular, we note that profitability depends crucially on the shadow wage, which itself suggests that frictions in agricultural labor markets have the potential to constrain adoption of dry season horticulture.[3]

We leverage the plausibly exogenous variation in our data to not only generate a new test for the presence of factor market failures but also to document directly the implications of these factor market failures for technology adoption, production, and ultimately the economic viability of these irrigation schemes. More specifically, irrigation represents a large increase in the production possibility frontier for farmers that is tied to a particular plot and not fungible across the household's plots. That increase is associated with a large increase in labor and input demands, as well as a potential change in the risk profile of cultivation. Intuitively, if a constraint on inputs or labor binds for these farmers, then when the demand for that input increases on irrigated plots, it should induce a budget allocation away from other plots. To formalize this intuition, we extend the agricultural household model presented in Benjamin (1992) to decision-making when farmers have access to multiple plots. The model suggests that an increase in labor and input demands on a newly-irrigated plot may constrain adoption of irrigation on other plots due to failures in insurance, labor, or input markets.

We then test this model using data on farmers' other plots. More specifically, we test whether farmers who have a plot just inside the irrigated area reduce input use on their other plots compared to farmers who have plots just outside the irrigated area. We find robust evidence that farmers change their input use and crop choice on other plots. That is, farmers inside the irrigated area sharply reduce their labor allocations, reduce purchased inputs, and cultivate bananas rather than horticulture on other plots, producing lower revenues. Together, these two findings suggest a clear inefficiency and failure in separation: efficient farming on a plot should be independent of a farmer's

[3]Labor constraints have been much less extensively studied in the recent RCT-based literature on technology adoption, in part due to challenges in designing interventions which would generate slack in these constraints.

practices on his other plots. We conclude that separation failures lower the profitability of the irrigation system.

Using our model, we generate tests to examine whether our findings can be explained solely by risk preferences or may be attributable to farmer constraints. We demonstrate that a household with two additional members experiences a 62-86% smaller negative adoption effect than an average size household. A common concern (also suggested by the model) is that larger households may be wealthier which may also lead to differential input use choices. In this context, however, a one standard deviation wealthier household who receives access to irrigation experiences a 40-80% larger negative adoption effect than an average wealth household, which is also consistent with the agricultural household model. These two trends together indicate that constraints in the labor market fundamentally reduce the potential for irrigation in Rwanda.

This paper is organized as follows. Section 2.2 describes the context we study and our sources of data. Section 2.3 presents our estimates of the impacts of irrigation in Rwanda. Section 2.4 presents our model of adoption of irrigation in the presence of separation failures, and we implement tests of separation failures and binding constraints suggested by the model in Section 2.5. Section 2.6 concludes.

2.2 Data and context

2.2.1 Irrigation in Rwanda

We study 3 hillside irrigation schemes, located in Karongi and Nyazna districts of Rwanda, that were constructed by the government in 2014; a timeline of construction and our surveys is presented in Figure 2.1. 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.[4] Absent irrigation, therefore, most farmers adopt either a rotation of staples, fallowing land in the dry season, or cultivate bananas.

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

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 seasonal 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.[5] 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 not unfamiliar but uncommon around these areas; at baseline 3.2% of plots outside of the command area are planted with at least some 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 follow 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 2.2. In each site, land was terraced in preparation for the irrigation works (as hillside irrigation would be unfeasible 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.[6] From these water sources, a main canal (visible in Figure 2.2) 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.[7] 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 water[8]

[5]Kigali is less than a 3 hour drive from these markets, facilitating trade.

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

[7]The lower pressure on these plots is attributable to the design of the pipes, which fill up with water before valves are opened; this means that pressure on the highest valves can fall when lower valves are opened. In practice, schedules of water usage are agreed upon to prevent this from happening.

[8]The government does have an objective of developing the financial self-sufficiency of the schemes. To do so, land taxes are intended to be applied to the plots in the command area, which (as land taxes)

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.

2.2.2 Data

Aerial sampling

Spatial discontinuity designs require a high density of observations near the boundary for statistical power. In our context, this required constructing a sample which featured a high density of plots very near the command area boundary. We used point-based sampling to over-sample plots cultivated just inside and just outside the command area. In practice, we constructed this aerial sample of plots by dropping a uniform grid of points across the site at 2-meter resolution, and then sampling points within the grid. After each point was sampled, we excluded any points within 10m of that that point (to keep from selecting multiple points too close together). In two of the three sites, there is a viable boundary of cultivable land both just inside and just outside the command area. In these sites, to guarantee a high density of observations near the canal, we over sampled points that were within 50m of command area boundary, both inside the command area and outside the command area.[9]

Enumerators were then given GPS devices with the locations of the points, and sent to 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,

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 ToT estimates presented in Section 2.3 are 397,000 - 411,000 RwF/ha. The first attempts to collect these taxes were made in 2017 Rainy 1. The survey team engaged in an experiment, described in ***work in progress***, to test whether these taxes were a barrier to use of the irrigation system by randomizing subsidies across farmers at up to 100%; we failed to find any evidence that the taxes changed farming practices (results available from authors). This is perhaps unsurprising as, based on the original schedule, tax compliance was very low, with less than 20% of taxes collected from farmers who did not receive full subsidies from the research team.

[9] In both sites, we additionally sampled some points further from the canal inside the command area (at a lower rate). We use these points along with data from the third site primarily to examine experimental treatments described below.

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 aerial 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, one of the points that fell on their plots was randomly selected to be that household's sample plot.

Survey

Our baseline survey survey was implemented in May 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 2.1. 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 2.3.2 we highlight that 2014 Dry adoption of irrigation is less than 25% of adoption in subsequent dry seasons, and in Section 2.3.1 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 the irrigation impacts household productive decisions more broadly.

Work in progress: Construction of key agricultural variables.

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, 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.

2.2.3 Stylized facts

To motivate our analysis of the impacts of hillside irrigation, we first introduce some stylized facts about irrigation in this context. Table 2.1 presents summary statistics for agricultural production from our four years of data, pooled across seasons.

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 wet season, and almost never use irrigation when cultivating staples or bananas (only 2% of plots cultivated with staples or bananas use irrigation in our data). In contrast, 92% of farmers who cultivate 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 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 460 days/ha of household labor, 60 days/ha of hired labor, and 46,000 RwF/ha of inputs, regardless of the season in which it is planted.[10] This contrasts to staple plots (300 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

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

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., 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).[11] As a result, ultimately the economic profitability of horticulture relative to bananas will depend critically on the constraints on household labor supply decisions.

2.3 Impacts of irrigation

2.3.1 Empirical strategy

For our benchmark approach, we estimate the plot level effects of access to irrigation using a regression discontinuity, comparing sample plots that are just inside of the command areas of the hillside irrigation schemes we study to sample plots that are just outside of the command areas. Specifically, we pool our data across time periods and estimate

$$y_{1ist} = \beta_0 + \beta_1 \text{CA}_{1is} + \beta_2 \text{Dist}_{1is} + \beta_3 \text{Dist}_{1is} * \text{CA}_{1is} + \alpha_{st} + X'_{1is}\gamma + \epsilon_{1ist} \tag{2.1}$$

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, Dist_{kis} is the distance of that plot from the command area boundary (positive for plots within the command area, negative for plots outside the command area), X_{kis} is a vector of controls that includes log area, and α_{st} are site-by-season fixed effects. We use $k = 1$ to indicate the household's sample plot, as opposed to the household's most important plot, and we restrict this and subsequent analysis to using sample plots within 50 meters of the discontinuity, consistent with our sampling strategy.

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

This approach is an example of a boundary spatial discontinuity design; as such it has the large advantage of being able to represent outcomes graphically. However, boundary designs need not represent spatially proximate comparisons if the density of observations within the bandwidth used for the analysis is not constant along the boundary. In other words, the identification assumption underlying a boundary spatial discontinuity design is $\lim_{\text{Dist}_{kist}\uparrow 0} E[\epsilon_{kist}|X_{kist}] = \lim_{\text{Dist}_{kist}\downarrow 0} E[\epsilon_{kist}|X_{kist}]$. This suggests that unobservable variation is constant among observations that are similar in being close to the boundary. However, this is somewhat distinct from assuming that unobservable variation is similar among *nearby* plots. That assumption would suggest that $\lim_{||k-k'||\to 0} E[\epsilon_{kist}|X_{kist}] = E[\epsilon_{k'i'st}|X_{k'i'st}]$ if $||k-k'||$ represents the distance between plot k and plot k'. When plots are not similarly dense along the length of the boundary, there is no guarantee that these two assumptions are equivalent (see Keele & Titiunik (2015); Cattaneo et al. (2018) for a discussion of how boundary designs map into conventional RDD frameworks).

In our context, this is indeed a concern. The depth of the cultivable area on either side of the main canal depends on characteristics of the hillside and terracing decisions. Hence, it is not the case that plots are identically distributed on either side of the canal. To establish robustness to this assumption, we present a second set of results using spatial fixed effects (SFE; see Goldstein & Udry (2008); Conley & Udry (2010); Magruder (2012, 2013)). 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 $\overline{z}_{kist} = (1/|\mathcal{N}_{kist}|)\sum_{k'\in\mathcal{N}_{kist}} z_{k'i'st}$. The SFE specification then estimates

$$y_{1ist} - \overline{y}_{1ist} = \beta_1(\text{CA}_{1is} - \overline{\text{CA}}_{1is}) + (V_{1is} - \overline{V}_{1is})'\gamma + (\epsilon_{1ist} - \overline{\epsilon}_{1ist}) \qquad (2.2)$$

where V_{kis} includes all controls from Equation 2.1, except the subsumed site-by-season fixed effects. The SFE design is consistent under the assumption that $E[\epsilon_{1ist} - \overline{\epsilon}_{1ist}]$ is uncorrelated with $\text{CA}_{1is} - \overline{\text{CA}}_{1is}$ and the differenced V's, that is, that differences in unobservable characteristics among nearby plots are uncorrelated with explanatory variables.[12]

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.

[12]This approach is similar to pairwise matching across the boundary (e.g., Dube et al. (2010)) and boundary designs with segment fixed effects (e.g., Dell (2010)).

Balance

We now use specifications (2.1) and (2.2) to examine whether the plots in our sample and the households who cultivate them are comparable at baseline.[13]

First, and crucially, Table 2.2 indicates that our sample plots are balanced in terms of ownership and rentals, and that the vast majority of sample plot owners on both sides of the canal owned the land over 5 years, or prior the start of the irrigation construction. There is, however, some imbalance on plot size; log area (measured in hundredths of hectares) is larger inside the command area than outside the command area. This imbalance is somewhat 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 2.3.2 on horticulture, input use, labor use, and yields, as all of these variables are much larger in smaller plots in both the command area and outside the command area. We do control for log area in all specifications that follow, but note that all our results are robust to its exclusion.

Following the ownership results, Table 2.3 examines the characteristics of households whose sample plots are just inside or just outside the command area. First, note that specifications that do not restrict to the discontinuity sample perform poorly here; we find significant imbalance on half of our variables, and the omnibus test rejects the null of balance. However, we observe balance for both RDD and SFE specifications; households with sample plots just inside the command area appear similar to households with sample plots just outside the command area. There are some marginally significant differences in whether the household head has completed primary schooling or not, though we note that 1 out of 10 variables significant at the 10% level is what one would expect due to chance.

Lastly, in Section 2.5.1, we consider the characteristics of households' most important plots; we show that these appear similarly balanced.

2.3.2 Estimating the effects of irrigation

Adoption Dynamics

Figure 2.3 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

[13] Note that when testing for balance, we do not include the controls X_{1is}.

area are approximately 5pp more likely to be irrigated than sample plots outside the command area. 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 18pp-25pp 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 (2.1) and (2.2) pooling across our three years of follow up surveys, splitting our results across dry and rainy seasons.

Impacts of irrigation

We first estimate the impact of access to irrigation on cultivation decisions. For parsimony, we present graphical evidence of the regression discontinuity only for the dry seasons (16C, 17C, and 18C).[14] In each of these 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. Figure 2.4 examines cultivation and irrigation decisions on sample plots. We observe a large increase in the use of irrigation across the boundary, confirming that the command area increased access to irrigation. However, we do not observe any meaningful increase in cultivation in the dry season.

We consider this result in regression form in Table 2.4. Columns 1 and 4 present the sample means and standard deviations of the outcome on plots just outside the command area in the regression discontinuity sample, as well as sample sizes used in regressions. Columns 2, 3, 5, and 6 present results from RDD and SFE specifications estimating (2.1) and (2.2). Columns 1, 2, and 3 present results in the dry season, while Columns 4, 5, and 6 present results in the rainy season. We confirm this pattern: command area plots are no more likely to be cultivated during the dry season than plots outside the command area. This result is surprising, given the anticipated role of irrigation in enabling horticulture cultivation in the dry season. Less surprisingly, we do not observe robust differences in cultivation inside relative to outside the command area during the rainy season.

Figure 2.5 explains this discrepancy. While farmers are similarly likely to cultivate land on either side of the canal, their crop choice in the dry season changes sharply. Outside the command area, most farmers who cultivate in the dry season are cultivating bananas. Inside the command area, most farmers who cultivate in the dry season are cultivating horticulture. Additionally, note that comparing Figure 2.4 with Figure 2.5, it is apparent that the decision to irrigate in the dry season in the command area

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

occurs in very similar proportions to the decision to cultivate horticulture. This is consistent with our data on dry season horticulture and irrigation use in Table 2.1, and is confirmed by our regression results. We estimate that dry season horticulture cultivation increases by 12.8-15.3pp relative to a base of 6.1% outside the command area, while dry season banana cultivation is reduced by 13.6-13.9pp relative to a 24.6% base outside the command area. We do not see robust effects on horticulture cultivation during the rainy season. In contrast, the decline in banana cultivation is of similar magnitude across seasons (15.8-17.0pp relative to a 27.3% base outside the command area in the rainy season), in line with the fact that bananas are a perennial crop.

We now examine the labor and input allocations and productive outcomes of farmers on their sample plots. Figure 2.6 shows that yields increase sharply at the boundary in the dry season. It also indicates a discrete jump in total expenditures (expenditures on inputs plus hired labor), though the magnitude of this increase is small relative to the increase in yields. Tables 2.5 and 2.6 put numbers on these increases: dry season yields increase by 62,500-71,000 RwF/ha, and as horticultural crops are relatively commercial, a 1 RwF/ha increase in yields corresponds to a 0.78-0.80 RwF/ha increase in sales. Total dry season expenditures increase much more modestly, by 5,400-8,600 RwF/ha, with 62-74% of this increase coming from inputs (rather than hired labor). In contrast, in the rainy seasons, we fail to reject that there is no effect on any input allocations or productive outcomes.

Figure 2.7 demonstrates that in addition to monetized inputs, household labor also increases sharply for plots in the command area. Table 2.5 shows that household labor increases by 67-78 days/ha in the dry season. This is an order of magnitude larger than any increase in hired labor. The magnitude of this effect of access to irrigation on labor suggests that the economic profitability of irrigation depends critically on the shadow wage. Figure 2.8 demonstrates that dry season "cash profits", defined as yields net of input and hired labor expenditures, increases sharply at the boundary. This measure of profits implicitly assumes a shadow wage of 0 RwF/day, which may be inaccurate. Alternatively, valuing household labor at the market wage of 800 RwF/day, we observe no increase in profits in the command area. Table 2.6 confirms that in the dry season, cash profits increase by 55,000-68,000 RwF/ha, a 77-95% increase over dry season cash profits outside the command area. In contrast, because of the large increase in household labor use, there are no significant increases in profits valuing household labor at the market wage.

Taken together, these results together 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 earn substantially higher cash profits in the dry season but similar cash profits in the other seasons; and they invest more in inputs and much more in household labor in the dry seasons. We also learn that the shadow wage, and therefore labor market failures, is likely to be important for the decision to cultivate horticulture. Building on this last result, in the next section we adapt the model from Benjamin (1992) to develop tests for the role of market failures in adoption of irrigation.

2.4 Testing for binding constraint

2.4.1 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 labor applied to plot k. The common production shock σ is a random variable such that $\sigma \sim \Psi(\sigma), E[\sigma] = 1$. 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 $F_{kM} > 0, F_{kL} > 0, F_{kML} > 0, F_{kMM} < 0, F_{kLL} < 0$.[15] 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 $\overline{L}$ which is allocated to labor on each plot, leisure, and up to $\overline{L^O}$ units of off farm labor L^O. Finally, we model irrigation access as an increase in A_1. In each case, we develop the necessary assumptions to imply the results above: that this increase in A_1 generates an increase in demand for inputs and labor on plot A_1.

Farmers maximize expected utility

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

[15] Among 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.

subject to the constraints enumerated above

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

In this framework, there are three crucial constraints farmers face that cause deviations from expected profit maximization. The first is that farmers have no access to insurance; this simplification captures the fact that access to income smoothing technologies has been shown to shift farmers towards riskier investments. The second is that farmers input use is constrained from above; this implies that constrained farmers underuse inputs relative to the optimum. This is consistent with evidence that providing farmers with improved access to credit leads to increased investment, and with the low input use we observe in our data. The third is that farmers off farm labor allocations are constrained from above; this implies that constrained farmers overuse on farm labor relative to the optimum when labor is priced at the market wage. We make this choice for two reasons. First, this is consistent with work in rural India finding that market wages are above the efficient wage due to nominal wage rigidities Kaur (2014), which Breza et al. (2018) suggests are consistent with social norms on wage setting. Potentially consistent with rigidities, we do not observe wages changing over time in our data, and potentially consistent with norms, 67% of our observed wages are either 700, 800, or 1000 RwF/person-day. Second, this is consistent with our stylized facts, which suggest that horticulture is less likely to be a profitable investment when labor is valued at the market wage; although this may explain why many farmers do not adopt horticulture, it is difficult to reconcile with the moderate adoption that we do observe.

After substituting in the constraints which bind with equality, we derive the following first order conditions[16]

$$(M_k) \quad \left(1 + \frac{\text{cov}(\sigma, u_c)}{\mathbf{E}[u_c]}\right) A_k F_{kM} = (1 + \lambda_M) r \tag{2.3}$$

$$(L_k) \quad \left(1 + \frac{\text{cov}(\sigma, u_c)}{\mathbf{E}[u_c]}\right) A_k F_{kL} = (1 - \lambda_L) w \tag{2.4}$$

$$(\ell) \quad \frac{\mathbf{E}[u_\ell]}{\mathbf{E}[u_c]} = (1 - \lambda_L) w \tag{2.5}$$

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

[16]The derivation is in Appendix B.2.

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. 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 do not 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).

2.4.2 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 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 constraints bind, 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 typically require additional assumptions on the shape of the utility function or on the distribution of σ; we flag 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.*[17]

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.[18] 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 insurance reduces the marginal utility from agricultural production relative to the marginal utility from consumption.[19] In turn, this causes labor and input allocations to the most important plot to fall.

Although this produces a test of separation, this does not allow us to test for which sets of constraints might generate separation failures. This is because 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,

[17] See proof in Appendix B.2.

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

[19] This 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 B.2.

fails, because inputs, labor, and horticulture are all complements in the production function.

2.4.3 Separating constraints

To shed some light on which 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 $\overline{L}$, the household's total available labor, and wealth to shift $\overline{M}$, 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.*[20]

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. When 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.

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.[21]

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*

[20] See proof in Appendix B.2.

[21] If *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.

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.[22]

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.[23] 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 households allocate more labor to their plots.

We focus on one particular case that builds on this intuition, presented in Figure 2.9. 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 $\overline{L}$), 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.[24] 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, summarised in Table 2.7, 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

[22] See proof in Appendix B.2.

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

[24] This 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).

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 evidence for the presence of labor constraints, and observing larger households as more responsive as inconsistent with our model.

2.5 Separation failures and adoption of irrigation

2.5.1 Empirical strategy

Our benchmark specification to test for separation failures mirrors our benchmark specification to estimating the impacts of irrigation. We still make use of the regression 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).

$$\begin{aligned} y_{2ist} = \beta_0 + \beta_1 \mathrm{CA}_{1is} + \beta_2 \mathrm{Dist}_{1is} + \beta_3 \mathrm{CA}_{1is} * \mathrm{Dist}_{1is} \\ + \beta_4 \mathrm{CA}_{2is} + X'_{1is}\gamma_1 + X'_{2is}\gamma_2 + \alpha_{st} + \epsilon_{2ist} \end{aligned} \tag{2.6}$$

This specification also includes controls CA_{2is}, an indicator for whether the most important plot is in the command area, and X_{2is}, the log area of the most important plot. 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 $\mathrm{CA}_{1is} * \mathrm{CA}_{2is}$ 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 2.7, we interpret negative coefficients on labor, inputs, irrigation use, and horticulture, as evidence of separation failures.

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.

$$\begin{aligned} y_{2ist} = \beta_0 + \beta_1 \mathrm{CA}_{1is} + W'_i\beta_2 + \mathrm{CA}_{1is} * W'_i\beta_3 + \beta_4 \mathrm{Dist}_{1is} + \beta_5 \mathrm{CA}_{1is} * \mathrm{Dist}_{1is} \\ + \beta_6 \mathrm{CA}_{2is} + X'_{1is}\gamma_1 + X'_{2is}\gamma_2 + \alpha_{st} + \epsilon_{2ist} \end{aligned} \tag{2.7}$$

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 2.7 presents which signs map to which market failures.

In all cases, we also report coefficients from SFE specifications. These specifications include the same controls, except the subsumed site-by-season fixed effects, and are described in detail in Section 2.3.1.

Balance

We now use specification (2.6) 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.[25] Balance tests for most important plots are reported in Table 2.8. 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 also located in the command area, as households' plots tend to be located near each other. Second, note that our benchmark RDD specification and SFE specification both correct for this imbalance. Although we have a p-value of less than 0.1 for one variable (an indicator for terracing), the omnibus test fails to reject the null of balance.

2.5.2 Results

A test for separation failures

First, the graphical intuition behind the test for separation failures is captured in Figures 2.10. In this figure, irrigation use on the sample plot and the most important plot is plotted against the distance of sample plot to the command area boundary. As presented in Figure 2.4, irrigation use on the sample plot is 17pp higher for sample plots just inside the command area compared to sample plots just outside the command area. However, we now see that on most important plots, irrigation use is 5pp *lower* when the sample plot is just inside the command area relative to when the sample plot is just outside the command area. This result represents a separation failure; as discussed in Section 2.4.2, the technology on the sample plot does not directly affect optimal allocations on the most important plot.

Note that this result is distinct from many other tests of separation failures, as it implies that in our context, the separation failure generates inefficiencies: we observe technologically identical most important plots, distinct only through the managing household and the technology of their sample plot, receiving different allocations of inputs. This

[25] Note that when testing for balance, as previously we do not include the controls $(X'_{1is}, X'_{2is}, \text{CA}_{2is})$.

contrasts with tests that consider differences in on farm labor allocations or land cultivated across households of different sizes, either statically or dynamically, or leveraging between or within household variation (e.g., Benjamin (1992); LaFave & Thomas (2016); Dillon & Barrett (2017); Dillon et al. (2019)); in particular, their tests provide evidence that at least one market has failed, which is known to be insufficient to show inefficiency. Alternatively, another literature has used production function estimates to infer marginal products of labor, land, and inputs from their allocations (Jacoby, 1993; Skoufias, 1994; Restuccia & Santaeulalia-Llopis, 2017); although heterogeneity in these marginal products is sufficient for the existence of market failures, these tests are typically not robust to the presence of unobserved heterogeneity across plots or to measurement error (Gollin & Udry, 2019).

We present results on separation failures from our benchmark specification in Tables 2.9, 2.10, 2.11, and 2.12. For interpretation, coefficients for sample plots are presented in Columns 1 and 2, and the mean outcome on the most important plot for sample plots just outside the command area is presented in Column 3. Columns 4 and 5 present our benchmark estimates of the effect of the sample plot shock on outcomes on the most important plot.

We discuss some key findings. First, irrigation use falls by 4.0-4.7pp on most important plots; this magnitude represents 39-45% of average irrigation use, and 26% of the command area effect on irrigation use.[26] In addition to being consistent with Figure 2.10, and with the presence of separation failures, the magnitude of this estimate is important, as it represents a within households negative spillover of the command area; we discuss how this affects our interpretation of our main reduced form estimates in Section 2.3 in the following paragraphs. Second, we observe similar decreases for horticulture (3.5-4.9pp), household labor (41.5-43.2 person-days/ha), and inputs (5,600-6,400 RwF/ha). However, although they are less robust, we observe increases in bananas (3.8-8.9pp); as these are a less labor and input intensive crop, this is consistent with our interpretation of the production function as the envelope of production functions across crop choices.

Next, we expect the results above to be driven primarily by most important plots located in the command area for most outcomes, as 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 and 7, 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 15-18pp increase in irrigation

[26] Although the p-value on this result is .120-.129, this specification loses power by considering irrigation use on most important plots outside the command area, which are almost never irrigated. 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 as an outcome.

use on sample plots in the command area coincides with a 10-12pp 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 2.3, 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 2.4 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 Table 2.13 we present results of the sample plot shock on labor and input use on sample plots and most important plots, controlling for crop fixed effects.[27] Columns 5 and 6 confirm that the effects we document in Section 2.3 are driven by the shift to dry season horticulture, as effects on sample plots disappear controlling for crop choice. However, Columns 3 and 4 suggest 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%-41%. Combined with our results on irrigation use and horticulture, this suggests that both intensive and extensive margin responses on most important plots are important in response to the sample plot shock.

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 input allocations on those plots. This implies we cannot interpret our reduced form estimates of the impacts of irrigation as the effect of building hillside irrigation schemes, or even as the effect of adding one additional plot to the irrigation scheme on adoption of irrigation. However, we can interpret them as the causal effect of access to irrigation on the sample plot on production on the sample plot. Additionally, our estimates on separation failures provide evidence on the negative spillover generated by separation failures: in particular, we know the impact of moving one plot into the command area on the household's likelihood of irrigating a second command area plot.

To quantify the degree to which separation failures affect our reduced form estimates of impacts of the command area, 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. To do so, we conduct a simple exercise where we increase adoption

[27] As 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.

of irrigation, on all command area plots held by households with multiple command area plots, by our point estimate for the effect of the sample plot shock on irrigation use on most important plots in the command area. This exercise suggests that adoption would be 21-24% higher with perfect insurance and if inputs to production flowed frictionlessly between househoulds. We interpret this estimate to be conservative for two reasons. First, we treat households with 3 or more command area plots the same as households with 2 command area plots; we do so because our research design has little to say about the impacts of two sample plot shocks as opposed to one sample plot shock on allocations to other plots. Second, this simple exercise abstracts from potential decreases in production driven by reduced labor and input allocations conditional on adopting irrigation; our specification with crop fixed effects provides suggestive evidence that accounting for these responses would decrease an estimate of counterfactual productivity without separation failures.

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 either wealthier or larger 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 2.14 and 2.15. 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 70-80% for irrigation use, 62-86% for horticulture, 54-58% for household labor, and 18-36% for inputs, with all but the input coefficient statistically significant and robust across specifications.[28] In contrast, wealthier households are more responsive to the sample shock across these same outcomes. A household with a one standard deviation higher asset index is more

[28]These percentages, and the remainder of percentages in this paragraph, are expressed relative to the average estimated sample plot shock.

responsive to the sample plot shock on its most important plot by 40-80% for irrigation use, 40-63% for horticulture, 35-58% for household labor, and 42-100% for input use; however, these results are less robust, as statistical significance drops for all outcomes except inputs in the RDD specification. 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 2.4.3, 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 inefficiently low adoption of irrigation.

Additional evidence on input and information constraints

Work in progress: Minikit RCT results.

2.6 Conclusion

Work in progress: Irrigation in Africa.

Work in progress: Separation failures and technology adoption.

Figure 2.1: Timeline

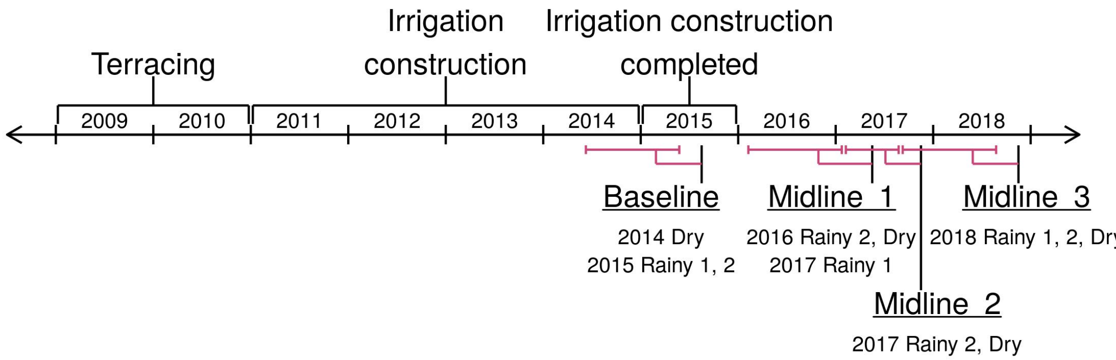

Notes: A timeline of events on the 3 hillside irrigation schemes we study is presented in this figure. Black lines are used to indicate when (or the period during which) events took place, while pink lines are used to indicate survey recall periods.

Figure 2.2: Hillside irrigation scheme

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

Figure 2.3: Adoption dynamics

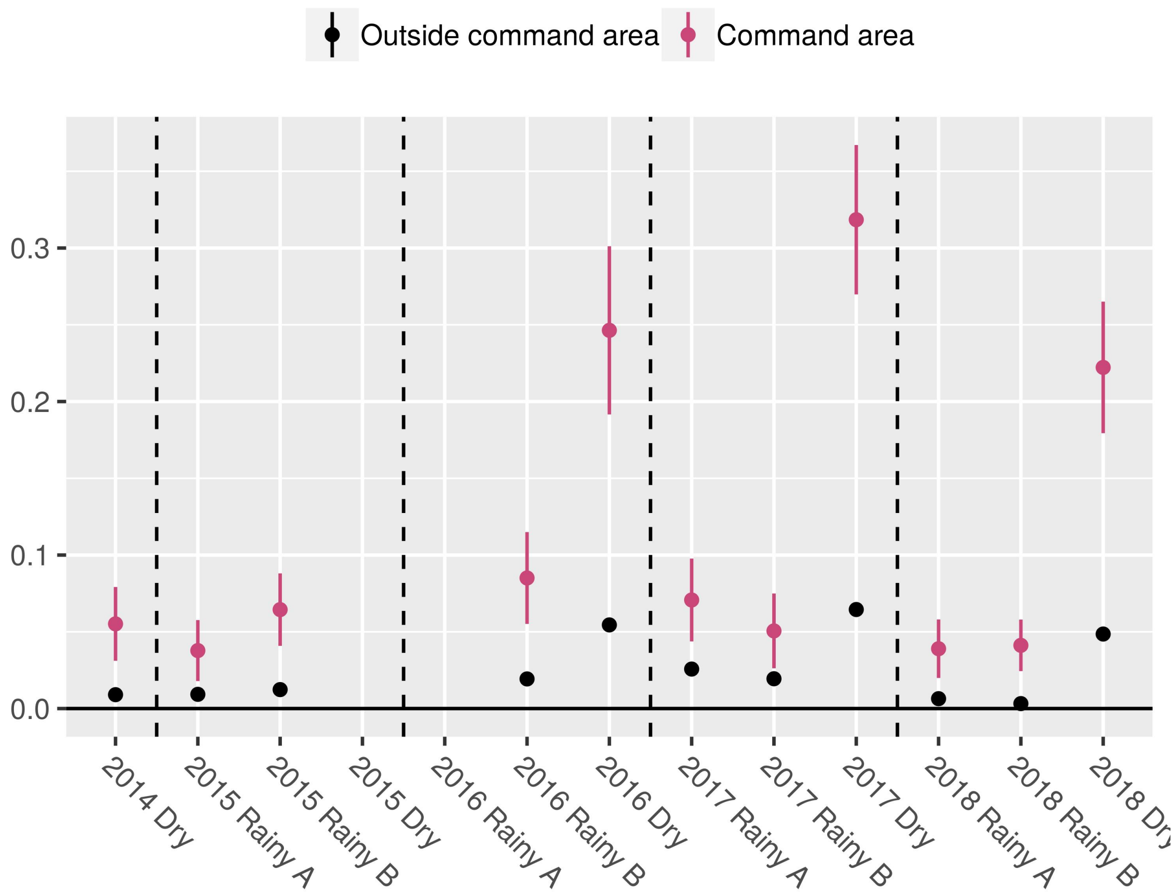

Notes: Average adoption of irrigation by season on sample plots in the main discontinuity sample, inside and outside the command area, is presented in this figure. Averages outside the command area are in black, while averages inside the command area and 95% confidence intervals for the difference are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

Figure 2.4: RDD: Irrigation

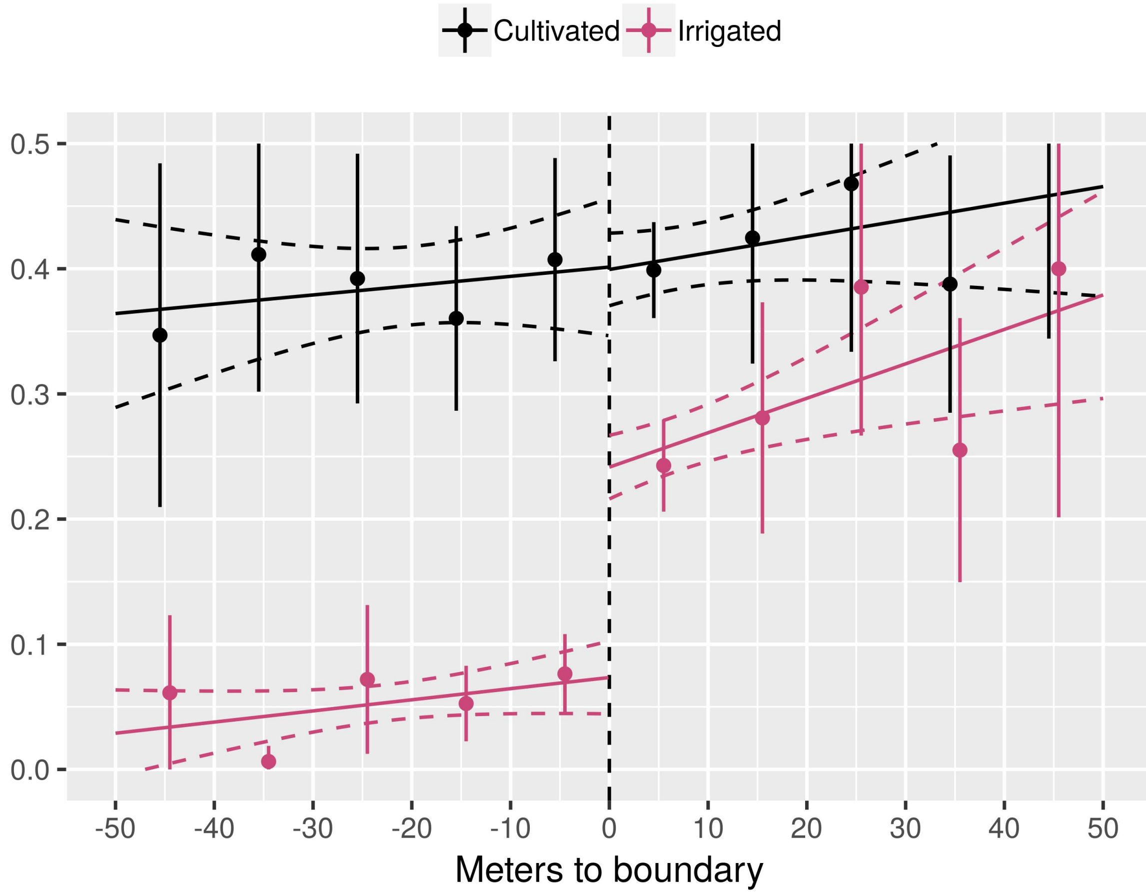

Notes: A visual regression discontinuity analysis for cultivation and irrigation, both on sample plots in the main discontinuity sample during the dry season, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes, and vertical lines through those points are 95% confidence intervals on the mean. 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 2.5: RDD: Crop choice

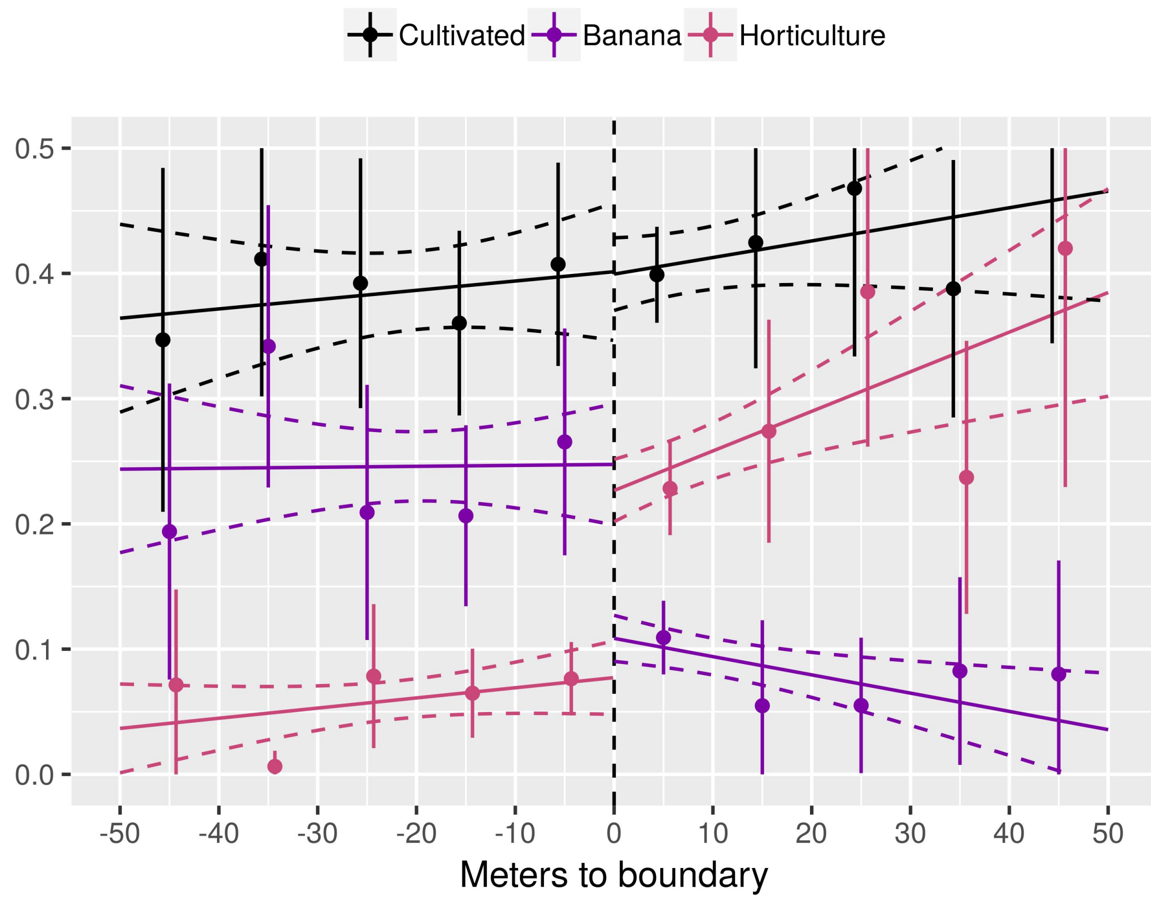

Notes: A visual regression discontinuity analysis for cultivation, banana cultivation, and horticulture cultivation, all on sample plots in the main discontinuity sample during the dry season, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes, and vertical lines through those points are 95% confidence intervals on the mean. 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 2.6: RDD: Yield

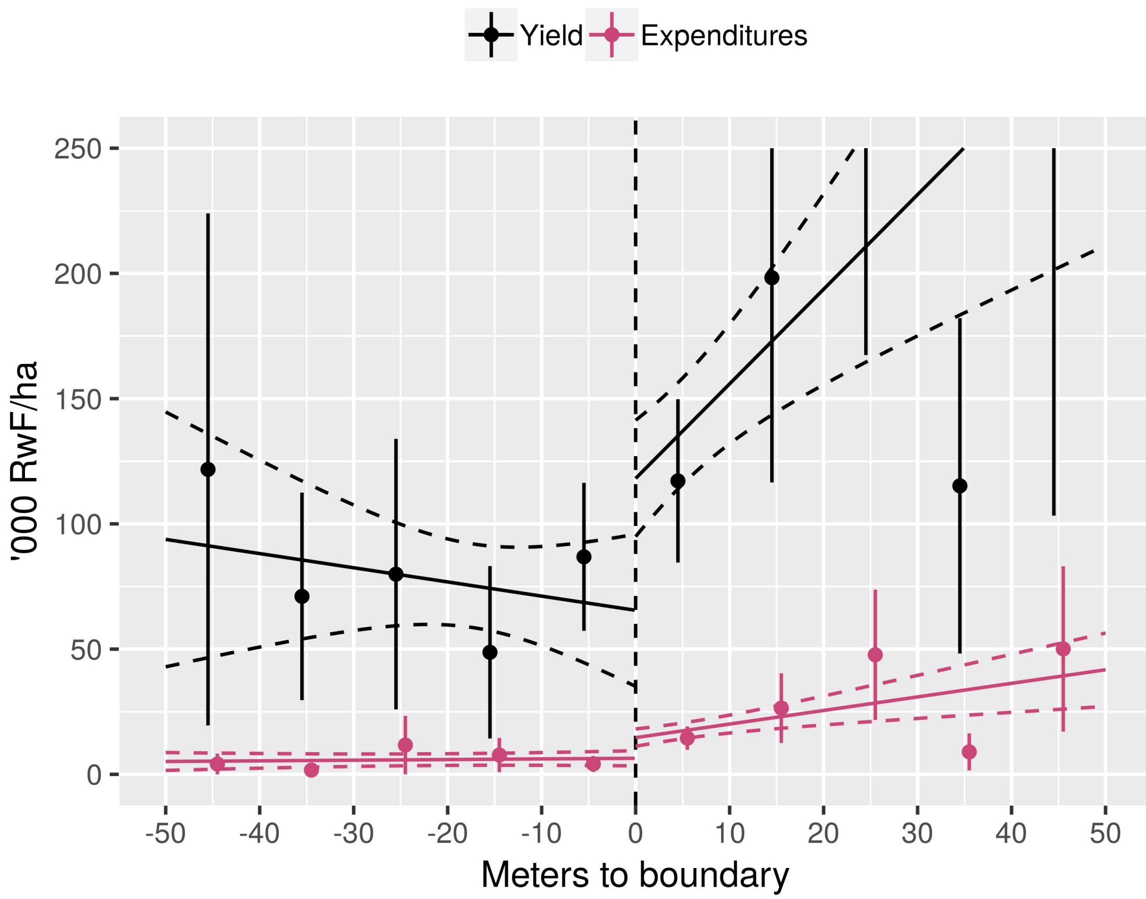

Notes: A visual regression discontinuity analysis for yields and expenditures, both on sample plots in the main discontinuity sample during the dry season, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes, and vertical lines through those points are 95% confidence intervals on the mean. 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 2.7: RDD: Labor

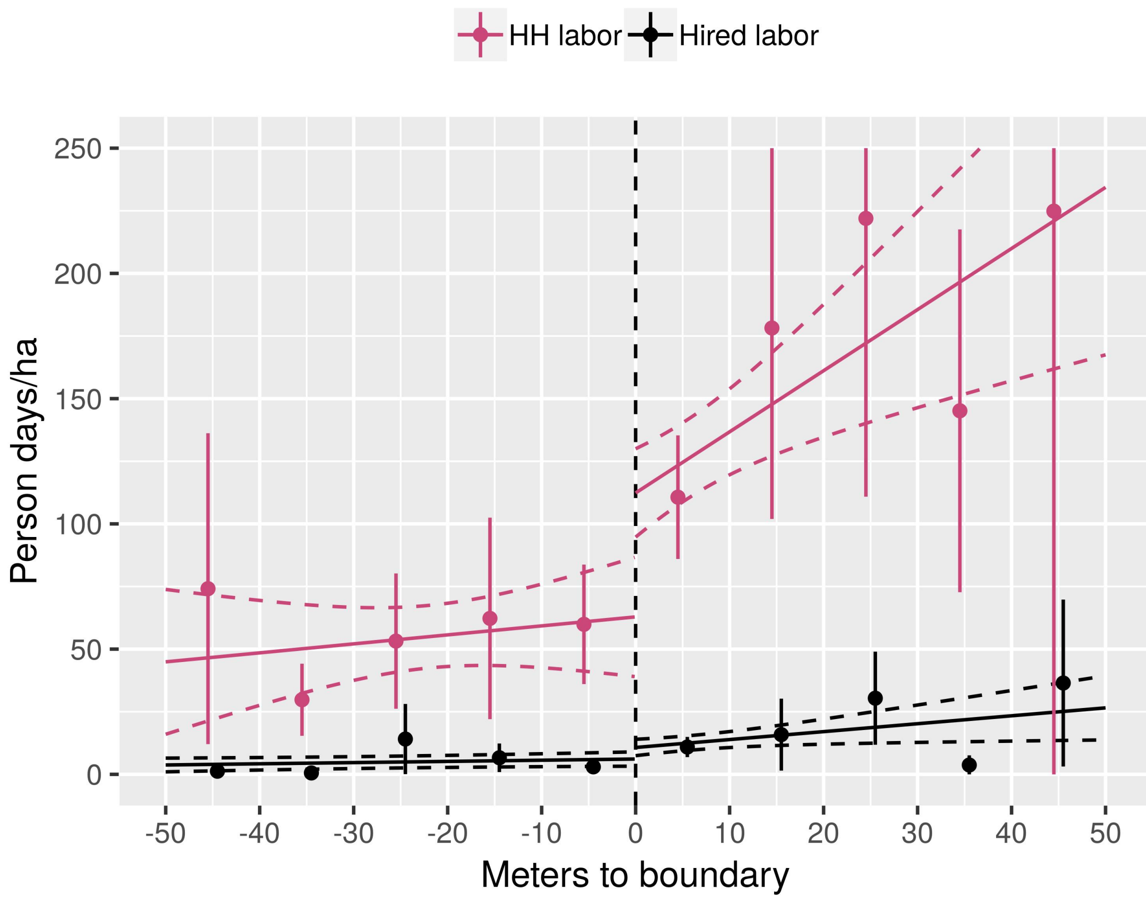

Notes: A visual regression discontinuity analysis for household labor and hired labor, both on sample plots in the main discontinuity sample during the dry season, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes, and vertical lines through those points are 95% confidence intervals on the mean. 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 2.8: RDD: Profits

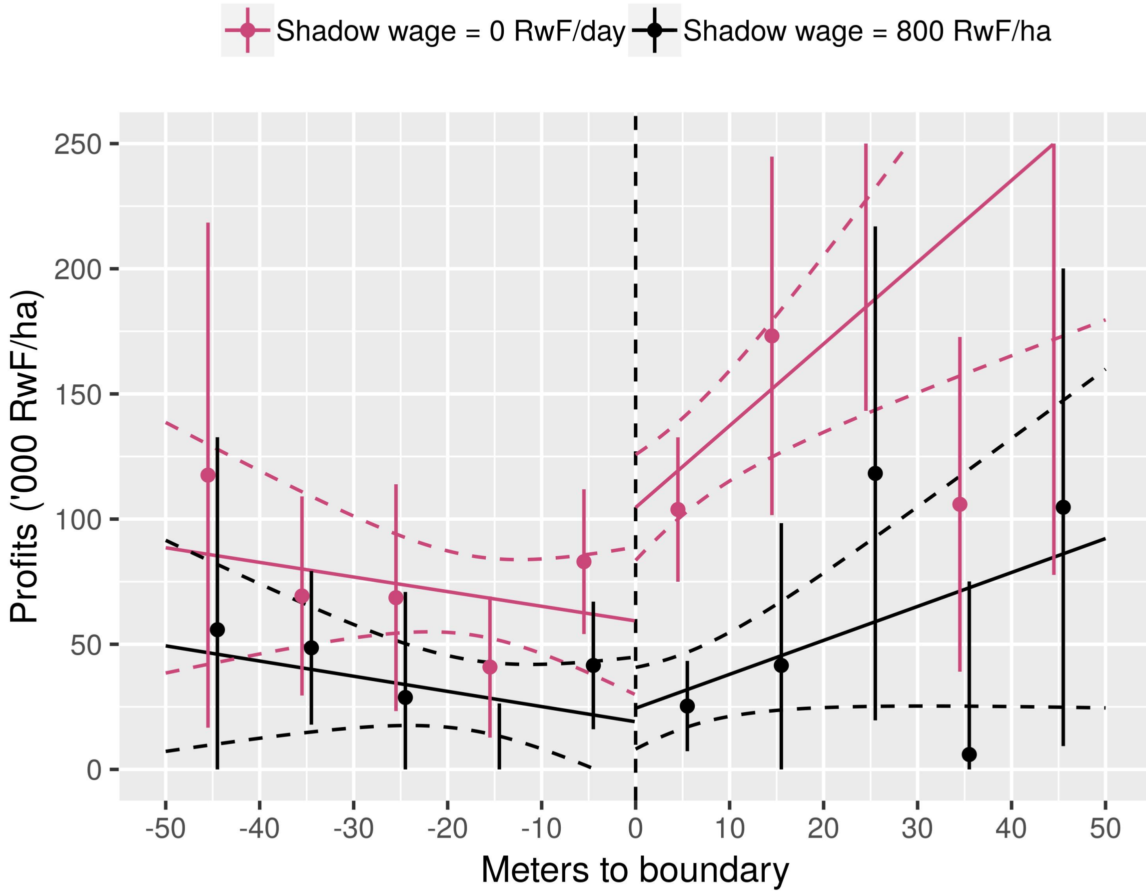

Notes: A visual regression discontinuity analysis for profits, on sample plots in the main discontinuity sample during the dry season, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes, and vertical lines through those points are 95% confidence intervals on the mean. 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 2.9: Differential responses under labor constraints

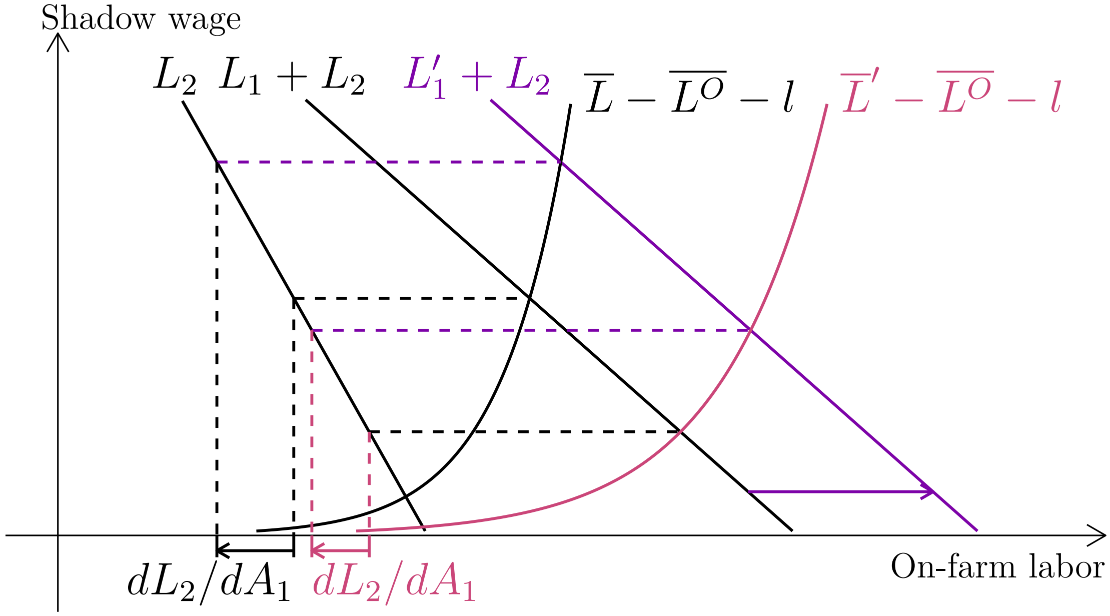

Notes: Households' labor allocations under a binding off farm labor constraint are presented in this figure. L_k and l are the household's labor allocation on plot k and choice of leisure, respectively, as a function of the shadow wage, with the argument suppressed. $L_1 + L_2$ is total household on 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 $L_1' + L_2$. $\overline{L} - \overline{L^O} - l$ is household on farm labor supply; for large households, on farm labor supply is shifted out to $\overline{L}' - \overline{L^O} - l$. 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 2.10: RDD: Most important plot, irrigation

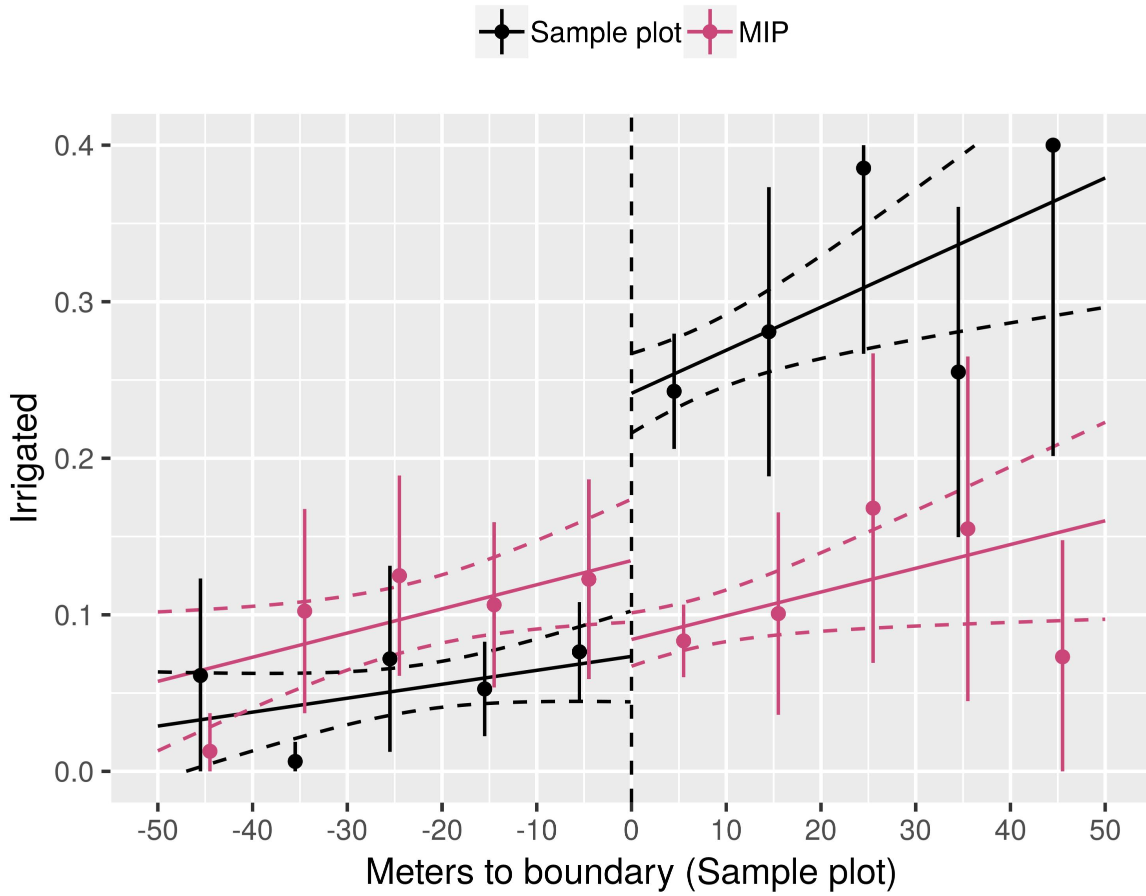

Notes: A visual regression discontinuity analysis for irrigation on sample plots and associated most important plots during the dry season, for sample plots in the main discontinuity sample, is presented in this figure. Distance to the boundary is reported in meters, with positive distance corresponding to sample plots inside the command area. Points are binned average outcomes, and vertical lines through those points are 95% confidence intervals on the mean. 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 2.1: Summary statistics on agricultural production

	Staples				Horticulture		
	All	Maize	Beans	Bananas	All	Rainy	Dry
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Yield	316	336	301	255	588	580	596
Hired labor (days)	39	39	40	11	61	62	60
HH labor (days)	304	307	291	104	465	451	478
Inputs	19	38	17	3	46	46	46
Profits							
Shadow wage = 0 RwF/day	269	269	254	243	496	486	506
Shadow wage = 800 RwF/day	26	23	21	160	124	125	124
Sales share	0.15	0.19	0.13	0.51	0.60	0.57	0.63
Irrigated	0.02	0.02	0.02	0.02	0.58	0.22	0.92
Rainy	0.99	1.00	1.00	0.49	0.48	1.00	0.00
log area	1.93	1.98	1.93	2.42	1.72	1.62	1.81
Share of obs.	0.68	0.10	0.45	0.17	0.12	0.06	0.06

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 hundredths of a hectare. All other variables are shares or indicators. For reference, the median wage in our data is 800 RwF/person-day.

Table 2.2: Balance: Sample plots

		Coef. (SE) [p]		
		Full Sample	RDD	SFE
	(1)	(2)	(3)	(4)
log area	2.090	0.045	0.425	0.252
	(1.179)	(0.077)	(0.121)	(0.147)
	969	[0.554]	[0.000]	[0.087]
Own plot	0.894	-0.012	0.004	-0.036
	(0.309)	(0.020)	(0.032)	(0.042)
	969	[0.535]	[0.907]	[0.391]
Owned plot >5 years	0.880	0.045	0.019	-0.008
	(0.326)	(0.019)	(0.037)	(0.035)
	686	[0.020]	[0.613]	[0.817]
Rented out to farmer	0.032	0.027	-0.003	0.005
	(0.177)	(0.012)	(0.023)	(0.030)
	969	[0.022]	[0.884]	[0.864]
Omnibus F-stat [p]		2.6	3.2	1.1
		[0.109]	[0.074]	[0.307]

Notes: Balance for sample plot characteristics is presented in this table. 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 4 present regression coefficients on a command area indicator, with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. 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 2 compares outcomes inside and outside the command area in the full sample. Column 3 uses the regression discontinuity specification in Equation (2.1), omitting controls X_{1is}. Column 4 uses the spatial fixed effects specification in Equation (2.2), omitting controls X_{1is}. Robust standard errors are clustered at the level of the nearest water user group.

Table 2.3: Balance: Households

		Coef. (SE) [p]		
		Full Sample	RDD	SFE
	(1)	(2)	(3)	(4)
HHH female	0.221	0.041	0.045	0.037
	(0.416)	(0.025)	(0.046)	(0.058)
	969	[0.094]	[0.326]	[0.522]
HHH age	47.5	0.5	2.1	0.3
	(14.5)	(0.8)	(1.4)	(1.8)
	967	[0.497]	[0.127]	[0.871]
HHH completed primary	0.287	0.069	0.128	0.129
	(0.453)	(0.025)	(0.047)	(0.061)
	966	[0.005]	[0.006]	[0.034]
HHH worked off farm	0.410	0.023	-0.039	0.004
	(0.493)	(0.027)	(0.051)	(0.063)
	969	[0.392]	[0.441]	[0.945]
# of plots	5.19	0.61	0.20	0.51
	(3.38)	(0.18)	(0.36)	(0.47)
	969	[0.001]	[0.582]	[0.280]
# of HH members	4.89	0.17	-0.00	0.07
	(2.16)	(0.11)	(0.21)	(0.28)
	969	[0.104]	[0.985]	[0.787]
# who worked off farm	0.77	0.10	0.01	0.11
	(0.85)	(0.05)	(0.08)	(0.11)
	969	[0.039]	[0.909]	[0.326]
Housing expenditures	49.2	-2.3	-5.6	-15.7
	(127.3)	(6.9)	(14.9)	(18.8)
	962	[0.739]	[0.707]	[0.406]
Asset index	-0.04	0.11	0.15	0.07
	(0.99)	(0.05)	(0.12)	(0.15)
	967	[0.034]	[0.203]	[0.633]
Omnibus F-stat [p]		3.6	1.8	1.3
		[0.059]	[0.185]	[0.260]

Notes: Balance for household characteristics is presented in this table. Column 1 presents, for households managing 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 present regression coefficients on a command area indicator, with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. 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 2 compares outcomes inside and outside the command area in the full sample. Column 3 uses the regression discontinuity specification in Equation (2.1), omitting controls X_{1is}. Column 4 uses the spatial fixed effects specification in Equation (2.2), omitting controls X_{1is}.

Table 2.4: Sample plots

	Dry season			Rainy seasons		
		Coef. (SE) [p]			Coef. (SE) [p]	
	Dep. var.	RDD*	SFE*	Dep. var.	RDD*	SFE*
	(1)	(2)	(3)	(4)	(5)	(6)
Cultivated	0.387	-0.009	0.031	0.814	-0.102	-0.056
	(0.487)	(0.041)	(0.046)	(0.389)	(0.031)	(0.041)
	2,442	[0.822]	[0.507]	4,080	[0.001]	[0.175]
Irrigated	0.056	0.152	0.179	0.015	0.035	0.066
	(0.230)	(0.023)	(0.030)	(0.121)	(0.009)	(0.014)
	2,442	[0.000]	[0.000]	4,080	[0.000]	[0.000]
Horticulture	0.061	0.128	0.153	0.072	0.015	0.060
	(0.240)	(0.023)	(0.028)	(0.258)	(0.018)	(0.021)
	2,441	[0.000]	[0.000]	4,079	[0.396]	[0.005]
Banana	0.246	-0.136	-0.139	0.273	-0.158	-0.170
	(0.431)	(0.037)	(0.040)	(0.446)	(0.039)	(0.043)
	2,441	[0.000]	[0.000]	4,079	[0.000]	[0.000]

Notes: Regression analysis is presented in this table. Columns 1 through 3 restrict to observations during the dry season, while columns 4 through 6 restrict to observations during the rainy season. Columns 1 and 4 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, 3, 5, and 6 present regression coefficients on a command area indicator, with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. Columns 2 and 5 use the regression discontinuity specification in Equation (2.1), omitting controls X_{1is}. Columns 3 and 6 use the spatial fixed effects specification in Equation (2.2), omitting controls X_{1is}.

Table 2.5: Sample plots

	Dry season			Rainy seasons		
		Coef. (SE) [p]			Coef. (SE) [p]	
	Dep. var.	RDD*	SFE*	Dep. var.	RDD*	SFE*
	(1)	(2)	(3)	(4)	(5)	(6)
HH labor/ha	55.8	67.1	78.0	219.3	-1.3	-1.1
	(197.0)	(17.5)	(23.4)	(320.2)	(25.4)	(28.9)
	2,438	[0.000]	[0.001]	4,072	[0.958]	[0.969]
Inputs/ha	2.3	5.3	4.0	15.4	-0.6	1.4
	(16.3)	(1.4)	(1.6)	(40.0)	(2.8)	(3.2)
	2,442	[0.000]	[0.014]	4,080	[0.821]	[0.651]
Hired labor exp./ha	3.6	3.3	1.4	14.9	2.8	0.7
	(25.6)	(2.1)	(2.8)	(45.7)	(3.3)	(4.3)
	2,442	[0.126]	[0.635]	4,080	[0.387]	[0.869]

Notes: Regression analysis is presented in this table. Columns 1 through 3 restrict to observations during the dry season, while columns 4 through 6 restrict to observations during the rainy season. Columns 1 and 4 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, 3, 5, and 6 present regression coefficients on a command area indicator, with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. Columns 2 and 5 use the regression discontinuity specification in Equation (2.1), omitting controls X_{1is}. Columns 3 and 6 use the spatial fixed effects specification in Equation (2.2), omitting controls X_{1is}.

Table 2.6: Sample plots

	Dry season			Rainy seasons		
		Coef. (SE) [p]			Coef. (SE) [p]	
	Dep. var.	RDD*	SFE*	Dep. var.	RDD*	SFE*
	(1)	(2)	(3)	(4)	(5)	(6)
Yield	76.8	62.5	71.0	257.1	-27.7	-20.4
	(260.6)	(22.9)	(30.5)	(445.2)	(31.6)	(38.1)
	2,317	[0.007]	[0.020]	3,942	[0.381]	[0.593]
Sales/ha	47.8	48.7	56.9	78.0	-11.6	19.5
	(174.5)	(14.5)	(21.1)	(218.2)	(18.1)	(23.2)
	2,442	[0.001]	[0.007]	4,080	[0.523]	[0.400]
Profits/ha						
Shadow wage = 0	71.0	54.9	67.5	227.0	-28.7	-21.0
	(246.8)	(20.9)	(28.0)	(421.4)	(29.4)	(36.0)
	2,317	[0.009]	[0.016]	3,942	[0.329]	[0.560]
Shadow wage = 800	31.1	4.6	18.3	52.7	-25.5	-18.6
	(217.0)	(17.0)	(24.8)	(351.9)	(25.8)	(34.3)
	2,315	[0.788]	[0.462]	3,935	[0.323]	[0.587]

Notes: Regression analysis is presented in this table. Columns 1 through 3 restrict to observations during the dry season, while columns 4 through 6 restrict to observations during the rainy season. Columns 1 and 4 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, 3, 5, and 6 present regression coefficients on a command area indicator, with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. Columns 2 and 5 use the regression discontinuity specification in Equation (2.1), omitting controls X_{1is}. Columns 3 and 6 use the spatial fixed effects specification in Equation (2.2), omitting controls X_{1is}.

Table 2.7: Model predictions

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

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

Table 2.8: Balance: Most important plot

		Coef. (SE) [p]		
		Full Sample	RDD	SFE
	(1)	(2)	(3)	(4)
log area	2.225	-0.108	0.094	0.027
	(1.041)	(0.068)	(0.128)	(0.164)
	784	[0.114]	[0.460]	[0.869]
Own plot	0.875	0.025	0.040	0.008
	(0.331)	(0.019)	(0.033)	(0.037)
	784	[0.174]	[0.226]	[0.823]
Owned plot >5 years	0.960	0.005	0.012	0.023
	(0.197)	(0.014)	(0.024)	(0.024)
	585	[0.728]	[0.617]	[0.344]
Rented out to farmer	0.033	0.013	-0.026	-0.025
	(0.179)	(0.010)	(0.022)	(0.024)
	784	[0.224]	[0.249]	[0.290]
Command area	0.399	0.187	-0.053	-0.090
	(0.491)	(0.032)	(0.058)	(0.069)
	784	[0.000]	[0.360]	[0.191]
Terraced	0.626	0.017	-0.099	-0.111
	(0.485)	(0.028)	(0.053)	(0.063)
	784	[0.539]	[0.063]	[0.077]
Rented out to investor	0.081	0.035	-0.042	-0.041
	(0.273)	(0.018)	(0.040)	(0.042)
	784	[0.054]	[0.292]	[0.323]
Omnibus F-stat [p]		5.6	1.3	0.9
		[0.019]	[0.265]	[0.336]

Notes: Balance for most important plot characteristics is presented in this table. Column 1 presents, for most important plots for which the associated sample plot is in the main discontinuity sample and located 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 present regression coefficients on a command area indicator, with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. 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 2 compares outcomes inside and outside the command area in the full sample. Column 3 uses the regression discontinuity specification in Equation (2.6), omitting controls $(\text{CA}_{2is}, X'_{1is}, X'_{2is})$; Column 4 uses a similar spatial fixed effects specification, omitting controls $(\text{CA}_{2is}, X'_{1is}, X'_{2is})$.

Table 2.9: Most important plot

	Sample plot		MIP				
	Coef. (SE) [p]			Coef. (SE) [p]			
	RDD* (1)	SFE* (2)	Dep. Var. (3)	RDD* (4)	SFE* (5)	RDD* (6)	SFE* (7)
Cultivated							
CA	-0.009	0.031	0.356	0.040	-0.032	0.082	0.013
	(0.041)	(0.046)	(0.479)	(0.041)	(0.048)	(0.044)	(0.054)
	[0.822]	[0.507]	2,123	[0.333]	[0.515]	[0.063]	[0.815]
CA * MIP in CA						-0.102	-0.099
						(0.054)	(0.060)
						[0.057]	[0.100]
Joint F-stat [p]						2.5	1.5
						[0.084]	[0.220]
Irrigated							
CA	0.152	0.179	0.104	-0.040	-0.047	0.003	0.008
	(0.023)	(0.030)	(0.306)	(0.026)	(0.031)	(0.020)	(0.029)
	[0.000]	[0.000]	2,123	[0.120]	[0.129]	[0.872]	[0.788]
CA * MIP in CA						-0.103	-0.123
						(0.036)	(0.041)
						[0.004]	[0.003]
Joint F-stat [p]						4.3	4.5
						[0.015]	[0.012]

Notes: Regression analysis is presented in this table. Columns 1 and 2 use outcomes on the sample plot (and replicate the analysis in Table 2.4), while Columns 3 through 7 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 3 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, 2, and 3 through 7, 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; robust standard errors clustered at the nearest water user group level are in parentheses, and p-values are in brackets. Column 4 uses the regression discontinuity specification in Equation 2.6; Column 5 uses a similar spatial fixed effects specification. Column 6 uses the regression discontinuity specification in Equation 2.7, where the interacted W_i is a command area indicator for the most important plot; Column 7 uses a similar spatial fixed effects specification.

Table 2.10: Most important plot

	Sample plot		MIP				
	Coef. (SE) [p]			Coef. (SE) [p]			
	RDD* (1)	SFE* (2)	Dep. Var. (3)	RDD* (4)	SFE* (5)	RDD* (6)	SFE* (7)
Horticulture							
CA	0.128	0.153	0.100	-0.035	-0.049	-0.000	-0.005
	(0.023)	(0.028)	(0.300)	(0.024)	(0.029)	(0.018)	(0.024)
	[0.000]	[0.000]	2,123	[0.147]	[0.091]	[0.982]	[0.837]
CA * MIP in CA						-0.082	-0.098
						(0.036)	(0.042)
						[0.024]	[0.021]
Joint F-stat [p]						2.6	2.9
						[0.079]	[0.057]
Banana							
CA	-0.136	-0.139	0.199	0.089	0.038	0.094	0.041
	(0.037)	(0.040)	(0.399)	(0.034)	(0.041)	(0.042)	(0.054)
	[0.000]	[0.000]	2,123	[0.008]	[0.351]	[0.027]	[0.442]
CA * MIP in CA						-0.013	-0.006
						(0.042)	(0.052)
						[0.764]	[0.902]
Joint F-stat [p]						3.7	0.5
						[0.027]	[0.636]

Notes: Regression analysis is presented in this table. Columns 1 and 2 use outcomes on the sample plot (and replicate the analysis in Table 2.4), while Columns 3 through 7 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 3 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, 2, and 3 through 7, 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; robust standard errors clustered at the nearest water user group level are in parentheses, and p-values are in brackets. Column 4 uses the regression discontinuity specification in Equation 2.6; Column 5 uses a similar spatial fixed effects specification. Column 6 uses the regression discontinuity specification in Equation 2.7, where the interacted W_i is a command area indicator for the most important plot; Column 7 uses a similar spatial fixed effects specification.

Table 2.11: Most important plot

	Sample plot		MIP				
	Coef. (SE) [p]			Coef. (SE) [p]			
	RDD* (1)	SFE* (2)	Dep. Var. (3)	RDD* (4)	SFE* (5)	RDD* (6)	SFE* (7)
HH labor/ha							
CA	67.1	78.0	66.6	-41.5	-43.2	-18.0	-18.2
	(17.5)	(23.4)	(219.6)	(21.2)	(23.6)	(15.4)	(19.0)
	[0.000]	[0.001]	2,120	[0.050]	[0.067]	[0.241]	[0.338]
CA * MIP in CA						-55.7	-55.8
						(25.0)	(30.7)
						[0.026]	[0.070]
Joint F-stat [p]						2.6	2.1
						[0.073]	[0.126]
Inputs/ha							
CA	5.3	4.0	5.4	-6.4	-5.6	-3.6	-2.7
	(1.4)	(1.6)	(28.3)	(2.8)	(3.0)	(1.9)	(2.4)
	[0.000]	[0.014]	2,123	[0.023]	[0.064]	[0.061]	[0.259]
CA * MIP in CA						-6.7	-6.5
						(3.2)	(3.3)
						[0.040]	[0.047]
Joint F-stat [p]						2.8	2.3
						[0.065]	[0.099]

Notes: Regression analysis is presented in this table. Columns 1 and 2 use outcomes on the sample plot (and replicate the analysis in Table 2.5), while Columns 3 through 7 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 3 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, 2, and 3 through 7, 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; robust standard errors clustered at the nearest water user group level are in parentheses, and p-values are in brackets. Column 4 uses the regression discontinuity specification in Equation 2.6; Column 5 uses a similar spatial fixed effects specification. Column 6 uses the regression discontinuity specification in Equation 2.7, where the interacted W_i is a command area indicator for the most important plot; Column 7 uses a similar spatial fixed effects specification.

Table 2.12: Most important plot

	Sample plot		MIP				
	Coef. (SE) [p]			Coef. (SE) [p]			
	RDD* (1)	SFE* (2)	Dep. Var. (3)	RDD* (4)	SFE* (5)	RDD* (6)	SFE* (7)
Hired labor exp./ha							
CA	3.3	1.4	3.8	-1.8	0.2	0.0	2.0
	(2.1)	(2.8)	(24.7)	(2.2)	(2.5)	(2.2)	(2.6)
	[0.126]	[0.635]	2,123	[0.420]	[0.935]	[0.997]	[0.451]
CA * MIP in CA						-4.2	-4.0
						(2.6)	(3.4)
						[0.109]	[0.246]
Joint F-stat [p]						1.4	0.7
						[0.254]	[0.487]

Notes: Regression analysis is presented in this table. Columns 1 and 2 use outcomes on the sample plot (and replicate the analysis in Table 2.5), while Columns 3 through 7 use outcomes on the associated most important plot. All columns restrict to observations during the dry season. Column 3 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, 2, and 3 through 7, 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; robust standard errors clustered at the nearest water user group level are in parentheses, and p-values are in brackets. Column 4 uses the regression discontinuity specification in Equation 2.6; Column 5 uses a similar spatial fixed effects specification. Column 6 uses the regression discontinuity specification in Equation 2.7, where the interacted W_i is a command area indicator for the most important plot; Column 7 uses a similar spatial fixed effects specification.

Table 2.13: Intensive margin effects

	MIP (Crop FE)			Sample plot (Crop FE)		
		Coef. (SE) [p]			Coef. (SE) [p]	
	Dep. var.	RDD*	SFE*	Dep. var.	RDD*	SFE*
	(1)	(2)	(3)	(4)	(5)	(6)
HH labor/ha	66.6	-34.0	-25.7	55.8	2.7	1.2
	(219.6)	(14.8)	(17.7)	(197.0)	(15.0)	(20.2)
	2,120	[0.022]	[0.146]	2,438	[0.859]	[0.952]
Inputs/ha	5.4	-4.9	-3.3	2.3	-0.1	-2.2
	(28.3)	(2.2)	(2.4)	(16.3)	(1.1)	(1.6)
	2,123	[0.022]	[0.160]	2,442	[0.923]	[0.157]
Hired labor exp./ha	3.8	-1.3	1.1	3.6	-0.1	-2.2
	(24.7)	(2.0)	(2.3)	(25.6)	(1.9)	(2.8)
	2,123	[0.508]	[0.628]	2,442	[0.967]	[0.432]

Notes: Regression analysis is presented in this table. Columns 1 through 3 use outcomes on the most important plot, while Columns 4 through 6 use outcomes on the sample plot. Columns 1 and 4 present, for sample plots (or associated 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, 3, 5, and 6 present regression coefficients on a command area indicator for the sample plot (or associated sample plot), with robust standard errors clustered at the nearest water user group level in parentheses, and p-values in brackets. Column 2 uses the regression discontinuity specification in Equation (2.2), with crop fixed effects included as controls; Column 3 uses a similar spatial fixed effects specification, with crop fixed effects included as controls. Column 5 uses the regression discontinuity specification in Equation (2.1), with crop fixed effects included as controls; Column 6 uses the similar spatial fixed effects specification in Equation (2.2), with crop fixed effects included as controls.

Table 2.14: Heterogeneity with respect to household size and wealth

	Coef. (SE) [p]			Coef. (SE) [p]	
	RDD* (1)	SFE* (2)		RDD* (1)	SFE* (2)
Cultivated			Horticulture		
CA	-0.093	-0.225	CA	-0.108	-0.120
	(0.087)	(0.096)		(0.045)	(0.049)
	[0.287]	[0.019]		[0.015]	[0.015]
CA * # of HH members	0.026	0.040	CA * # of HH members	0.015	0.015
	(0.013)	(0.015)		(0.007)	(0.008)
	[0.051]	[0.007]		[0.034]	[0.077]
CA * Asset index	-0.018	-0.052	CA * Asset index	-0.014	-0.031
	(0.027)	(0.032)		(0.016)	(0.019)
	[0.501]	[0.107]		[0.382]	[0.117]
Joint F-stat [p]	2.1	2.8	Joint F-stat [p]	2.0	2.3
	[0.100]	[0.043]		[0.120]	[0.081]
Irrigated			Banana		
CA	-0.109	-0.135	CA	0.036	-0.130
	(0.047)	(0.050)		(0.067)	(0.075)
	[0.022]	[0.008]		[0.594]	[0.082]
CA * # of HH members	0.014	0.019	CA * # of HH members	0.010	0.033
	(0.007)	(0.009)		(0.011)	(0.013)
	[0.051]	[0.031]		[0.365]	[0.008]
CA * Asset index	-0.016	-0.037	CA * Asset index	-0.009	-0.022
	(0.015)	(0.019)		(0.023)	(0.026)
	[0.310]	[0.048]		[0.688]	[0.385]
Joint F-stat [p]	1.8	2.8	Joint F-stat [p]	2.8	2.7
	[0.151]	[0.043]		[0.041]	[0.048]

Notes: Regression analysis is presented in this table. All columns use outcomes on most important plots. 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; robust standard errors clustered at the nearest water user group level are in parentheses, and p-values are in brackets. 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. Column 1 uses the regression discontinuity specification in Equation 2.7; Column 2 uses a similar spatial fixed effects specification.

Table 2.15: Heterogeneity with respect to household size and wealth

	Coef. (SE) [p]	
	RDD*	SFE*
	(1)	(2)
HH labor/ha		
CA	-96.3	-101.6
	(37.6)	(36.0)
	[0.011]	[0.005]
CA * # of HH members	11.2	12.5
	(4.7)	(4.6)
	[0.016]	[0.007]
CA * Asset index	-14.6	-25.1
	(10.5)	(12.5)
	[0.164]	[0.046]
Joint F-stat [p]	2.3	3.1
	[0.083]	[0.027]

	Coef. (SE) [p]	
	RDD*	SFE*
	(1)	(2)
Inputs/ha		
CA	-9.2	-9.1
	(4.4)	(4.6)
	[0.034]	[0.045]
CA * # of HH members	0.6	0.9
	(0.5)	(0.7)
	[0.244]	[0.199]
CA * Asset index	-2.7	-5.6
	(1.6)	(1.9)
	[0.090]	[0.003]
Joint F-stat [p]	2.0	3.2
	[0.109]	[0.025]
Hired labor exp./ha		
CA	-5.1	-2.5
	(3.7)	(3.0)
	[0.169]	[0.404]
CA * # of HH members	0.6	0.6
	(0.5)	(0.4)
	[0.192]	[0.181]
CA * Asset index	0.2	-1.3
	(1.4)	(1.3)
	[0.887]	[0.307]
Joint F-stat [p]	0.7	0.8
	[0.527]	[0.505]

Notes: Regression analysis is presented in this table. All columns use outcomes on most important plots. 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; robust standard errors clustered at the nearest water user group level are in parentheses, and p-values are in brackets. 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. Column 1 uses the regression discontinuity specification in Equation 2.7; Column 2 uses a similar spatial fixed effects specification.

Chapter 3

Marshallian consumer surplus from intertemporal substitution: Applications to savings, credit, and index insurance

3.1 Introduction

The envelope theorem is commonly applied in economics to derive Roy's identity. In a static model, the change in utility from a marginal price change of a good, normalized by the marginal utility of income, is the quantity consumed of the good. Marshallian consumer surplus, the welfare trapezoid yielded by the integral of quantity consumed over the price change, provides a useful summary metric of the effect on consumer welfare of the price change. In a dynamic model with uncertainty, constructing a single measure of the consumer surplus from a permanent price change is more challenging, since it depends on the household's discount factors in addition to its quantity consumed for each period-state. As a result, a common simplification is to instead calculate the average per-period consumer surplus, effectively assuming a static context.

I argue that Marshallian consumer surplus is a similarly good metric for per-period consumer surplus from a price change to risky assets in a dynamic model as it is for consumer goods in a static model. I do so using a model which nests approaches from Deaton (1991) and Carroll (2012), who focus on a special class of the risky assets I study to understand saving and borrowing behavior. To gain tractability in this setting, I avoid calculating the household's discount factors for each period-state, or trying to sum welfare gains across potentially unobserved period-states. Instead, I collapse each time period into a static setting, admitting the result that consumer demand for the risky asset in each period is a sufficient statistic for the welfare effects of a change in the price of the asset. As a result, the intuition is the same as in the static model: on the margin, neoclassical households must be indifferent to small changes in quantity purchased of the asset, so the first order effect of the price change on welfare dominates. Unfortunately, this approach cannot recover exact consumer surplus: Blackorby et al. (1984) prove that this holds in general in a dynamic model, even with standard consumer goods, and Hausman & Newey (2016) demonstrate a non-identification result for exact consumer surplus even in a static setting with exogenous variation in budget constraints. However, Marshallian consumer surplus has the advantage of being straightforward and transparent to calculate while preserving much of the ease of interpretation of exact measures. Additionally, this framework suggests that dynamic elasticities of the demand curve to a permanent unexpected price change conditional on a set of shocks are the key moments to match if one is interested in estimating a structural model of borrowing, saving, and investment, with the objective of measuring the welfare gains from access to a particular technology.

I demonstrate how to calculate the consumer surplus[1] from price changes to a risky asset using this approach by reanalyzing results from 3 papers studying experimental and quasi-experimental variation in prices. In each setting, the key step is to model the intertemporal substitution technology as a risky asset, either purchased (for savings or index insurance) or sold (for credit) by the household; after this, the model based approach can be applied directly. In the first of the 3 papers, Karlan et al. (2014a) randomize prices of index insurance in Ghana to estimate market demand, and additionally compare the effects on households of grants of index insurance to cash grants of the actuarially fair value of the index insurance. If one is concerned exclusively with welfare, one could make this comparison simply by comparing the consumer surplus from index insurance to the actuarially fair value of the grants of index insurance. I demonstrate that on average, households value the index insurance 55% as much as the cash grants; since index insurance pays out after cultivation, while the cash grants pay out before, this is potentially consistent with a high marginal product of investment during cultivation. In the second, Karlan & Zinman (2013) randomize the interest rates on loans from Compartamos, a leading microfinance institution in Mexico. Combining their estimated effects with additional assumptions specific to their context, I produce an estimate of the consumer surplus from access to Compartamos bank of $1 per month per capita in each municipality where Compartamos operates. This effect is economically meaningful, but too small to be detectable using the welfare proxy approaches

[1]For convenience, for the remainder of the paper I use consumer surplus to refer to Marshallian consumer surplus, except when explictly noted. However, note that in contrast to compensating or equivalent variation - which, following Hausman (1981), are measures of exact consumer surplus - Marshallian consumer surplus may not be uniquely defined when multiple prices change, and does not represent the exact consumer surplus of the price change unless the marginal utility of income is constant. I discuss how these points relate to my proposed measure in Section 3.2.3.

that have been used in the literature on the welfare impacts of microcredit. In the third, Duflo et al. (2006) randomize incentives for retirement savings in the United States, and compare the elasticity of the response to quasi-experimental variation in incentives to save from the Savers' Credit. By comparing the two estimated demand curves, I lower bound the lost consumer surplus from inattention to the Savers' Credit, which I estimate to be more than 50% of the total surplus available if savers were not inattentive.

This paper contributes primarily to the literature that measures the welfare gains from new technologies from intertemporal substitution. In this literature it is closest to Einav et al. (2010), who estimate consumer and producer surplus in insurance markets estimating response of demand and claims to variation in prices.[2] The key contribution of this paper is to extend their approach to a setting where concerns of intertemporal substitution, and not just heterogeneity and uncertainty, are first order. In contrast to their work, the literature has primarily relied on a mix of reduced form approaches relying on proxies for welfare (Kaboski & Townsend (2012), Bryan et al. (2014), Banerjee et al. (2015), Breza & Kinnan (2017)) and structural approaches (Kaboski & Townsend (2011), Bryan et al. (2014), Breza & Kinnan (2017)); I argue that a flexible sufficient statistics approach is a valuable complement to the existing literature.[3] Consistent with this, even among studies employing reduced form approaches, many make the point that estimating consumer demand is important for understanding the welfare gains from the technology (Dupas et al. (2016), Karlan & Zinman (2013)), an intuition which I formalize.

The remainder of the paper is organized as follows. Section 3.2 introduces the model, demonstrates how a few simple extensions can add important flexibility and realism, and builds a feasible welfare measure. Section 3.3 presents each of the 3 empirical applications, with each subsection describing the reanalysis applied to each paper, and presenting consumer surplus graphically and numerically. Section 3.4 concludes.

[2] Other similar contributions include Einav & Finkelstein (2011) and Chetty & Finkelstein (2012).

[3] Many existing papers have applied sufficient statistics approaches to measure the welfare gains from an intertemporal substitution technology. To cite a few examples, Chetty (2008) estimates optimal unemployment insurance, Annan (2017) estimates the welfare effects of a policy restricting the sale of auto insurance on credit, and Auclert (2017) estimates heterogeneous effects of changes in asset prices through shocks to interest rates on household wealth. See Chetty (2009) for a review of this literature.

3.2 Model

3.2.1 Setup

The setup closely follows Deaton (1991) and Carroll (2012). Each period, consumers solve

$$V(x, h_t, \rho) = \max_{b_t} u(x - \rho b_t, s_t) + \delta \mathbf{E}_{s_t}\left[V(y(s_{t+1}) + f(b_t, s_{t+1}), h_{t+1}, \rho)|s_t\right] \tag{3.1}$$

subject to the constraint

$$b_t \in [\underline{b_t}(h_t), \overline{b_t}(h_t)] \tag{3.2}$$

x is the consumer's cash on hand at the start of period t. u is the household's utility index; it may depend on the realized state s_t as in Dean & Sautmann (2014), which evolves following a Markov process.[4] b_t is the household's choice variable, the number of units of the risky asset they purchase at price ρ. h_t is the full history of states and asset purchase decisions before period t. $[\underline{b_t}(h_t), \overline{b_t}(h_t)]$ is a closed interval of the real line which constrains choices of b_t. That it depends on h_t enables certain forms of dynamic incentives, a key feature of many credit products. $y(s_{t+1})$ is an income shock, and f is a risky concave production technology. When $u(x - \rho b_t, s_t) \equiv u(x - \rho b_t)$, $f_{t+1}(b_t, s_{t+1}) = b_t$, and s_t has no serial dependence, this model reduces to the Deaton (1991) buffer stock savings model, with ρ as the inverse interest rate factor.

The comparative static of interest is a permanent shock to ρ on welfare, normalized by the marginal utility of income. Let b_t^* be the solution to the household optimization problem in period t. For convenience here, I write $V_t \equiv V(x, h_t, \rho)$, and $u_t \equiv u(x - \rho b_t^*, s_t)$; $V_{t+1}(s_{t+1})$ and $u_{t+1}(s_{t+1})$ are defined analogously. Let $\gamma(s_{t+1}) \equiv \frac{\delta u'_{t+1}(s_{t+1})}{u'_t}$ be the intertemporal marginal rate of substitution associated with s_{t+1}. Then

$$\frac{dV_t/d\rho}{dV_t/dx} = \frac{\partial V_t/\partial \rho}{\partial V_t/\partial x} = \underbrace{-b_t^*}_{\text{"static" gains}} + \underbrace{\mathbf{E}_{s_t}[\gamma(s_{t+1})]}_{\text{discount factor}} \mathbf{E}_{s_t}\left[\underbrace{\frac{\gamma(s_{t+1})}{\mathbf{E}_{s_t}[\gamma(s_{t+1})]}}_{\text{weight}} \underbrace{\frac{\partial V_{t+1}(s_{t+1})/\partial \rho}{\partial V_{t+1}(s_{t+1})/\partial x}}_{\text{"future" gains}}\right] \tag{3.3}$$

The first step, that the total derivatives equal the partial derivatives, is just the envelope theorem, and the second step involves evaluating each derivative. We can break the expression up into four components.

First, the "static" gains is the standard expression one recovers in a static model,

[4] Modeling state as following a Markov process, combined with the dependence of u, y, and f on s_{t+1}, means that this infinite time horizon approach strictly generalizes discrete time horizon approaches; one simply fixes $u = y = f = 0$ by allowing state to collapse to a degenerate distribution for all $t > T$.

quantity purchased of the risky asset. However, the gains are not static in the sense that there is no guarantee of which period-states' contemporaneous utility changes in value. In particular, this means that this expression holds only in expectation; a household with a price elasticity of demand for the risky asset above 1 will see its welfare fall in response to a price decrease, for example, if it encounters an s_{t+1} such that $f(b, s_{t+1}) = 0$.

Second, the discount factor used is the expected intertemporal marginal rate of substitution. To facilitate intuition, if $f(b, s) = b$ and households are unconstrained in their choice of b, then the first order condition implies that the discount factor equals ρ; it will be lower (higher) if their choice of b_t is constrained from below (above). The household discounts the weighted expectation of future gains, where the weights are proportional to each period-state's intertemporal marginal rate of substitution; gains are weighted more heavily in periods in which households are relatively poorer. The "future" gains are recursively defined; they will again contain the four terms in the above expression.

Next, consider a discrete change in prices from ρ_0 to ρ_1. When $\rho_0 = \infty$ and $\rho_1 = \rho$, this represents the introduction of the technology. Let $\text{CS} \equiv \int_{\rho_0}^{\rho_1} \frac{dV_t/d\rho}{dV_t/dx} d\rho$ be the consumer surplus of the price change. As is the case in a static model, this is strictly in between the compensating variation and the equivalent variation of the price change, and therefore represents a useful metric for the consumer welfare effects of the price change when the marginal utility of income does not change too much as ρ changes. One can show that there exist $\gamma^*(s_{t+1})$, where $\gamma^*(s_{t+1}) \in \text{ConvexHull}(\{\gamma(s_{t+1})|\rho \in [\rho_1, \rho_0]\})$, such that

$$\text{CS} = \int_{\rho_1}^{\rho_0} b_t^* d\rho + \mathbf{E}_{s_t}[\gamma^*(s_{t+1})]\mathbf{E}_{s_t}\left[\frac{\gamma^*(s_{t+1})}{\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]} \int_{\rho_0}^{\rho_1} \frac{\partial V_{t+1}(s_{t+1})/\partial\rho}{\partial V_{t+1}(s_{t+1})/\partial x} d\rho\right] \tag{3.4}$$

The first term is now the standard welfare trapezoid we get in the static case, and the second term is the discounted weighted average of consumer surplus in future period-states.

3.2.2 Extensions

The model can be extended in a number of key ways. First, to see how the model can not be extended, note that the key assumption underlying the envelope theorem trick is that the only channel through which ρ entered the household's decision is through its budget constraint. In many realistic settings, this is violated to a minor degree. For example, when borrowing is constrained, in practice the constraint is frequently defined in terms of ρb_t, not in terms of b_t. In this case, there is an additional effect through the change in the constraint; however, this is likely to be second order. However, it can also be violated severely, such as in contexts where moral hazard is directly affected by ρ (as Karlan & Zinman (2009) find in South Africa).

First, a menu of technologies can be permitted. The increases in the quantity of the risky asset b_t may substitute or complement other risky assets available to the household. More generally, a more flexible income process can be implemented, where the next period's income can depend flexibly on the full history of past purchases and sales of the full menu of risky assets and a multidimensional random shock; among other things, this flexibly allows for investments to pay out stochastically over the course of multiple periods. Additionally, any non-convexities in these technologies are permitted.

Second, heterogeneity across households is permitted; in this case, aggregate consumer surplus is the sum of consumer surplus across households.[5] This adds realism, and with a continuum of households it potentially generates smooth aggregate demand for b_t despite household non-convexities, which can simplify estimation. To cite one example, Kaboski & Townsend (2011) study demand for credit in rural Thailand with a generalization of the buffer stock savings model. Additionally they allow for households to make discrete investments; however, with a continuum of households, despite this nonconvexity their model still results in smooth aggregate demand for credit. In fact, their model is fully nested within this framework, with the exception of how they allow default. Despite this, their model is likely well approximated within this framework, since certain forms of default are allowed in this framework through the dependence of f, $\underline{b_t}$, and $\overline{b_t}$ on h_t.

Third, this approach allows for certain forms of general equilibrium effects. If markets are perfectly competitive, then an application of the first welfare theorem allows converting the general equilibrium problem to a planner's problem, after which the envelope theorem can once again be applied. As an example, Wright & Williams (1984) calculate the welfare effects of the introduction of competitive storage in a closed economy with competitive output markets. In this framework, the joint distribution by period-state of γ and demand for storage as a function of the cost of storage is sufficient to calculate the surplus from the introduction of the technology. However, a cost of using the approach in general equilibrium is that an increase in consumer surplus for a single household does not necessarily correspond to an increase in utility for that household; a share of that surplus may be passed on to other households through price effects, to cite one possibility.

3.2.3 A feasible measure

From Equation 3.4, we can see that aggregate consumer surplus will depend on the joint distribution, for each period-state, of marginal utility of consumption and demand

[5] Moreover, with complete insurance across households, $\gamma(s_{t+1})$ will be equal across households, a point made by Ligon et al. (2002) among others; this would permit summing within period-state before summing across period-states.

for the risky asset as a function of the price of the risky asset. Without imposing additional structure on the problem, this is infeasible; however, this statement holds more generally for the estimation of consumer surplus even for traditional consumer goods. In particular, note that estimates of consumer surplus for consumer goods using static models, using either exact or Marshallian approaches, will be biased when consumption of those goods is correlated with the marginal utility of consumption within household. Despite this, what I will call flow or per-period consumer surplus is still a potentially useful metric of the welfare effects of the price change, and a well estimated structural model potentially allows the decomposition of differences between structural estimates of exact consumer surplus and reduced form estimates of flow Marshallian consumer surplus into a discount factor, correlation between the marginal utility of consumption and consumption of the good within household, and income effects of the price change on consumption of the good. In this sense, estimates of flow Marshallian consumer surplus complement, as opposed to substitute, other approaches used to estimate the welfare gains from a new consumer good.

I argue similarly for the use of flow consumer surplus[6] as a metric for the welfare effects of the price change for a risky asset, which I define to be

$$CS_{t+k} \equiv \int_{\rho_1}^{\rho_0} b_{t+k}^* d\rho \tag{3.5}$$

which is feasible to estimate. In particular, demand for the risky asset in period $t+k$ is sufficient for CS_{t+k}, which can be estimated from exogenous permanent shocks to ρ in period t.

As discussed above, this measure is imperfect relative to traditional Marshallian consumer surplus from Equation 3.4 in two senses. First, we are discarding $\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]$, which may vary across households. However, we can potentially bound this term; to give one simple example, if households are never constrained, and $f(b_t, s_{t+1}) = b_t$, then we know from the first order condition that $\mathbf{E}_{s_t}[\gamma^*(s_{t+1})] \in [\rho_1, \rho_0]$.[7] Second, we are discarding the weights $\frac{\gamma^*(s_{t+1})}{\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]}$. To understand the impact of this, consider the Deaton (1991) saving model. In this model, $\frac{\gamma^*(s_{t+1})}{\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]}$ will be highest in periods where the households receives a negative income shock, and as a result saves less. Because of this, $\int_{\rho_0}^{\rho_1} \frac{\partial V_{t+1}(s_{t+1})/\partial \rho}{\partial V_{t+1}(s_{t+1})/\partial x} d\rho$ will be lower in periods where $\frac{\gamma^*(s_{t+1})}{\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]}$ is high, which will bias this approach's measurement upwards. In contrast, with a credit technology, the measures from this approach will be biased downwards. Third, if shocks are correlated across households, or if households have biased beliefs, then effectively the realized distribution of states may be different from the anticipated distribution of states. As

[6] Hereafter, I will use consumer surplus to refer to flow Marshallian consumer surplus for convenience, which I contrast with traditional consumer surplus measures.

[7] A similar argument can be made to bound $\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]$ for constrained households, either from above or from below. Other investments made by the household can potentially refine these bounds.

a result, households may happen to save into periods in which their marginal utility from consumption is high, or borrow from periods in which their marginal utility from consumption is low. In that sense, we should interpret this measure as anticipated consumer surplus from the perspective of period $t + k$.[8]

3.3 Empirical applications

To make concrete the applications of this approach, I reanalyze results from 3 experiments on savings, credit, and insurance which randomly vary prices or interest rates. The first is Karlan et al. (2014a) (hereafter KOOU), who compare the effects of index insurance and cash grants on agricultural investment in Ghana, and randomly vary the prices of index insurance to estimate demand. I compare the consumer surplus of grants of index insurance to cash grants equal to the actuarially fair value of the index insurance. The second is Karlan & Zinman (2013) (hereafter KZ), who study the long run effects of a change in microfinance interest rates by Compartamos, a Mexican MFI. I use their experiment to estimate the consumer surplus from Compartamos, and I compare this estimate to estimates from the literature measuring the welfare impacts of microfinance using RCTs. The third is Duflo et al. (2006) (hereafter DGLOS), who study retirement savings in the United States, and demonstrate that households respond much less to the Savers' Credit, effectively a subsidy for retirement savings, than to more salient but equivalent deposit matching treatments. I use their estimates to lower bound the foregone consumer surplus due to inattention to the Savers' Credit.

3.3.1 Index insurance: KOOU

Empirical strategy

KOOU study and compare the effects of index insurance and cash grants on farmer investment decisions and welfare. To do so, they conduct two sets of experiments. In the first, they randomly assign either grants of index insurance or cash grants equal to the actuarially fair value of the index insurance to households. In the second, they randomly vary the prices of index insurance to estimate demand. In particular, I focus on two price experiments that they run in the second year of the study. In the first, a random subset of households from what they call Sample Frames 1 and 2 were randomly offered the opportunity to buy an index insurance product at 12% or 50% of

[8] A similar caveat may be made for behavioral households. In all these scenarios, as is typical in models with internalities, the bias in the measure will be driven by the gap between anticipated $\mathbf{E}_{s_t}[\gamma^*(s_{t+1})]$ and the true expectation of the realized distribution of $\gamma^*(s_{t+1})$, along with the price elasticity of demand for the risky asset.

the actuarially fair price, with prices randomized across communities. In the second, in what they call Sample Frame 3, similar in characteristics but not drawn from the same population as Sample Frame 2, communities were offered index insurance with prices randomized across communities to 100% or 150% of the actuarially fair price. Although 150% is above actuarially fair price, it was chosen to reflect the market price typically offered for index insurance.

To proceed, I make the assumption that the demand for the index insurance product is the same in both experiments. This is not necessarily the case, since the products and the communities are not identical; however, with a functional form assumption on demand this is testable, and I impose linearity. Additionally, I bound quantity demanded of index insurance from above, since households were prohibited from purchasing index insurance that covers more than their total landholdings. Finally, with the full demand curve estimated, I compare the actuarially fair value of the insurance grants to the consumer surplus of insurance grants.

To simplify comparison, prices of actuarially fair insurance are normalized to 1, so quantities at actuarially fair prices are expenditures, and quantities are normalized to be per household.[9] Having heterogeneity in the product across households is not a problem, since this is flexibly modeled by variation in f. However, the assumption that demand is the same in the lower price experiment (with Sample Frame 1 and 2) and in the higher price experiment (with Sample Frame 3) is nontrivial, but is used primarily for extending the demand curve estimated using Sample Frame 3 to lower prices.

Results

Figure 3.1 plots the fitted demand from the price experiment. A formal statistical test rejects a linear fit,[10] however the degree to which this assumption could potentially bias the exercise is easy to see from the figure. The consumer surplus from grants of index insurance is A + B (4.4 GHC/acre), while grants of value equal to the actuarially fair price of the index insurance is B + C (8 GHC/acre). Roughly, on average, households would be indifferent between receiving 4.4 GHC/acre and receiving the grants of index insurance that pays out 8 GHC/acre on average after shocks are realized, despite the fact that the index insurance is more likely to pay out when households have experienced a negative shock. This is potentially consistent with the fact that the index insurance

[9] Across the households from the price experiments I focus on, the actuarially fair price of the index insurance varied from 7.65 GHC/acre to 9.5 GHC/acre, and prices in the experiment were set based on that. I work through this exercise as if all of them were priced at 8 GHC/acre, but the results of the exercise do not meaningfully change with alternative approaches to calculating prices and quantities to enable comparison across households that received products priced slightly differently.

[10] The estimated slopes and intercepts from the two price experiments are statistically significantly different.

does not pay out until after cultivation is already completed, while households purchase the index insurance product before cultivation, a period during which households likely have a high marginal product of capital unless they have access to low interest credit (Casaburi & Willis (2016)).

3.3.2 Microcredit: KZ

Empirical strategy

KZ study the effects of a decrease in interest rates by Compartamos on demand for microcredit in Mexico. 132 branches were randomly assigned to tiered monthly pricing of either 3.0%/3.5%/4.0% or 3.5%/4.0%/4.5% interest. In all cases, this was a decrease from the 4.0%/4.5%/5.0% interest offered at baseline, which was comparable to the terms offered by competitors. The experiment lasted 29 months, and was perceived by households to be a permanent change, allowing estimates of demand from 1 to 29 months out from the shock. Loan terms were 4 months, monthly interest payments are calculated based on the initial loan amount, and 15% VAT is charged on top of each month's loan payment.[11] Letting b_t represent a single week's payment, then in the model $\rho b_t(1 + (1.15)16r) = 16b_t$; the income the household receives for selling the risky asset (the sequence of future payments to Compartamos) times 1 + the share of this income paid as interest equals the total payments. As a result, the relevant $\rho = \frac{16}{1+18.4r}$, or 10.31, 9.73, 9.22, and 8.75 for 3.0%, 3.5%, 4.0%, and 4.5% respectively. Comparing the ratios in equivalent tiers in the treatment and control group, the intervention can be thought of approximately as a 5.5% increase in ρ.

Note that in this case, an important aspect of the model is that f is allowed to depend on h_{t+1} and not just s_{t+1}; past borrowing and realizations of state can influence future interest rates. However, this assumes away many forms of moral hazard, in particular interest rates cannot affect repayment probabilities except through selection. This is potentially reasonable for Compartamos clients, for whom default rates are about 1%. I also ignore an additional requirement that the household leaves 10% of its initial loan balance in a savings account with Compartamos. The initial loan balance is ρb_t, so this means the future sequence of payments depends partially on ρ (since the household receives $0.1\rho b_t$ after 16 weeks of payments). However, this introduces a negligible amount of bias into this approach. Additionally, KZ note that, anecdotally, this requirement is not vigorously enforced.

To estimate consumer surplus, of the price change, I use KZ results on the effects of the reduced interest rates on loan amount. I normalize prices to 1 in the high interest

[11] Additionally, households must deposit 10% of the initial loan balance with Compartamos, which is released after the 4 months.

rate group, so loan amounts are quantities, while loan amounts divided by ρ in the low interest group (1.055) are quantities in that group. Following their Table 3, monthly demand curves are plotted separately for each year. As a result, the demand curve can be interpreted as average monthly demand per branch, and one can multiply by 132 to get total demand for Compartamos if a particular price prevailed at all of the experimental branches.

Finally, to estimate total consumer surplus of Compartamos, I assume that demand is linear until $\rho = 0.95$, roughly the prevailing market price at other MFIs, beyond which point it falls to 0. This would be consistent with the patterns of demand for credit found by Karlan & Zinman (2008), who find that price elasticities spike above market rates. As a test of linearity, I compare the predicted borrowing at $\rho = 0.95$ to borrowing before the experiment, at which point $\rho = 0.95$ for Compartamos loans.

Results

Figure 3.2 graphs demand curves for credit from Compartamos, which increase over each of the 3 years of the experiment, with quantities normalized such that prices (in USD) are 1 in the high interest rate arm. First, note that the linear model appears to perform well, with predicted borrowing at the market interest rate roughly constant and equal to the baseline level of borrowing (the red dot), when Compartamos had $\rho = 0.95$. Calculating monthly consumer surplus of Compartamos loans in each year of the experiment simply requires calculating the area of the trapezoid (A for year 1, A+B for year 2, and A+B+C for year 3) and multiplying by 132 (the number of experimental branches).[12] This yields a monthly consumer surplus of 5.3 million USD, 6.3 million USD, and 7.5 million USD in year 1, 2, and 3, respectively. This is approximately 9% of the volume of loans disbursed at low interest rates (which is easy to see from Figure 3.2).[13] Alternatively phrased, on average, clients would have been indifferent between receiving 9% less balance today, but making the same payments, and not taking out the loan at all. This is about 30 USD/loan, or about 1 USD/capita/month in the branch's municipality.

Although this estimate is small, it is potentially meaningful. A typical borrower's household from their sample earns about 300 USD/month, so 30 USD surplus from a 4 month loan represents 2.5% of a typical household's income over that period. As an alternative to this approach, Angelucci et al. (2015) use a randomized rollout of branches

[12] Note that this is under the counterfactual that Compartamos offered the lower interest rate to all consumers. I consider this the relevant counterfactual since following the experiment, Compartamos rolled this lower interest rate out to all its branches.

[13] Note that this is driven to a large degree by the assumption that demand is perfectly elastic at above market interest rates, and relaxing this assumption would increase this percentage; therefore, one could alternatively interpret this result as a lower bound.

of Compartamos to estimate welfare effects of microcredit. Consistent with much of the microfinance literature, they fail to find significant effects on consumption. However, their standard error on the effect on food expenditures, to choose one commonly used category, is 3 USD/household/month. Due to moderate take up of microcredit, it is challenging to sufficiently power a cluster RCT to estimate the welfare impacts of microcredit, even aggregating across multiple RCTs.[14] In contrast, the standard error on the consumer surplus estimate (1 USD/capita/month) is closer to 0.2 USD/capita/month, since demand can be estimated much more precisely with sufficient variation in prices.

3.3.3 Savings: DGLOS

Empirical strategy

DGLOS study the effects of inattention to the Savers' Credit, a subsidy for retirement savings in the United States, on household savings. First, they use quasi-experimental variation in eligibility for the Savers' Credit to document small responses to what is effectively a 100% match for retirement savings.[15] In contrast, in an experiment with H&R Block in which households were randomly assigned to receive either 0%, 20%, or 50% matches for retirement savings, they document much more elastic responses to the more salient experimental matches. Although this dynamic setting is ostensibly more complicated than static settings in which inattention is studied, such as Chetty et al. (2009), this framework permits using the same tools from behavioral public finance to analyze both settings. To produce the policy relevant demand curves, I fit linear (with respect to ρ) demand curves to their estimates of the effects of matches and the Savers' Credit. I make the assumption that the elasticity to matched contributions estimated from the experiment is bias free, and use this to construct a bias free demand curve for retirement savings deposits, which I contrast with the demand curve affected by consumer bias estimated using quasi-experimental variation from the Savers' Credit. With the two demand curves, we can calculate the potential consumer surplus under the unbiased demand curve, and put a lower bound on the potential foregone surplus caused by inattention to the Savers' Credit.

[14] See Meager (2015) for a metaanalysis of 7 RCTs evaluating the effects of microcredit. Similarly, 95% intervals for their posterior distributions of the average effect of microcredit on monthly consumption cover 0 to 25 USD/household/month.

[15] DGLOS also study households that receive effective 25% and 12.5% matches; I drop them from this analysis because the estimated slope of demand using these households is significantly less precise.

Results

The estimated demand curves, along with the average deposits conditional on ρ used to estimate the demand curve, are plotted in Figure 3.3.[16] The shaded area A, 20 USD/year/household, is a lower bound on the foregone consumer surplus caused by inattention to the Savers' Credit per year per household. In comparison, the total area under the demand curve estimated using the match experiment is 37 USD/year/household; at least 53% of this potential consumer surplus is foregone due to inattention to the Savers' Credit.

Although 20 USD/year seems small for the United States, this comes with 3 caveats. First, this eligible sample is a relatively poor sample of households, with annual incomes less than 30,000 USD. Second, a relatively small percentage of sample households save, even with a match; in the experimental (quasi-experimental) sample, 2.5% (2.1%) saved without a match, while 14.0% (3.3%) saved with a 50% match (100% match). Since this surplus is 0 for households that would not have saved even if they were attentive to the Savers' Credit, one could calculate a conditional consumer surplus by scaling 20 USD/year/household up by the inverse of share of households that would have saved if they were attentive to the Savers' Credit if that share was observable. Third, households retirement savings may be suboptimal due to present bias or other behavioral biases. As a result, changes in savings induced by price changes may have first order welfare effects that this approach does not take into account.

3.4 Conclusion

This paper argues that demand for intertemporal substitution technologies is almost sufficient for the exact consumer surplus from a price change to the technology under a broad class of neoclassical models. This suggests that tools traditionally used in the analysis of static demand for consumer goods may also be used to understand demand for savings and credit.

A few key insights from this approach to understanding demand for intertemporal substitution technologies are particularly important. First, structural exercises estimating the welfare gains from an intertemporal substitution technology should plot demand as a function of a permanent price shock, as in Figure 3.2. The welfare effects of the technology estimated from the structural exercise will be similar to the area under this demand curve, so this provides both a convenient visualization of their estimates and

[16] Average deposits were higher for the experimental sample in the year of the experiment; I divide experimental deposits by a constant multiple to recover these points. This approach is consistent with the assumption that the experiment recovers the bias-free elasticity.

an overidentification test based on the reasonableness of the elasticity. Second, RCTs of these technologies can bound neoclassical welfare gains from data on take-up alone, and can potentially produce relatively tight bounds with reasonable assumptions on the shape of demand. Additionally, these demand based bounds will often be significantly more precise than what is feasible from estimates based on welfare proxies. Third, large deviations between estimated consumer surplus from this approach and money metric equivalents to reduced form effects on welfare proxies (such as business profits or consumption) are unlikely to be explained by neoclassical models. As an example, this insight is helpful to understand the result from Bryan et al. (2014), that implausible levels of risk aversion are needed to rationalize the large responses to small incentives to migrate and the large effects of migration on consumption. This is because the consumer surplus from the subsidies is small (since the subsidies are small), but the average effect on a money metric welfare proxy is large, something that can only be rationalized with a small discount factor, which they rule out since households are observed to save, or a large degree of curvature in utility. Finally, and importantly, Marshallian consumer surplus approaches admit simple graphical illustrations of welfare. Dynamic models often require complicated numerical approaches to estimation, but these graphical representations facilitate transparency of how welfare estimates are constructed and accessibility of results to more general audiences.

Figure 3.1: Index insurance consumer surplus

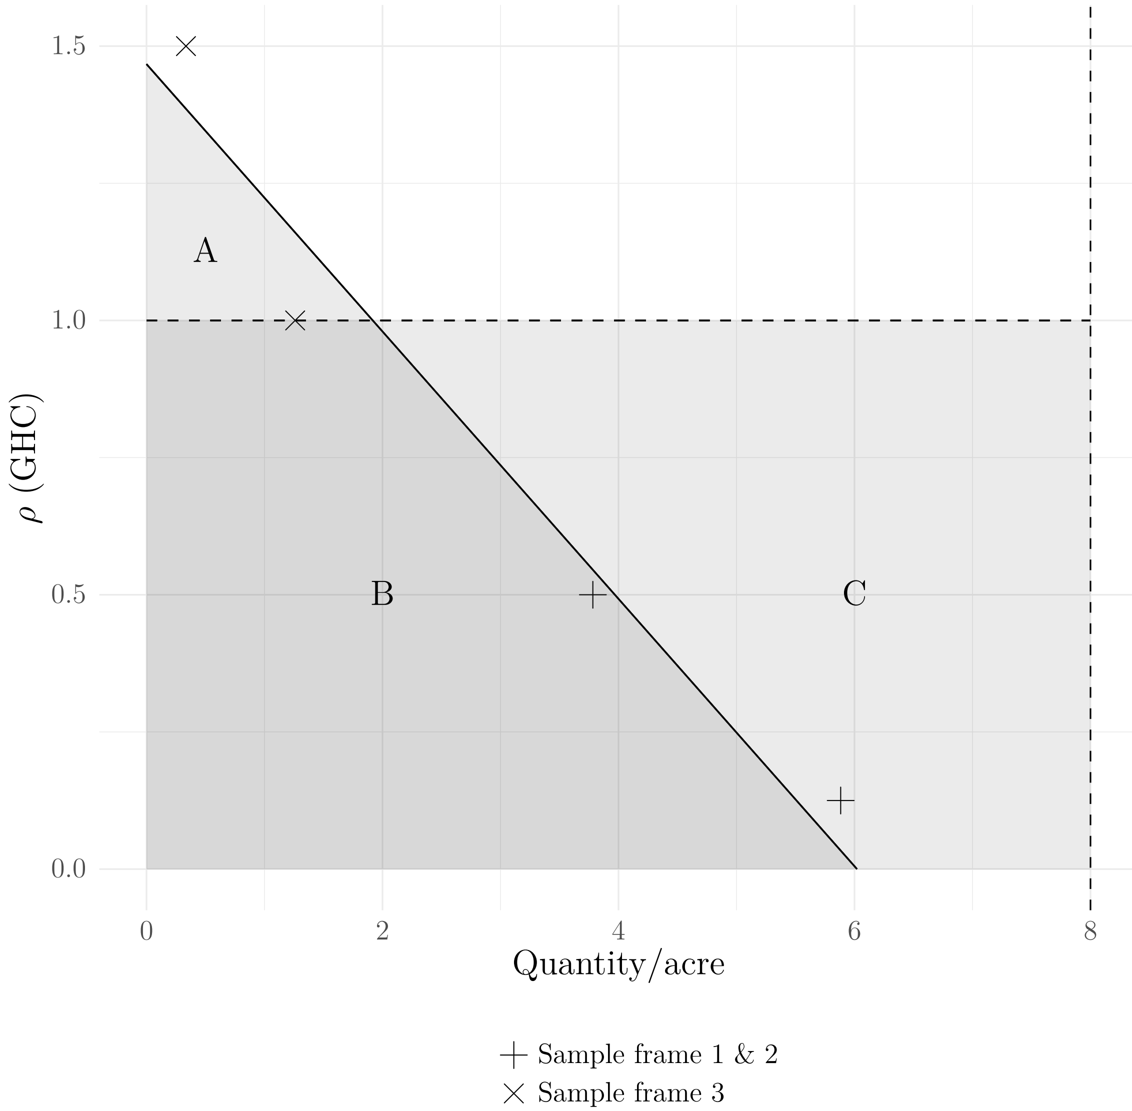

Figure 3.2: Monthly consumer surplus from Compartamos loans

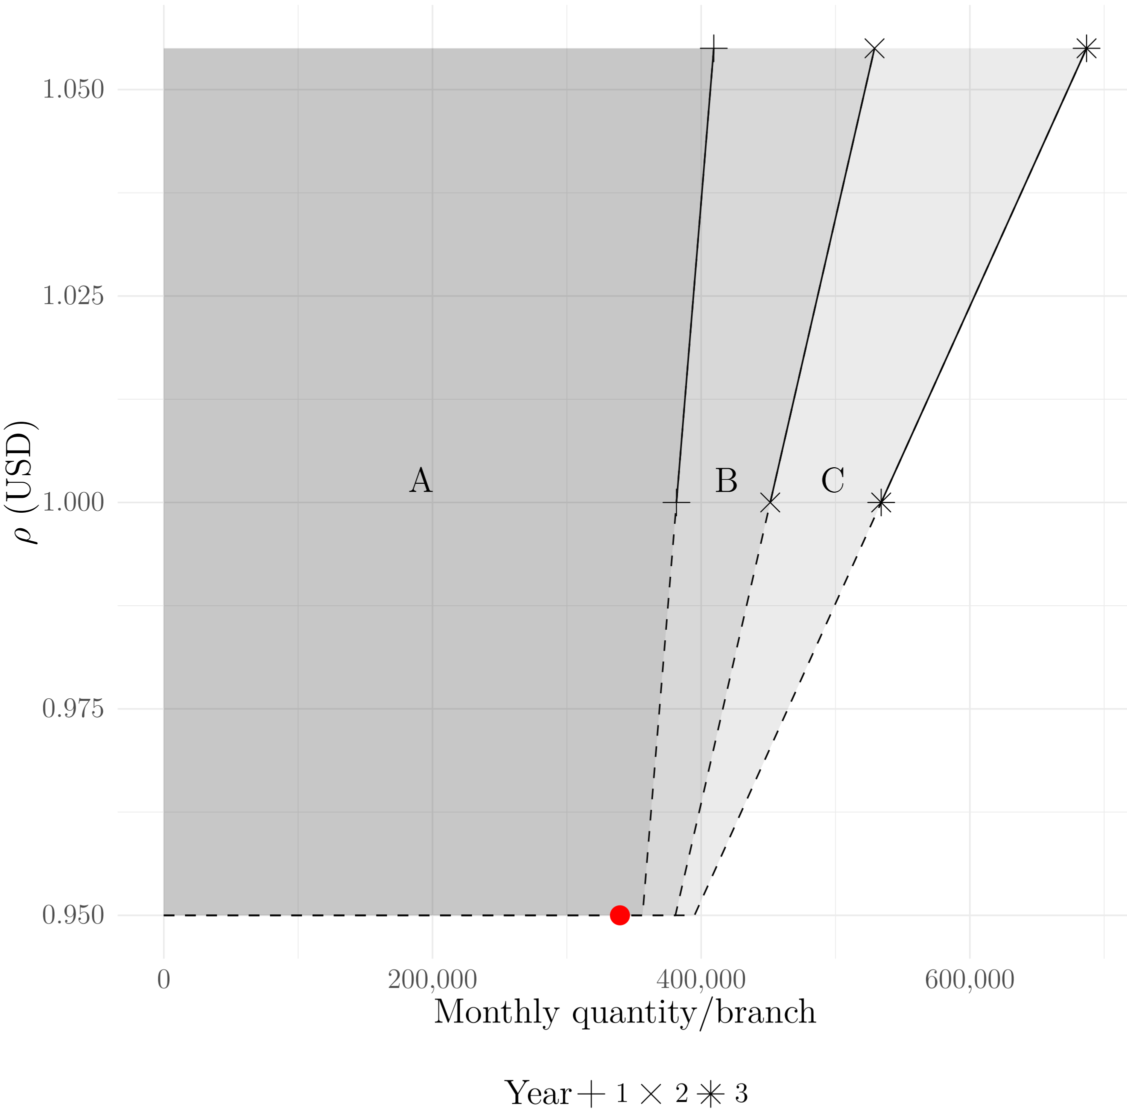

The red dot is Year 0 average quantity and price across all experimental branches.

Figure 3.3: Biased and debiased retirement savings demand

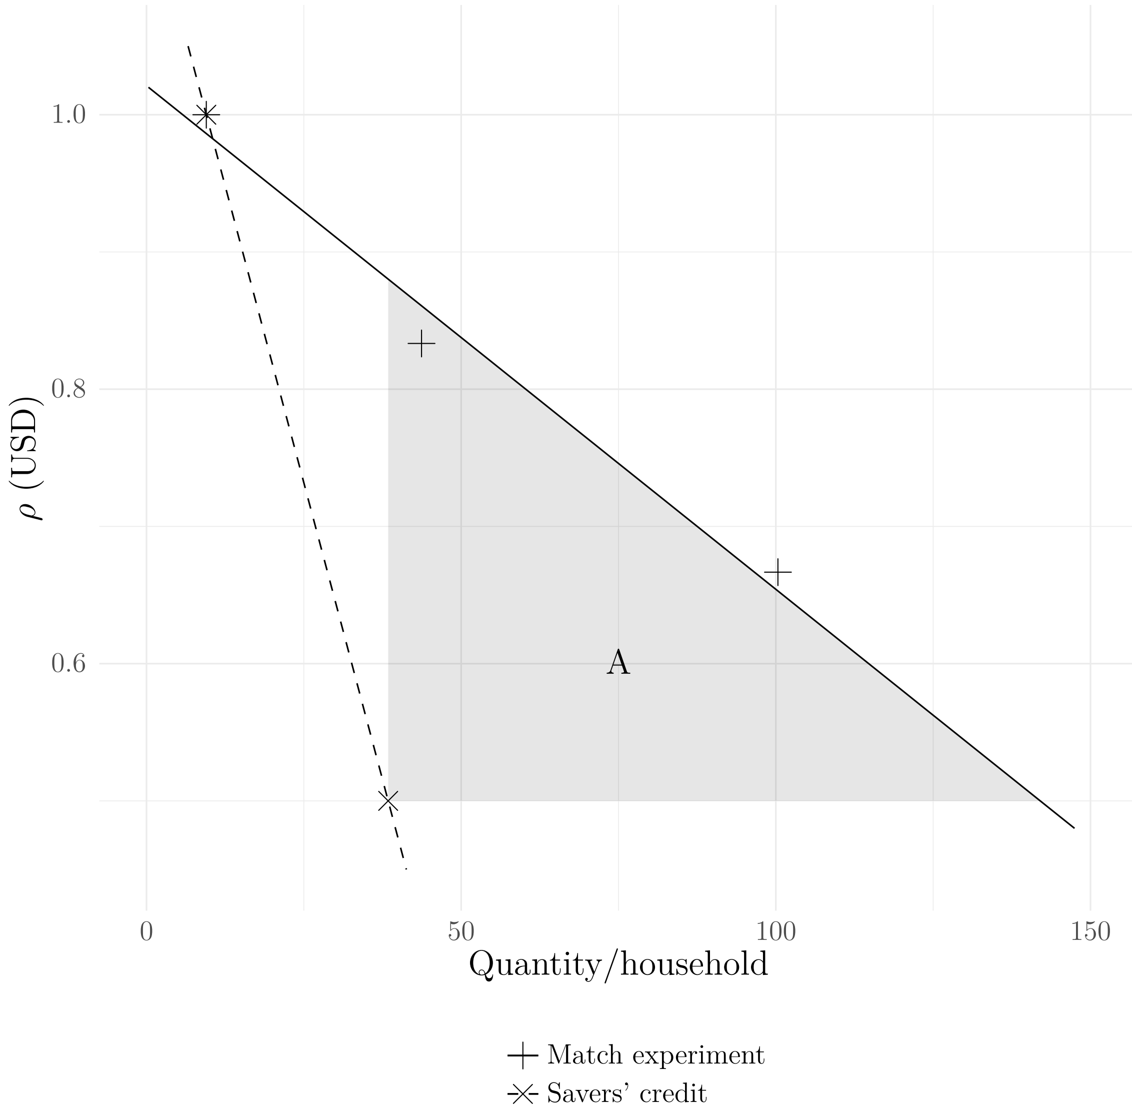

Bibliography

Acemoglu, D. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *The Quarterly Journal of Economics*, 113(4), 1055–1089.

Adão, R. (2016). Worker heterogeneity, wage inequality, and international trade: Theory and evidence from brazil.

Ahn, H. & Powell, J. L. (1993). Semiparametric estimation of censored selection models with a nonparametric selection mechanism. *Journal of Econometrics*, 58(1-2), 3–29.

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.

Allcott, H., Mullainathan, S., & Taubinsky, D. (2014). Energy policy with externalities and internalities. *Journal of Public Economics*, 112, 72–88.

Allen, T. & Atkin, D. (2015). Volatility, insurance and the gains from trade.

Anderson, M. L. (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.

Andrews, I., Gentzkow, M., & Shapiro, J. M. (2018). On the informativeness of descriptive statistics for structural estimates.

Angelucci, M., Karlan, D., & Zinman, J. (2015). Microcredit impacts: Evidence from a randomized microcredit program placement experiment by compartamos banco. *American Economic Journal: Applied Economics*, 7(1), 151–182.

Angrist, J. & Fernandez-Val, I. (2010). *Extrapolate-ing: External validity and overidentification in the late framework*. Technical report, National Bureau of Economic Research.

Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434), 444–455.

Angrist, J. D. & Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Annan, F. (2017). Credit-induced moral hazard in insurance: Theory and evidence from Ghana.

Arnold, D., Dobbie, W., & Yang, C. S. (2018). Racial bias in bail decisions. *The Quarterly Journal of Economics*.

Auclert, A. (2017). Monetary policy and the redistribution channel.

Badiani, R. & Jessoe, K. (2017). Electricity prices, groundwater and agriculture: the environmental and agricultural impacts of electricity subsidies in india.

Banerjee, A., Karlan, D., & Zinman, J. (2015). Six randomized evaluations of microcredit: Introduction and further steps. *American Economic Journal: Applied Economics*, 7(1), 1–21.

Barrett, C. B. & Bevis, L. E. (2015). The self-reinforcing feedback between low soil fertility and chronic poverty. *Nature Geoscience*, 8(12), 907.

Beaman, L., BenYishay, A., Magruder, J., & Mobarak, A. M. (2018). Can network theory based targeting increase technology adoption.

Beaman, L., Karlan, D., Thuysbaert, B., & Udry, C. (2014). Self-selection into credit markets: Evidence from agriculture in mali. *NBER Working Paper 20387*.

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).

Benjamini, Y. & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, (pp. 289–300).

BenYishay, A. & Mobarak, A. (2014). Social learning and communication. *NBER Working Paper 20139*.

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.

Bhattacharya, D. (2017). *Empirical welfare analysis for discrete choice: Some General Results*. Technical report, Cambridge University working paper.

Björklund, A. & Moffitt, R. (1987). The estimation of wage gains and welfare gains in self-selection models. *The Review of Economics and Statistics*, (pp. 42–49).

Blackorby, C., Donaldson, D., & Moloney, D. (1984). Consumer's surplus and welfare change in a simple dynamic model. *The Review of Economic Studies*, 51(1), 171–176.

Bleakley, H. & Ferrie, J. (2014). Land openings on the georgia frontier and the coase theorem in the short-and long-run.

Bold, T., Kaizzi, K. C., Svensson, J., & Yanagizawa-Drott, D. (2017). Lemon technologies and adoption: measurement, theory and evidence from agricultural markets in uganda. *The Quarterly Journal of Economics*, 132(3), 1055–1100.

Breza, E. & Kinnan, C. (2017). Measuring the equilibrium impacts of credit: Evidence from the indian microfinance crisis.

Breza, E., Krishnaswamy, N., & Kaur, S. (2018). Scabs: The social suppression of labor supply.

Brinch, C. N., Mogstad, M., & Wiswall, M. (2017). Beyond late with a discrete instrument. *Journal of Political Economy*, 125(4), 985–1039.

Brozović, N., Sunding, D. L., & Zilberman, D. (2010). On the spatial nature of the groundwater pumping externality. *Resource and Energy Economics*, 32(2), 154–164.

Bryan, G., Chowdhury, S., & Mobarak, A. M. (2014). Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. *Econometrica*, 82(5), 1671–1748.

Burke, M. & Lobell, D. B. (2017). Satellite-based assessment of yield variation and its determinants in smallholder african systems. *Proceedings of the National Academy of Sciences*, 114(9), 2189–2194.

Bustos, P., Caprettini, B., & Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from brazil. *American Economic Review*, 106(6), 1320–65.

Cai, J., De Janvry, A., & Sadoulet, E. (2015). Social networks and the decision to insure. *American Economic Journal: Applied Economics*, 7(2), 81–108.

Carneiro, P., Heckman, J. J., & Vytlacil, E. J. (2011). Estimating marginal returns to education. *American Economic Review*, 101(6), 2754–81.

Carroll, C. (2012). Theoretical foundations of buffer stock saving.

Carter, M. R., Laajaj, R., & Yang, D. (2013). The impact of voucher coupons on the uptake of fertilizer and improved seeds: evidence from a randomized trial in mozambique. *American Journal of Agricultural Economics*, 95(5), 1345–1351.

Casaburi, L. & Willis, J. (2016). Time vs. state in insurance: Experimental evidence from contract farming in kenya.

Cattaneo, M. D., Idrobo, N., & Titiunik, R. (2018). A practical introduction to regression discontinuity designs: Volume ii.

Chatterjee, S., Lamba, R., & Zaveri, E. (2017). The water gap: Environmental effects of agricultural subsidies in india.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2016). Double/debiased machine learning for treatment and causal parameters. *arXiv preprint arXiv:1608.00060*.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of political Economy*, 116(2), 173–234.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1), 451–488.

Chetty, R. & Finkelstein, A. (2012). Social insurance: Connecting theory to data.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *The American Economic Review*, 99(4), 1145–1177.

Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R., & Vickery, J. (2013). Barriers to household risk management: Evidence from india. *American Economic Journal: Applied Economics*, 5(1), 104–35.

Cole, S. A. & Fernando, A. (2016). 'mobile'izing agricultural advice: Technology adoption, diffusion and sustainability. *Harvard Business School Finance Working Paper*, (13-047).

Conley, T. G. & Udry, C. R. (2010). Learning about a new technology: Pineapple in ghana. *The American Economic Review*, (pp. 35–69).

Costinot, A., Donaldson, D., & Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, 124(1), 205–248.

Crépon, B., Devoto, F., Duflo, E., & Parienté, W. (2015). Estimating the impact of microcredit on those who take it up: Evidence from a randomized experiment in morocco. *American Economic Journal: Applied Economics*, 7(1), 123–50.

D'Agostino, A. L. (2017). Technical change and gender wage inequality: Long-run effects of india's green revolution.

Das, M., Newey, W. K., & Vella, F. (2003). Nonparametric estimation of sample selection models. *The Review of Economic Studies*, 70(1), 33–58.

Dean, M. & Sautmann, A. (2014). Credit constraints and the measurement of time preferences. *Available at SSRN 2423951*.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica*, 59(5), 1221–1248.

Deininger, K. & Feder, G. (2009). Land registration, governance, and development: Evidence and implications for policy. *The World Bank Research Observer*, 24(2), 233–266.

Dell, M. (2010). The persistent effects of peru's mining mita. *Econometrica*, 78(6), 1863–1903.

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, 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.

Dubash, N. K. (2007). The electricity-groundwater conundrum: Case for a political solution to a political problem. *Economic and Political Weekly*, (pp. 45–55).

Dube, A., Lester, T. W., & Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *The review of economics and statistics*, 92(4), 945–964.

Duflo, E., Gale, W., Liebman, J., Orszag, P., & Saez, E. (2006). Saving incentives for low-and middle-income families: Evidence from a field experiment with h&r block. *The Quarterly Journal of Economics*, 121(4), 1311–1346.

Duflo, E. & Pande, R. (2007). Dams. *The Quarterly Journal of Economics*, 122(2), 601–646.

Dupas, P., Karlan, D., Robinson, J., & Ubfal, D. (2016). Banking the unbanked? evidence from three countries.

Einav, L. & Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *The Journal of Economic Perspectives*, 25(1), 115–138.

Einav, L., Finkelstein, A., & Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *The quarterly journal of economics*, 125(3), 877–921.

Eisenhauer, P., Heckman, J. J., & Vytlacil, E. (2015). The generalized roy model and the cost-benefit analysis of social programs. *Journal of Political Economy*, 123(2), 413–443.

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.

Evenson, R. E. & Gollin, D. (2003). Assessing the impact of the green revolution, 1960 to 2000. *science*, 300(5620), 758–762.

FAO (2019). Faostat database. Data retrieved February 15, 2019 from `http://www.fao.org/faostat/`.

Fishman, R., Jain, M., & Kishore, A. (2017). When water runs out: Adaptation to gradual environmental change in indian agriculture.

Fishman, R., Lall, U., Modi, V., & Parekh, N. (2016). Can electricity pricing save india's groundwater? field evidence from a novel policy mechanism in gujarat. *Journal of the Association of Environmental and Resource Economists*, 3(4), 819–855.

Foster, A. D. & Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annu. Rev. Econ.*, 2(1), 395–424.

French, E. & Taber, C. (2011). Identification of models of the labor market. In *Handbook of Labor Economics*, volume 4 (pp. 537–617). Elsevier.

Gandhi, V. P. & Bhamoriya, V. (2011). Groundwater irrigation in india: Growth, challenges, and risks.

Gine, X. & Jacoby, H. G. (2016). *Markets, contracts, and uncertainty in a groundwater economy*. The World Bank.

Giné, X. & Yang, D. (2009). Insurance, credit, and technology adoption: Field experimental evidence from malawi. *Journal of development Economics*, 89(1), 1–11.

Gisser, M. & Sanchez, D. A. (1980). Competition versus optimal control in groundwater pumping. *Water resources research*, 16(4), 638–642.

Goldstein, M., Houngbedji, 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*.

Government of India (2018). Economic survey of india (2017-18).

Hagerty, N. (2018). Liquid constrained: Estimating the potential gains from water markets.

Hausman, J. A. (1981). Exact consumer's surplus and deadweight loss. *The American Economic Review*, 71(4), 662–676.

Hausman, J. A. & Newey, W. K. (2016). Individual heterogeneity and average welfare. *Econometrica*, 84(3), 1225–1248.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1), 153–161.

Heckman, J. J., Urzua, S., & Vytlacil, E. (2006). Understanding instrumental variables in models with essential heterogeneity. *The Review of Economics and Statistics*, 88(3), 389–432.

Heckman, J. J. & Vytlacil, E. (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica*, 73(3), 669–738.

Heckman, J. J. & Vytlacil, E. J. (1999). Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the national Academy of Sciences*, 96(8), 4730–4734.

Heckman, J. J. & Vytlacil, E. J. (2007a). Econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics*, 6, 4779–4874.

Heckman, J. J. & Vytlacil, E. J. (2007b). Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. *Handbook of econometrics*, 6, 4875–5143.

Heisey, P. W. & Norton, G. W. (2007). Fertilizers and other farm chemicals. *Handbook of agricultural economics*, 3, 2741–2777.

Hendren, N. (2014). *Efficient Welfare Weights*. Technical report, National Bureau of Economic Research.

Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy*, 30(1), 51–89.

Hull, P. (2018). Isolateing: Identifying counterfactual-specific treatment effects with cross-stratum comparisons. *Available at SSRN 2705108*.

Hussam, R., Rigol, N., & Roth, B. (2017). Targeting high ability entrepreneurs using community information: Mechanism design in the field.

Imbens, G. W. & Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica (1986-1998)*, 62(2), 467.

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.

Kaboski, J. P. & Townsend, R. M. (2011). A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica*, 79(5), 1357–1406.

Kaboski, J. P. & Townsend, R. M. (2012). The impact of credit on village economies. *American Economic Journal: Applied Economics*, 4(2), 98–133.

Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014a). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics*, 129(2), 597–652.

Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014b). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics*, 129(2), 597–652.

Karlan, D. & Zinman, J. (2009). Observing unobservables: Identifying information asymmetries with a consumer credit field experiment. *Econometrica*, 77(6), 1993–2008.

Karlan, D. & Zinman, J. (2013). Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in mexico.

Karlan, D. S. & Zinman, J. (2008). Credit elasticities in less-developed economies: Implications for microfinance. *American Economic Review*, 98(3), 1040–68.

Kaur, S. (2014). Nominal wage rigidity in village labor markets. *NBER Working Paper 20770*.

Keele, L. J. & Titiunik, R. (2015). Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1), 127–155.

Key, N., Sadoulet, E., & Janvry, A. D. (2000). Transactions costs and agricultural household supply response. *American journal of agricultural economics*, 82(2), 245–259.

Kirkeboen, L. J., Leuven, E., & Mogstad, M. (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics*, 131(3), 1057–1111.

Kline, P. & Walters, C. R. (2017). Through the looking glass: Heckits, late, and numerical equivalence. *arXiv preprint arXiv:1706.05982*.

Kondylis, F., Mueller, V., & Zhu, J. (2017). Seeing is believing? evidence from an extension network experiment. *Journal of Development Economics*, 125, 1–20.

Koundouri, P. (2004a). Current issues in the economics of groundwater resource management. *Journal of Economic Surveys*, 18(5), 703–740.

Koundouri, P. (2004b). Potential for groundwater management: Gisser-sanchez effect reconsidered. *Water Resources Research*, 40(6).

Kowalski, A. E. (2016). Doing more when you're running late: Applying marginal treatment effect methods to examine treatment effect heterogeneity in experiments.

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.

Ligon, E., Thomas, J. P., & Worrall, T. (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *The Review of Economic Studies*, 69(1), 209–244.

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.

Matzkin, R. L. (2007). Nonparametric identification. *Handbook of Econometrics*, 6, 5307–5368.

Meager, R. (2015). Understanding the impact of microcredit expansions: A bayesian hierarchical analysis of 7 randomised experiments.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.

Mogstad, M., Santos, A., & Torgovitsky, A. (2017). *Using instrumental variables for inference about policy relevant treatment effects*. Technical report, National Bureau of Economic Research.

Mountjoy, J. (2018). Community colleges and upward mobility.

Narain, P., Khan, M., & Singh, G. (2006). *Potental for Water Conservation and Havesting Against Drought in Rajasthan*, volume 104. IWMI.

Niehaus, P. & Sukhtankar, S. (2013). The marginal rate of corruption in public programs: Evidence from india. *Journal of public Economics*, 104, 52–64.

Rajasthan Directorate of Economics and Statistics (2011). Basic statistics 2011.

Restuccia, D. & Santaeulalia-Llopis, R. (2017). Land misallocation and productivity. *NBER Working Paper 23128*.

Rodell, M., Famiglietti, J., Wiese, D., Reager, J., Beaudoing, H., Landerer, F., & Lo, M.-H. (2018). Emerging trends in global freshwater availability. *Nature*, (pp.1).

Rodell, M., Velicogna, I., & Famiglietti, J. S. (2009). Satellite-based estimates of groundwater depletion in india. *Nature*, 460(7258), 999.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2), 135–146.

Sayre, S. & Taraz, V. (2018). Groundwater externalities with endogenous well deepening.

Schlenker, W., Hanemann, W. M., & Fisher, A. C. (2007). Water availability, degree days, and the potential impact of climate change on irrigated agriculture in california. *Climatic Change*, 81(1), 19–38.

Sekhri, S. (2013). Sustaining groundwater: Role of policy reforms in promoting conservation in india. *India Policy Forum 2012-13, Volume 9*.

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.

Shah, F. A., Zilberman, D., & Chakravorty, U. (1995). Technology adoption in the presence of an exhaustible resource: the case of groundwater extraction. *American Journal of Agricultural Economics*, 77(2), 291–299.

Shah, T. (2009). Climate change and groundwater: India's opportunities for mitigation and adaptation. *Environmental Research Letters*, 4(3), 035005.

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.

Small, K. A. & Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica: Journal of the Econometric Society*, (pp. 105–130).

Suhag, R. (2016). Overview of ground water in india.

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

Tarozzi, A., Desai, J., & Johnson, K. (2015). The impacts of microcredit: Evidence from ethiopia. *American Economic Journal: Applied Economics*, 7(1), 54–89.

Theis, C. V. (1935). The relation between the lowering of the piezometric surface and the rate and duration of discharge of a well using ground-water storage. *Eos, Transactions American Geophysical Union*, 16(2), 519–524.

Udry, C. (1996a). Efficiency and market structure: testing for profit maximization in african agriculture.

Udry, C. (1996b). Gender, agricultural production, and the theory of the household. *Journal of political Economy*, 104(5), 1010–1046.

Willig, R. D. (1978). Incremental consumer's surplus and hedonic price adjustment. *Journal of Economic Theory*, 17(2), 227–253.

Willis, R. J. & Rosen, S. (1979). Education and self-selection. *Journal of political Economy*, 87(5, Part 2), S7–S36.

Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2), 420–445.

World Bank (2007). World development report 2008: Agriculture for development.

Wright, B. D. & Williams, J. C. (1984). The welfare effects of the introduction of storage. *The Quarterly Journal of Economics*, 99(1), 169–192.

Young, A. (2018). Consistency without inference: Instrumental variables in practical application.

Appendix A

The treatment effect elasticity of demand: Estimating the welfare losses from groundwater depletion in India

A.1 Data appendix

A.1.1 Construction of W_{ns}

I construct two variables using potential crop yield: log relative potential irrigated crop yield, and log potential rainfed crop yield. Define A^I_{nsc} and A^R_{nsc} to be the FAO GAEZ potential crop yield in district n in state s for crop c under the intermediate irrigated and rainfed scenarios, respectively, which I calculate by averaging the values across FAO GAEZ 5 arc-minute cells to the district level. Let L_{nsct} be the land allocated to crop c in district n in state s in year t, observed in Ag '56-'11. Let $L_{sc} = \sum_{n,t} L_{nsct}$ be the total area, across all years in Ag '56-'11, allocated to crop c in state s. I define

$$W_{ns} \equiv \log \frac{\sum_c L_{sc} \min\{A^I_{nsc}, 10A^R_{nsc}\}}{\sum_c L_{sc} A^R_{nsc}}$$

$$\log \text{RF yield}_{ns} \equiv \log \frac{\sum_c L_{sc} A^R_{nsc}}{\sum_c L_{sc}}$$

where W_{ns} is the log relative potential irrigated crop yield, and RF yield$_{ns}$ is the log potential rainfed crop yield. A few notes on the construction. First, the weights L_{sc} are constant within state; this ensures that variation in W_n is caused by variation across districts in the potential yield increase from irrigation, and not variation across districts in weights. Since these weights vary across states, I control flexibly for state in all analysis. It is important to allow the weights to vary across states; there is large variation across states in crop choice. Second, applying $\min\{A^I_{nsc}, 10A^R_{nsc}\}$ is similar to winsorizing W_{ns} at $\log 10$ for each crop. This is almost exclusively necessary for a few desert districts in Rajasthan and Gujarat; dropping these districts does not meaningfully change results, and the weighted instrumental variables estimator already places very little weight on these districts. However, not implementing this winsorization puts very high weight on these districts in estimation of the coefficient on W_{ns}, since these districts' predicted rainfed yield is close to 0. Since these districts are very dependent on irrigation and have relatively high yields, this increases the first stage and reduced form coefficients on W_{ns}. Third, controlling for $\log \text{RF yield}_{ns}$ and a state fixed effect, the coefficient on W_{ns} would be the same if instead $W_{ns} = \log \frac{\sum_c L_{sc} \min\{A^I_{nsc}, 10A^R_{nsc}\}}{\sum_c L_{sc}}$, or log potential irrigated crop yield.

A.2 Model appendix

A.2.1 Proofs and derivations appendix

Proof of generality of functional form.. Under weak separability of unobserved heterogeneity, and imposing the exclusion restrictions, agent surplus under treatment $Y_{1i}(w) - C_{1i}(z) = U(h(w, z), \tilde{V}_i)$, following Bhattacharya (2017) in defining weak separability. Taking derivatives with respect to w and z yields

$$\frac{\partial Y_{1i}}{\partial w} = \frac{\partial U(h(z,w); \tilde{V}_i)}{\partial h}\frac{\partial h(z,w)}{\partial w}$$
$$\frac{\partial C_{1i}}{\partial z} = \frac{\partial U(h(z,w); \tilde{V}_i)}{\partial h}\frac{\partial h(z,w)}{\partial z}$$
$$\frac{\partial^2 Y_{1i}}{\partial w \partial z} = \frac{\partial^2 U(h(z,w); \tilde{V}_i)}{\partial h^2}\frac{\partial h(z,w)}{\partial w}\frac{\partial h(z,w)}{\partial z} + \frac{\partial U(h(z,w); \tilde{V}_i)}{\partial h}\frac{\partial^2 h(z,w)}{\partial w \partial z}$$

A few restrictions appear here. First, $\frac{\partial^2 Y_{1i}}{\partial w \partial z} = 0$ (exclusion restriction). Second, $\frac{\partial h(z,w)}{\partial z} > 0$ and $\frac{\partial h(z,w)}{\partial w} > 0$ (monotonicity). Third, $\frac{\partial U(h(z,w);\tilde{V}_i)}{\partial h} > 0$ (monotonicity). Therefore, excluding edge cases, $\frac{\partial^2 U(h(z,w);\tilde{V}_i)}{\partial h^2} = 0$ and $\frac{\partial^2 h(z,w)}{\partial w \partial z} = 0$. The latter implies $h(z, w) = h_W(w) + h_Z(z) + V_{hi}$. The former implies $\frac{\partial U(h(z,w);\tilde{V}_i)}{\partial h} = V_{\gamma i}$ for some constant which is a function of $\tilde{V}_i$. Making these substitutions implies

$$Y_{1i}(w) - C_{1i}(z) = V_{\gamma i}(h_W(w) + h_Z(z) + V_{hi}) + \tilde{v}_i$$

which is equivalent to

$$\begin{aligned} Y_{1i}(w) &= V_{\gamma i}\gamma_W(w) + V_{1i} \\ C_{1i}(z) &= V_{\gamma i}\gamma_Z(z) + V_{Ci} \end{aligned}$$

■

Derivation of Equation 1.4 and 1.5.. Calculating each derivative,

$$\begin{aligned} \frac{d\mathbf{E}[Y_i(z,w)]}{dz} &= f_V(F_V^{-1}(\mathbf{E}[D_i(z,w)]))\frac{d\gamma_Z(z)}{dz}\mathbf{E}[Y_{1i}(w) - Y_{0i}|U_i = \mathbf{E}[D_i(z,w)]] \\ \frac{d\mathbf{E}[\pi_i(z,w)]}{dz} &= -\mathbf{E}[D_i(z,w)]\mathbf{E}[V_{\gamma i}|U_i < \mathbf{E}[D_i(z,w)]]\frac{d\gamma_Z(z)}{dz} \\ \frac{d\mathbf{E}[\pi_i(z,w)]}{dw} &= -\mathbf{E}[D_i(z,w)]\mathbf{E}[V_{\gamma i}|U_i < \mathbf{E}[D_i(z,w)]]\frac{d\gamma_W(w)}{dw} \\ \frac{d\mathbf{E}[D_i(z,w)]}{dz} &= f_V(F_V^{-1}(\mathbf{E}[D_i(z,w)]))\frac{d\gamma_Z(z)}{dz} \\ \frac{d\mathbf{E}[D_i(z,w)]}{dw} &= f_V(F_V^{-1}(\mathbf{E}[D_i(z,w)]))\frac{d\gamma_W(w)}{dw} \end{aligned}$$

Some algebra then yields the desired result. ■

Derivation of Equation 1.10. Calculating the derivative of $\text{TOT}(u;w)$ yields

$$\frac{d\text{TOT}(u;w)}{dw} = \mathbf{E}[V_{\gamma i}|U_i < u]\frac{d\gamma_W(w)}{dw}$$

Some algebra, and results from the proof of Equation 1.4 and 1.5, yields the desired result. ■

Derivation of Equation 1.11.. Calculating each derivative,

$$\begin{aligned} \frac{d\mathbf{E}[Y_i(z,w)]}{dw} = f_V(F_V^{-1}(\mathbf{E}[D_i(z,w)]))\frac{d\gamma_W(w)}{dw}\mathbf{E}[Y_{1i}(w) - Y_{0i}|U_i = \mathbf{E}[D_i(z,w)]] + \\ \mathbf{E}[D_i(z,w)]\mathbf{E}[V_{\gamma i}|U_i < \mathbf{E}[D_i(z,w)]]\frac{d\gamma_W(w)}{dw} \end{aligned}$$

Some algebra, and results from the proof of Equation 1.4 and 1.5, yields the desired result. ■

Proof of Equation 1.19.. It suffices to show that $\beta_Z^{WIV} + \text{LASE}_W = \beta_W^{IV}$. Let $Z_i^{\perp} \equiv Z_i - \mathbf{E}[Z_i|W_i, X_i]$, and $W_i^{\perp} \equiv W_i - \mathbf{E}[W_i|Z_i, X_i]$. Note that

$$\beta_Z^{WIV} = \sum_s \frac{\mathbf{E}[Y_i Z_i^{\perp}|S_i = s]}{\mathbf{E}[D_i Z_i^{\perp}|S_i = s]} \frac{\mathbf{E}[1\{S_i = s\}(\varpi_W(S_i)/\varpi_Z(S_i))D_i Z_i^{\perp}]}{\mathbf{E}[(\varpi_W(S_i)/\varpi_Z(S_i))D_i Z_i^{\perp}]}$$

I then proceed in two steps. First, I show that

$$\frac{\mathbf{E}[1\{S_i = s\}(\overline{\omega}_W(S_i)/\overline{\omega}_Z(S_i))D_i Z_i^{\perp}]}{\mathbf{E}[(\overline{\omega}_W(S_i)/\overline{\omega}_Z(S_i))D_i Z_i^{\perp}]} = \frac{\mathbf{E}[1\{S_i = s\}D_i W_i^{\perp}]}{\mathbf{E}[D_i W_i^{\perp}]}$$

Second, I consider conditions under which Assumption 5a holds. Written in terms of the natural estimators of $\text{LATE}_{Z|s}$ and $\text{LATE}_{W|s} + \text{LASE}_{W|s}$, with $\text{LASE}_{W|s}$ defined similarly,

$$\frac{\mathbf{E}[Y_i Z_i^{\perp}|S_i = s]}{\mathbf{E}[D_i Z_i^{\perp}|S_i = s]} = \frac{\mathbf{E}[Y_i W_i^{\perp}|S_i = s]}{\mathbf{E}[D_i W_i^{\perp}|S_i = s]} - \text{LASE}_{W|s}$$

Substituting each of these expressions into the original equation yields

$$\beta_Z^{WIV} + \text{LASE}_W = \sum_s \frac{\mathbf{E}[Y_i W_i^{\perp}|S_i = s]}{\mathbf{E}[D_i W_i^{\perp}|S_i = s]} \frac{\mathbf{E}[1\{S_i = s\}D_i W_i^{\perp}]}{\mathbf{E}[D_i W_i^{\perp}]} = \beta_W^{IV}$$

which completes the proof.

For the first step, I use the result that $\overline{\omega}_W(s) = \frac{\mathbf{E}[1\{S_i=s\}D_i W_i^{\perp}]}{\mathbf{E}[D_i W_i^{\perp}]}$ and $\overline{\omega}_Z(s) = \frac{\mathbf{E}[1\{S_i=s\}D_i Z_i^{\perp}]}{\mathbf{E}[D_i Z_i^{\perp}]}$, which can be shown by rewriting the IV estimator as a weighted average of IV estimators conditional on $S_i = s$. Substituting these expressions in immediately completes the first step.

For the second step, I impose some additional assumptions. First, $Z_i^{\perp} \perp (W_i, \tilde{X}_i)$ and $W_i^{\perp} \perp (Z_i, \tilde{X}_i)$ conditional on $S_i = s$. These are strong assumptions, but can be achieved by reweighting. Second, I assume marginal treatment effects and the propensity score are linear conditional on $S_i = s$. Third, I assume $\mathbf{E}[(Z_i^{\perp})^3|S_i = s] = 0$ and $\mathbf{E}[(W_i^{\perp})^3|S_i = s] = 0$. Again, these are strong assumptions, but can be achieved by reweighting.

I then proceed using

$$\frac{\mathbf{E}[Y_i Z_i^{\perp}|S_i = s]}{\mathbf{E}[D_i Z_i^{\perp}|S_i = s]} = \frac{\mathbf{E}[(\mathbf{E}[Y_i|Z_i^{\perp}, W_i, X_i] - \mathbf{E}[Y_i|Z_i = \mathbf{E}[Z_i|W_i, X_i], W_i, X_i])Z_i^{\perp}|S_i = s]}{\mathbf{E}[(\mathbf{E}[D_i|Z_i^{\perp}, W_i, X_i] - \mathbf{E}[D_i|Z_i = \mathbf{E}[Z_i|W_i, X_i], W_i, X_i])Z_i^{\perp}|S_i = s]}$$

This requires a few steps. I focus on the numerator; the approach is the same for the denominator. First, I project Y_i onto $Z_i^{\perp}$, yielding $\mathbf{E}[Y_i Z_i^{\perp}|S_i = s] = \mathbf{E}[\mathbf{E}[Y_i|Z_i^{\perp}, S_i]Z_i^{\perp}|S_i = s]$. Second, I apply the law of iterated expectations. Since $Z_i^{\perp} \perp (W_i, \tilde{X}_i)$ conditional on $S_i = s$, $\mathbf{E}[\mathbf{E}[Y_i|Z_i^{\perp}, S_i]Z_i^{\perp}|S_i = s] = \mathbf{E}[\mathbf{E}[Y_i|Z_i^{\perp}, W_i, X_i]Z_i^{\perp}|S_i = s]$. Lastly, using $Z_i^{\perp} \perp (W_i, \tilde{X}_i)$, and $\mathbf{E}[Z_i^{\perp}|S_i = s] = 0$, we complete the equality.

Next, I substitute these differences with integrals over marginal treatment effects and the propensity score. Here, I use the linearization of both. Let $\text{MTE}(u; w, \tilde{x}, s) = m_{1s}u + m_{2s}w + \tilde{x}'m_{3s}$ and $\mathbf{E}[D_i(z, w; \tilde{x}, s)] = d_{1s}z + d_{2s}w + \tilde{x}'d_{3s}$. Then, some calculus yields

$$\frac{\mathbf{E}[Y_i Z_i^{\perp}|S_i = s]}{\mathbf{E}[D_i Z_i^{\perp}|S_i = s]} = \frac{\mathbf{E}[d_{1s}(Z_i^{\perp})^2(m_{1s}\mathbf{E}[D_i|W_i, X_i] + m_{2s}W_i + \tilde{X}_i' m_{3s}) + \frac{1}{2}d_{1s}m_{1s}(Z_i^{\perp})^3|S_i = s]}{\mathbf{E}[d_{1s}(Z_i^{\perp})^2|S_i = s]}$$

Two simplifications can be made here. First, I use $\mathbf{E}[(Z_i^{\perp})^3|S_i = s] = 0$. Second, I use $Z_i^{\perp} \perp (W_i, \tilde{X}_i)$. Together, these yield

$$\frac{\mathbf{E}[Y_i Z_i^{\perp}|S_i = s]}{\mathbf{E}[D_i Z_i^{\perp}|S_i = s]} = m_{1s}\mathbf{E}[D_i|S_i = s] + m_{2s}\mathbf{E}[W_i|S_i = s] + \mathbf{E}[\tilde{X}_i'|S_i = s]m_{3s}$$

A symmetric proof shows the same result holds for $\frac{\mathbf{E}[Y_i W_i^{\perp}|S_i=s]}{\mathbf{E}[D_i W_i^{\perp}|S_i=s]} - \text{LASE}_{W|s}$, which completes the proof.

■

A.2.2 Weights

LATE and LASE weights

I start with the result from Heckman & Vytlacil (2005) on OLS.

$$\frac{\text{Cov}(Q, T - \mathbf{E}[T|X])}{\text{Var}(T - \mathbf{E}[T|X])} = \int\int \frac{\partial \mathbf{E}[Q|T = t, X = x]}{\partial t}\omega(t,x)dtdx$$

$$\omega(t,x) = \frac{\Pr[T > t, X = x]\mathbf{E}[T - \mathbf{E}[T|X]|T > t, X = x]}{\int\int \Pr[T > t', X = x']\mathbf{E}[T - \mathbf{E}[T|X]|T > t', X = x]dt'dx'}$$

The first expression shows that the coefficient on T, controlling for X, estimates a weighted average of derivatives of the conditional expectation function of Q given $T = t$ and $X = x$ with respect to t. The second expression shows that the weights $\omega(t,x)$ are the partial expectation, conditional on $X = x$, of $T - \mathbf{E}[T|X]$ given $T > t$, times the probability that $X = x$. Note this partial expectation approaches 0 at the edges of the conditional support of T conditional on $X = x$, which is consistent with our intuition that OLS estimates should not depend on derivatives of the conditional expectation function outside the support of the covariates. Additionally, it is helpful to note that

$$\int \omega(t,x)dt = \frac{\Pr[X = x]\text{Var}(T|X = x)}{\int \Pr[X = x']\text{Var}(T|X = x')dx'}$$

The weights placed on each x depend on the probability $X = x$ and the conditional variance of T given $X = x$.

Still following Heckman & Vytlacil (2005), we can now apply this to the IV estimator $\beta_Z^{IV} = \frac{\text{Cov}(Y_i, Z_i - \mathbf{E}[Z_i|W_i, X_i])}{\text{Cov}(D_i, Z_i - \mathbf{E}[Z_i|W_i, X_i])} = \text{LATE}_Z$. For the definition of these weights, it will be useful to define the propensity score $P(z, w; x) = \mathbf{E}[D_i|Z_i = z, W_i = w, X_i = x]$. Note that just identified linear instrumental variables is just a ratio of OLS estimators, so we can simply apply the formula above. Additionally, we make the substitution that $\frac{\partial \mathbf{E}[Y_i(z,w;x)]}{\partial z} = \frac{\partial P(z,w;x)}{\partial z} \text{MTE}(P(z, w; x); w, x)$. Applying these results yields

$$\text{LATE}_Z = \int \text{MTE}(u; w, x)\omega_Z(u; w, x) du dw dx$$

$$\begin{aligned}\omega_Z(u; w, x) = &(\Pr[P(Z_i, W_i; X_i) > u, W_i = w, X_i = x] \cdot \\ &\mathbf{E}[Z_i - \mathbf{E}[Z_i|W_i, X_i]|P(Z_i, W_i; X_i) > u, W_i = w, X_i = x]) \,/ \\ &\left(\int\int\int \Pr[P(Z_i, W_i; X_i) > u', W_i = w', X_i = x'] \cdot \right. \\ &\left. \mathbf{E}[Z_i - \mathbf{E}[Z_i|W_i, X_i]|P(Z_i, W_i; X_i) > u', W_i = w', X_i = x'] du' dw' dx' \right)\end{aligned}$$

Once again, the weights on MTE are in terms of partial expectation functions; weight is placed on latent propensities to adopt u within the support of the propensity score $P(Z_i, W_i; X_i)$. Again, for interpretation it is helpful to integrate over u to estimate the weight placed on observations with $(W_i, X_i) = (w, x)$. When the propensity score is linear in z conditional on (W_i, X_i), one can show

$$\int \omega_Z(u; w, x) du = \frac{\text{Var}(P(Z_i, W_i; X_i)|W_i = w, X_i = x)\Pr[W_i = w, X_i = x]}{\int\int \text{Var}(P(Z_i, W_i; X_i)|W_i = w', X_i = x')\Pr[W_i = w', X_i = x'] dw' dx'}$$

The most weight is placed on values of (W_i, X_i) which have the highest conditional variance of the propensity score and which are observed the most frequently.

Finally, we can apply this to instrumental variables using W_i as an instrument, $\beta_W^{IV} = \frac{\text{Cov}(Y_i, W_i - \mathbf{E}[W_i|Z_i, X_i])}{\text{Cov}(D_i, W_i - \mathbf{E}[W_i|Z_i, X_i])} = \text{LASE}_W + \text{LATE}_W$. Once again, we represent this as the ratio of OLS estimators, and we apply the result above for OLS. Here, we make use of the fact that $\frac{\partial \mathbf{E}[Y_i(z,w;x)]}{\partial w} = \frac{\partial P(z,w;x)}{\partial w}(\text{MSE}(P(z, w; x); x) + \text{MTE}(P(z, w; x); w, x))$. It will also be necessary to define implicitly define $\check{Z}(u; w, x)$ by $u = P(\check{Z}(u; w, x), w; x)$; $\check{Z}$ inverts the propensity score to recover the value of z that will set the propensity score equal to u

given $(W_i, X_i) = (w, x)$. Then,

$$\begin{aligned}
\text{LATE}_W &= \int \text{MTE}(u; w, x)\omega_W(u; w, x)dudwdx \\
\text{LASE}_W &= \int \text{MSE}(u; x)\omega_W(u; w, x)dudwdx \\
\omega_W(u; w, x) &= \Bigg(\frac{\partial P(\check{Z}(u; w, x), w; x)/\partial w}{\partial P(\check{Z}(u; w, x), w; x)/\partial z} \cdot \\
&\quad \Pr[W_i > w, P(Z_i, W_i; X_i) = u, X_i = x] \cdot \\
&\quad \mathbf{E}[W_i - \mathbf{E}[W_i|Z_i, X_i]|W_i > w, P(Z_i, W_i; X_i) = u, X_i = x]\Bigg) / \\
&\quad \Bigg(\int\int\int \frac{\partial P(\check{Z}(u'; w', x'), w'; x')/\partial w}{\partial P(\check{Z}(u'; w', x'), w'; x')/\partial z} \cdot \\
&\quad \Pr[W_i > w', P(Z_i, W_i; X_i) = u', X_i = x'] \cdot \\
&\quad \mathbf{E}[W_i - \mathbf{E}[W_i|Z_i, X_i]|W_i > w', P(Z_i, W_i; X_i) = u', X_i = x']du'dw'dx'\Bigg)
\end{aligned}$$

Although these expressions appear more complicated, integrating over u and w, once again we can interpret them roughly as variances of the propensity score conditional on the controls Z_i and X_i; this is exact when the propensity score is linear in z and w conditional on $X_i = x$.

Finally, these expressions are all functions of $P(z, w; x)$ and the joint distribution of (Z_i, W_i, X_i), all of which are nonparametrically identified, so the weights are nonparametrically identified. In practice, estimation of the weights may involve placing parametric restrictions on $P(z, w; x)$.

Efficient reweighting

Define

$$\beta_Z^{WIV}(w_Z) = \frac{\text{Cov}(w_Z(S_i)Y_i, Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)])}{\text{Cov}(w_Z(S_i)D_i, Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)])}$$

and $\beta_W^{WIV}(w_W)$ analagously. Let $\overline{\omega}_W(s) = \int \omega_W(u; w, (x, s))dudwdx$ and $\overline{\omega}_Z(s) = \int \omega_Z(u; w, (x, s))dudwdx$. Given this, for $\beta_W^{WIV}(w_W)$ and $\beta_Z^{WIV}(w_Z)$ to place the same weight on compliers with $S_i = s$, it must be the case that

$$w_Z(s)\overline{\omega}_Z(s) = w_W(s)\overline{\omega}_W(s)$$

Efficient weights solve

$$w = \arg\min_{m} \mathrm{Var}\left[\hat{\beta}_W^{WIV}(w_W) - \hat{\beta}_Z^{WIV}(w_Z)\right]$$
$$\text{s.t. } w_Z(s)\overline{\omega}_Z(s) = w_W(s)\overline{\omega}_W(s)$$

I assume the propensity score is linear in (z, w). Under this assumption, $\overline{\omega}_W$ and $\overline{\omega}_Z$ simplify to

$$\overline{\omega}_W(s) = \frac{\mathrm{Var}(W_i - \mathbf{E}[W_i|Z_i, (X_i, S_i)]|S_i = s)\Pr[S_i = s]}{\mathrm{Var}(W_i - \mathbf{E}[W_i|Z_i, (X_i, S_i)])}$$
$$\overline{\omega}_Z(s) = \frac{\mathrm{Var}(Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)]|S_i = s)\Pr[S_i = s]}{\mathrm{Var}(Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)])}$$

Define $g_Z \equiv \frac{\mathrm{Cov}(D_i, Z_i - \mathbf{E}[Z_i|W_i, X_i])}{\mathrm{Var}(Z_i - \mathbf{E}[Z_i|W_i, X_i])}$ and $g_W \equiv \frac{\mathrm{Cov}(D_i, W_i - \mathbf{E}[W_i|Z_i, X_i])}{\mathrm{Var}(W_i - \mathbf{E}[W_i|Z_i, X_i])}$; that g_Z and g_W are constants follows from the assumption that the propensity score is linear in (z, w). Suppose further that the structural errors in the outcome equation are homoskedastic. Then the optimal weights satisfy

$$w_Z(s) = \frac{g_W^2 \mathrm{Var}(W_i - \mathbf{E}[W_i|Z_i, (X_i, S_i)])\overline{\omega}_W(s)}{g_Z^2 \mathrm{Var}(Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)])\overline{\omega}_Z(s) + g_W^2 \mathrm{Var}(W_i - \mathbf{E}[W_i|Z_i, (X_i, S_i)])\overline{\omega}_W(s)}$$
$$w_W(s) = \frac{g_Z^2 \mathrm{Var}(Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)])\overline{\omega}_Z(s)}{g_Z^2 \mathrm{Var}(Z_i - \mathbf{E}[Z_i|W_i, (X_i, S_i)])\overline{\omega}_Z(s) + g_W^2 \mathrm{Var}(W_i - \mathbf{E}[W_i|Z_i, (X_i, S_i)])\overline{\omega}_W(s)}$$

To interpret this expression, note that the realized equivalent of $\frac{g_Z^2 \mathrm{Var}(Z_i|W_i, (X_i, S_i))}{g_W^2 \mathrm{Var}(W_i|Z_i, (X_i, S_i))}$ is just the ratio of the first stage F-stats. As one F-stat grows arbitrarily large relative to the other, the weights essentially reweight observations in the regression with the larger F-stat so that the weights on observables in that regression are the same as the weights on observables in the unweighted regression with the smaller F-stat.

A.2.3 Control function

The control function approach is predicated on the normality assumption

$$\begin{pmatrix} Y_{1i} \\ C_{1i} \\ Y_{0i} \end{pmatrix} \sim N\left(\begin{pmatrix} (g_W + c_0)W_i + X_i'\mu_1 \\ g_Z Z_i + X_i'\mu_C \\ c_0 W_i + X_i'\mu_0 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{1c} & \Sigma_{10} \\ \Sigma_{1c} & \Sigma_{cc} & \Sigma_{c0} \\ \Sigma_{10} & \Sigma_{c0} & \Sigma_{cc} \end{pmatrix}\right)$$

Under this model,

$$\mathbf{E}[D_i(z,w;x)] = \Phi\left(\frac{-x'\mu_V + g_W w - g_Z z}{\sigma_V}\right)$$

where Φ is the normal CDF, $\mu_V = -\mu_1 + \mu_C + \mu_0$,[1] and $\sigma_V^2 = \text{Var}[V_i|X_i]$. I estimate this with a first step probit; conventionally, σ_V would not be identified. However, as noted by Björklund & Moffitt (1987), the generalized Roy structure allows it to be identified here, since we can estimate the direct effect of w on treatment effects. I do this in the second step, using the identity

$$\mathbf{E}\left[Y_{di}|D_i = d, Z_i = z, W_i = w, X_i = x\right] = X_i'\mu_d + c_d w + b_d \lambda_d(\mathbf{E}[D_i(z,w;x)])$$

where $c_0 = 0$, $c_1 - c_0 = g_W$, $b_0 = \frac{\text{Cov}(V_{0i},V_i|X_i)}{\sigma_V}$, $b_1 = -\frac{\text{Cov}(V_{1i},V_i|X_i)}{\sigma_V}$, $\lambda_0(u) = \frac{\phi(\Phi^{-1}(u))}{1-u}$, and $\lambda_1(u) = \frac{\phi(\Phi^{-1}(u))}{u}$. I estimate this conditional expectation function by OLS. Note the exclusion restriction that Z_i does not directly enter the conditional expectation function for Y_{di}. Although this is not required to estimate the model under normality, without this exclusion restriction identification depends strongly on functional form assumptions.

In Table 1.4 and Table 1.6, I construct control function estimates of local average treatment effects and local average surplus effects. Let $Z_i^{\perp} = Z_i - \mathbf{E}[Z_i|W_i, X_i]$ For a local average treatment effect, I use

$$\begin{aligned}
\frac{\mathbf{E}[Y_i Z_i^{\perp}]}{\mathbf{E}[D_i Z_i^{\perp}]} &= \frac{\mathbf{E}[(Y_i - \mathbf{E}[Y_i|Z_i = \mathbf{E}[Z_i|W_i,X_i],W_i,X_i])Z_i^{\perp}]}{\mathbf{E}[(D_i - \mathbf{E}[D_i|Z_i = \mathbf{E}[Z_i|W_i,X_i],W_i,X_i])Z_i^{\perp}]} \\
&= \frac{\mathbf{E}[(\mathbf{E}[Y_i|Z_i,W_i,X_i] - \mathbf{E}[Y_i|Z_i = \mathbf{E}[Z_i|W_i,X_i],W_i,X_i])Z_i^{\perp}]}{\mathbf{E}[(D_i - \mathbf{E}[D_i|Z_i = \mathbf{E}[Z_i|W_i,X_i],W_i,X_i])Z_i^{\perp}]} \\
&= \frac{\mathbf{E}\left[\int_{\mathbf{E}[D_i|Z_i=\mathbf{E}[Z_i|W_i,X_i],W_i,X_i]}^{\mathbf{E}[D_i|Z_i,W_i,X_i]} \text{MTE}(u;W_i,X_i)du Z_i^{\perp}\right]}{\mathbf{E}[(D_i - \mathbf{E}[D_i|Z_i = \mathbf{E}[Z_i|W_i,X_i],W_i,X_i])Z_i^{\perp}]}
\end{aligned}$$

Focusing on the numerator in each expression. The first step follows from $\mathbf{E}[\mathbf{E}[Y_i|Z_i = \mathbf{E}[Z_i|W_i,X_i],W_i,X_i]Z_i^{\perp}] = 0$, which follows from an application of the law of iterated expectations conditioning on (W_i, X_i). The second step follows from $\mathbf{E}[Y_i Z_i^{\perp}] = \mathbf{E}[\mathbf{E}[Y_i|Z_i,W_i,X_i]Z_i^{\perp}]$. This again follows from an application of the law of iterated expectations conditioning on (Z_i, W_i, X_i). The third step is just the fundamental theorem of calculus, and that the marginal treatment effect equals the derivative of the conditional expectation of Y_i with respect to z. I therefore use the plug-in estimator of this as my control function estimate of the local average treatment effect. Nearly identical calculations hold for the local average surplus effect, and bias from exclusion restriction violations. Standard errors are calculated using the delta method, and derivatives with

[1] This implies $\mathbf{E}[V_i|X_i] = X_i'\mu_V$.

respect to control function parameters are estimated numerically.

A.3 Figures

Figure A.1: Productivity growth and groundwater withdrawals

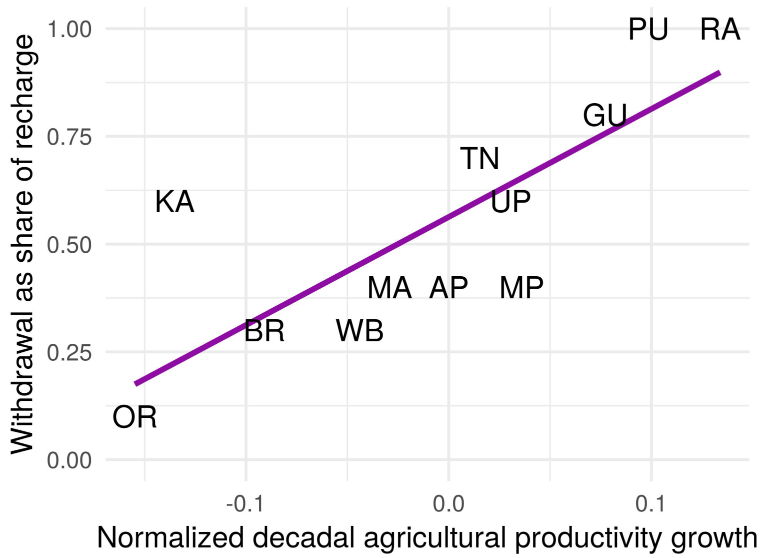

Notes: This figure plots, for each state, the lower bound estimate of its groundwater withdrawals as a share of recharge rate, as reported in Rodell et al. (2009), against its normalized decadal agricultural productivity growth, calculated in a regression of log agricultural productivity on state fixed effects interacted with year dummies, relative to Andhra Pradesh (AP). The purple line is the line of best fit, with a slope of 2.5 and $R^2 = 0.63$.

Appendix B

Irrigation in Rwanda: Farmers' Responses to a Massive Expansion of the Production Possibility Frontier

B.1 Sampling

Work in progress: Sampling appendix.

B.2 Model appendix

Derivation of first order conditions. Substitute for L^O using the household labor constraint, $L_1 + L_2 + \ell + L^O = \overline{L}$, and substitute for c in the household's maximization problem. This leaves two constraints, $M_1 + M_2 \leq \overline{M}$, and $\overline{L} - L_1 - L_2 - \ell \leq \overline{L^O}$; call the multipliers on these constraints $\widetilde{\lambda_M}$ and $\widetilde{\lambda_L}$, respectively. Taking first order conditions yields

$$\begin{array}{rcl} (M_k) & \mathbf{E}[u_c\sigma]A_k F_{kM} - \mathbf{E}[u_c]r & = \widetilde{\lambda_M} \\ (L_k) & \mathbf{E}[u_c\sigma]A_k F_{kL} - \mathbf{E}[u_c]w & = -\widetilde{\lambda_L} \\ (\ell) & \mathbf{E}[u_\ell] - \mathbf{E}[u_c]w & = -\widetilde{\lambda_L} \end{array}$$

To ease interpretation, normalize $\lambda_M \equiv \widetilde{\lambda_M}/r\mathbf{E}[u_c]$ and $\lambda_L \equiv \widetilde{\lambda_L}/w\mathbf{E}[u_c]$, and substitute $\text{cov}(\sigma, u_c) = \mathbf{E}[u_c\sigma] - \mathbf{E}[u_c]\mathbf{E}[\sigma] = \mathbf{E}[u_c\sigma] - \mathbf{E}[u_c]$. This yields

$$\begin{aligned}
(M_k) \quad & \left(1 + \tfrac{\text{cov}(\sigma,u_c)}{\mathbf{E}[u_c]}\right) A_k F_{kM} && = (1 + \lambda_M) r \\
(L_k) \quad & \left(1 + \tfrac{\text{cov}(\sigma,u_c)}{\mathbf{E}[u_c]}\right) A_k F_{kL} && = (1 - \lambda_L) w \\
(\ell) \quad & \tfrac{\mathbf{E}[u_\ell]}{\mathbf{E}[u_c]} && = (1 - \lambda_L) w
\end{aligned}$$

No constraints. When no constraints bind, as discussed the first order conditions simplify to

$$\begin{aligned}
(M_k) \quad & A_k F_{kM} && = r \\
(L_k) \quad & A_k F_{kL} && = w \\
(\ell) \quad & \tfrac{u_\ell}{u_c} && = w
\end{aligned}$$

Note that the first order conditions for M_2 and L_2 are functions only of (M_2, L_2), and exogenous (A_2, r, w). Therefore, $\frac{dM_2}{dA_1} = \frac{dL_2}{dA_1} = 0$.

Insurance market failure. Consider the case when insurance markets fail. To abstract fully from labor supply, we temporarily remove leisure from the model. To further simplify, we drop other inputs from the production function; when the production function is homogeneous in labor and other inputs, this is without loss of generality. Let $\gamma(c) = \frac{\mathbf{E}[u_c(c)]}{\mathbf{E}[\sigma u_c(c)]}$; $\gamma > 1$ is the ratio of the marginal utility from consumption to the marginal utility from agricultural production. This yields the first order conditions

$$\begin{aligned}
(L_1) \quad & A_1 F_{1L}(L_1) - \gamma(\sigma A_1 F_1(L_1) + \sigma A_2 F_2(L_2) + w(\overline{L} - L_1 - L_2) + r\overline{M}) w && = 0 \\
(L_2) \quad & A_2 F_{2L}(L_2) - \gamma(\sigma A_1 F_1(L_1) + \sigma A_2 F_2(L_2) + w(\overline{L} - L_1 - L_2) + r\overline{M}) w && = 0
\end{aligned}$$

The central intuition for this case can be captured from just the first order conditions: $\overline{L}$ and $\overline{M}$ enter symmetrically into the model, so larger households should respond similarly to richer households. If absolute risk aversion decreases sufficiently quickly (e.g. with CRRA preferences), then for sufficiently high levels of consumption $\mathbf{E}[\sigma u_c] = \mathbf{E}[\sigma]\mathbf{E}[u_c] = \mathbf{E}[u_c] \Rightarrow \gamma = 1$. Therefore, sufficiently wealthy or sufficiently large households should not respond to the sample plot shock. Below, we will maintain the assumption that preferences exhibit decreasing absolute risk aversion.

Some additional comments for the the derivations. The substitution that $A_k F_{kL} = \gamma w$ will be used frequently to simplify. Additionally, we use γ_c to represent the derivative of γ with respect to consumption (so $\gamma_c = \frac{partial}{partialc}\gamma(c)$), and γ_F to represent the derivative of γ with respect to agricultural production (so $\frac{\partial}{\partial F}\gamma(\sigma F)$). We define γ_{Fc} and γ_{FF} analagously. Additionally, the substitution $\frac{1}{A_k F_{kLL}} = \frac{\partial L_k}{\partial w^*}$, the partial derivative of labor demand on plot k with respect to the uncertainty adjusted wage γw, is useful for interpretation. It is assumed that $\frac{\partial L_k}{\partial w^*} < 0$, so labor demand is downward sloping. Finally, two multipliers frequently emerge in the math. The first, $M_N = 1 - \gamma w \frac{\partial L_1}{\partial w^*} \frac{\gamma w}{A_1 F_1}$, is a production multiplier: it captures when A_1 increases, how much does more does F_1 increase due to increased labor allocations to plot 1, holding fixed risk adjusted wages. The second, $M_D = 1 - \gamma_F w \gamma w \frac{\partial (L_1 + L_2)}{\partial w^*}$, is a risk multiplier: its inverse captures when agricultural production is increased, how much less does agricultural production increase due to reduced labor allocations caused by risk aversion. Note that $M_N > 1$, and we will present conditions below sufficient for $\gamma_F > 0$, which implies $M_D > 1$.

Stack the left hand sides of the first order conditions into the vector FOC_I. Define the Jacobian $J_I \equiv D_{(L_1,L_2)} \text{FOC}_I$. Applying the implicit function theorem yields $D_{(A_1)}(L_1, L_2)' = -J_I^{-1} D_{(A_1)} \text{FOC}_I$ and $D_{(\overline{L})}(L_1, L_2)' = -J_I^{-1} D_{(\overline{L})} \text{FOC}_I$. Some algebra yields

$$
\begin{aligned}
J_I &= \begin{pmatrix} A_1 F_{1LL} - \gamma_F w \gamma w & -\gamma_F w \gamma w \\ -\gamma_F w \gamma w & A_2 F_{2LL} - \gamma_F w \gamma w \end{pmatrix} \\
J_I^{-1} &= \frac{\frac{\partial L_1}{\partial w^*} \frac{\partial L_2}{\partial w^*}}{M_D} \begin{pmatrix} A_2 F_{2LL} - \gamma_F w \gamma w & \gamma_F w \gamma w \\ \gamma_F w \gamma w & A_1 F_{1LL} - \gamma_F w \gamma w \end{pmatrix} \\
D_{(A_1)} \text{FOC}_I &= \left(\frac{\gamma w}{A_1} - \gamma_F w \frac{A_1 F_1}{A_1}, -\gamma_F w \frac{A_1 F_1}{A_1} \right)' \\
D_{(\overline{L})} \text{FOC}_I &= \left(-\gamma_c w^2, -\gamma_c w^2 \right)' \\
D_{(A_1)}(L_1, L_2)' &= -\frac{\frac{\partial L_1}{\partial w^*} \frac{\partial L_2}{\partial w^*}}{M_D} \begin{pmatrix} \frac{\gamma w}{A_1 \frac{\partial L_2}{\partial w^*}} - \left(\frac{\gamma_F w (\gamma w)^2}{A_1} - \frac{\gamma_F w A_1 F_1}{A_1 \frac{\partial L_2}{\partial w^*}} \right) \\ \frac{\gamma_F w (\gamma w)^2}{A_1} - \frac{\gamma_F w A_1 F_1}{A_1 \frac{\partial L_1}{\partial w^*}} \end{pmatrix} \\
D_{(\overline{L})}(L_1, L_2)' &= \frac{\gamma_c w^2}{M_D} \begin{pmatrix} \frac{\partial L_1}{\partial w^*} \\ \frac{\partial L_2}{\partial w^*} \end{pmatrix}
\end{aligned}
$$

Some additional simplification will be useful for $\frac{dL_2}{dA_1}$; note that

$$
\frac{dL_2}{dA_1} = \left(\frac{\gamma_F w A_1 F_1}{A_1} \right) \frac{\partial L_2}{\partial w^*} \frac{M_N}{M_D}
$$

Under conditions assumed, a sufficient condition for $\frac{dL_2}{dA_1} < 0$ is that $\gamma_F > 0$: that is, an exogenous increase in agricultural production would increase the household's marginal utility of consumption relative to marginal utility from agricultural production. That this should hold seems intuitive, but it need not hold in general: increases in agricultural production, by increasing consumption, may move households to a less risk averse portion of their utility function, and in turn increase marginal utility from agricultural production relative to marginal utility of consumption. Some assumptions on the distribution of σ rule this out: we follow Karlan et al. (2014b) and, for some $k > 1$, assume $\sigma = k$ with probability $\frac{1}{k}$ ("the good state") and $\sigma = 0$ with probability $\frac{k-1}{k}$ ("the bad state"); i.e., there is a crop failure with probability $\frac{k-1}{k}$.

Under this assumption, some additional simplification and signing is possible. Define $\overline{R} = -\frac{\mathbf{E}[u_c \frac{u_{cc}}{u_c}]}{\mathbf{E}[u_c]}$ to be the household's average risk aversion, $R_k = -\mathbf{E}[\frac{u_{cc}}{u_c}|\sigma = k]$ to be the household's risk aversion in the good state, and $R_{kc} = -\mathbf{E}[\frac{u_{ccc}}{u_c}|\sigma = k] + R_k^2$ to be the derivative of the household's risk aversion in the good state. Note that by assumption, $R_{kc} < 0$ and $R_k < \overline{R}$. From this, it follows that

$$\gamma_c = \frac{\mathbf{E}[u_{cc}]}{\mathbf{E}[\sigma u_c]} - \frac{\mathbf{E}[\sigma u_{cc}]\mathbf{E}[u_c]}{\mathbf{E}[\sigma u_c]^2} = \gamma(R_k - \overline{R}) < 0$$

$$\gamma_F = \frac{\mathbf{E}[\sigma u_{cc}]}{\mathbf{E}[\sigma u_c]} - \frac{\mathbf{E}[\sigma^2 u_{cc}]\mathbf{E}[u_c]}{\mathbf{E}[\sigma u_c]^2} = (k-1)\frac{\mathbf{E}[u_c|\sigma = 0]}{\mathbf{E}[u_c|\sigma = k]} R_k = (k\gamma - 1)R_k > 0$$

$$\gamma_{Fc} = \left((k\gamma - 1)\frac{R_{kc}}{R_k} + k\gamma(R_k - \overline{R})\right) < 0$$

$$\gamma_{FF} = (k^2\gamma - 1)R_{kc} + k(k\gamma - 1)R_k^2 = k\gamma_F R_k \left(-\frac{R_{kc}}{R_k} + R_k\right) > 0$$

Since, $\gamma_F > 0$, it follows that $\frac{dL_2}{dA_1} < 0$, so households substitute labor away from their other plots in response to the sample plot shock.

These results are almost sufficient to sign our cross partial of interest. Note that

$$\frac{\frac{1}{w}\frac{d^2 L_2}{d\overline{L}dA_1}}{\frac{dL_2}{dA_1}} = \frac{\frac{1}{w}\frac{d\gamma_F}{d\overline{L}}}{\gamma_F} + \frac{\frac{1}{w}\frac{dF_1}{d\overline{L}}}{F_1} + \frac{\frac{1}{w}\frac{d(\partial L_2/\partial w^*)}{d\overline{L}}}{(\partial L_2/\partial w^*)} + \frac{\frac{1}{w}\frac{dM_N}{d\overline{L}}}{M_N} - \frac{\frac{1}{w}\frac{dM_D}{d\overline{L}}}{M_D}$$

The following substitutions are now useful

$$
\begin{aligned}
\frac{\frac{1}{w}\frac{d\gamma_F}{d\overline{L}}}{\gamma_F} &= \frac{\gamma_{Fc}}{\gamma_F} + \left(\frac{1}{w}\frac{dF}{d\overline{L}}\right)\frac{\gamma_{FF}}{\gamma_F} \\
\frac{1}{w}\frac{dF}{d\overline{L}} &\equiv \frac{1}{w}\frac{d(A_1F_1 + A_2F_2)}{d\overline{L}} = \frac{\gamma w}{w}\left(\frac{dL_1}{d\overline{L}} + \frac{dL_2}{d\overline{L}}\right) = \frac{M_D - 1}{M_D}\left(-\frac{\gamma_c}{\gamma_F}\right) \\
\frac{\gamma_{Fc}}{\gamma_F} &= \frac{R_{kc}}{R_k} + \frac{k\gamma}{k\gamma - 1}\frac{\gamma_c}{\gamma} \\
\frac{\frac{1}{w}\frac{dF_1}{d\overline{L}}}{F_1} &= \frac{M_N - 1}{M_D}\left(-\frac{\gamma_c}{\gamma}\right) \\
\frac{\frac{1}{w}\frac{d(\partial L_k/\partial w^*)}{d\overline{L}}}{(\partial L_k/\partial w^*)} &= \frac{\gamma_c w \frac{\partial L_k}{\partial w^*}}{M_D}\left(-\frac{F_{kLLL}}{F_{kLL}}\right) \\
\frac{\frac{1}{w}\frac{dM_N}{d\overline{L}}}{M_N} &= \frac{M_N - 1}{M_N}\left(2\frac{\gamma_c}{\gamma} + \frac{\frac{1}{w}\frac{d(\partial L_1/\partial w^*)}{d\overline{L}}}{(\partial L_1/\partial w^*)}\right) \\
\frac{\frac{1}{w}\frac{dM_D}{d\overline{L}}}{M_D} &= \frac{M_D - 1}{M_D}\left(\frac{\gamma_{Fc}}{\gamma_F} + \frac{\gamma_c}{\gamma} + \frac{\frac{1}{w}\frac{d(\partial(L_1+L_2)/\partial w^*)}{d\overline{L}}}{(\partial(L_1 + L_2)/\partial w^*)}\right)
\end{aligned}
$$

Making these substitutions, and rearranging terms, yields

$$
\begin{aligned}
\frac{\frac{1}{w}\frac{d^2L_2}{d\overline{L}dA_1}}{\frac{dL_2}{dA_1}} = {} & \frac{1}{M_D}\frac{R_{kc}}{R_k} + \frac{M_D - 1}{M_D}\left(\left(-\frac{\gamma_c}{\gamma_F}\right)\frac{\gamma_{FF}}{\gamma_F} + \frac{k\gamma}{k\gamma - 1}\frac{\gamma_c}{\gamma}\right) + \\
& \left(\frac{k\gamma}{k\gamma - 1}\frac{1}{M_D} + 2\frac{M_N - 1}{M_N} - \left(1 + \frac{k\gamma}{k\gamma - 1}\right)\frac{M_D - 1}{M_D} - \frac{M_N - 1}{M_N M_D}\right)\frac{\gamma_c}{\gamma} + \\
& \frac{M_N - 1}{M_N}\frac{\frac{1}{w}\frac{d(\partial L_1/\partial w^*)}{d\overline{L}}}{(\partial L_1/\partial w^*)} + \frac{\frac{1}{w}\frac{d(\partial L_2/\partial w^*)}{d\overline{L}}}{(\partial L_2/\partial w^*)} - \frac{M_D - 1}{M_D}\frac{\frac{1}{w}\frac{d(\partial(L_1+L_2)/\partial w^*)}{d\overline{L}}}{(\partial(L_1 + L_2)/\partial w^*)}
\end{aligned}
$$

To sign this derivative, I now make the assumption that $F_{kLLL} \approx 0$. Essentially, this is making as assumption about the relative magnitudes of channels through which the effect of increasing $\overline{L}$, through reduced risk aversion, affects the responsiveness of the household to the sample plot shock. In particular, it assumes that changes in the labor demand elasticity caused by increased labor allocations do not dominate the direct effects of reduced risk aversion. This simplifies this expression to

$$
\begin{aligned}
\frac{\frac{1}{w}\frac{d^2L_2}{d\overline{L}dA_1}}{\frac{dL_2}{dA_1}} = {} & \frac{1}{M_D}\frac{R_{kc}}{R_k} + \frac{M_D - 1}{M_D}\left(\left(-\frac{\gamma_c}{\gamma_F}\right)\frac{\gamma_{FF}}{\gamma_F} + \frac{k\gamma}{k\gamma - 1}\frac{\gamma_c}{\gamma}\right) + \\
& \left(\frac{k\gamma}{k\gamma - 1}\frac{1}{M_D} + 2\frac{M_N - 1}{M_N} - \left(1 + \frac{k\gamma}{k\gamma - 1}\right)\frac{M_D - 1}{M_D} - \frac{M_N - 1}{M_N M_D}\right)\frac{\gamma_c}{\gamma}
\end{aligned}
$$

We now sign each term individually. First, $\frac{1}{M_D}\frac{R_{kc}}{R_k} < 0$, as we assumed decreasing absolute risk aversion. For the second term, note that $\left(-\frac{\gamma_c}{\gamma_F}\right)\frac{\gamma_{FF}}{\gamma_F} + \frac{k\gamma}{k\gamma-1}\frac{\gamma_c}{\gamma} = \left(-\frac{\gamma_c}{\gamma_F}\right)\left(\frac{\gamma_{FF}}{\gamma_F} - \frac{k\gamma}{k\gamma-1}\frac{\gamma_F}{\gamma}\right)$. Therefore,

$$\text{sign}\left[\frac{M_D-1}{M_D}\left(\left(-\frac{\gamma_c}{\gamma_F}\right)\frac{\gamma_{FF}}{\gamma_F} + \frac{k\gamma}{k\gamma-1}\frac{\gamma_c}{\gamma}\right)\right] = \text{sign}\left[\frac{\gamma_{FF}}{\gamma_F} - \frac{k\gamma}{k\gamma-1}\frac{\gamma_F}{\gamma}\right]$$

Substituting yields

$$\frac{\gamma_{FF}}{\gamma_F} - \frac{k\gamma}{k\gamma-1}\frac{\gamma_F}{\gamma} = \frac{k^2-1}{k\gamma^2}\frac{R_{kc}}{R_k} + kR_k - kR_k = \frac{k^2-1}{k\gamma^2}\frac{R_{kc}}{R_k} < 0$$

Therefore, $\frac{M_D-1}{M_D}\left(\left(-\frac{\gamma_c}{\gamma_F}\right)\frac{\gamma_{FF}}{\gamma_F} + \frac{k\gamma}{k\gamma-1}\frac{\gamma_c}{\gamma}\right) < 0$.

For the third term, note that

$$\begin{aligned}\frac{k\gamma}{k\gamma-1}\frac{1}{M_D} + 2\frac{M_N-1}{M_N} - \left(1+\frac{k\gamma}{k\gamma-1}\right)\frac{M_D-1}{M_D} - \frac{M_N-1}{M_N M_D} \\ = \frac{2(M_N-M_D) + \frac{1}{k\gamma-1}M_N(2-M_D)+1}{M_N M_D}\end{aligned}$$

This third term is positive when M_D is sufficiently small. Explained alternatively, M_D sufficiently small means that an exogenous increase in agricultural production cannot cause households to decrease their labor allocations by too much; i.e., households cannot be too risk averse.[1] When this holds, then

$$\left(\frac{k\gamma}{k\gamma-1}\frac{1}{M_D} + 2\frac{M_N-1}{M_N} - \left(1+\frac{k\gamma}{k\gamma-1}\right)\frac{M_D-1}{M_D} - \frac{M_N-1}{M_N M_D}\right)\frac{\gamma_c}{\gamma} < 0$$

Since each term is negative, we have that $\frac{\frac{1}{w}\frac{d^2L_2}{d\overline{L}dA_1}}{\frac{dL_2}{dA_1}} < 0$. As $\frac{dL_2}{dA_1} < 0$, we have $\frac{d^2L_2}{d\overline{L}dA_1} > 0$. Since $w\overline{L}$ and $r\overline{M}$ enter the household's problem symmetrically, then $\frac{d^2L_2}{d\overline{M}dA_1} > 0$.

Input constraint. When only the input constraint binds, the first order conditions simplify to

$$\begin{aligned}(M_k)\quad & A_kF_{kM} &= (1+\lambda_M)r \\ (L_k)\quad & A_kF_{kL} &= w \\ (\ell)\quad & \frac{\mathbf{E}[u_\ell]}{\mathbf{E}[u_c]} &= w\end{aligned}$$

[1]Note that this condition is sufficient, but not necessary, to sign this term and also to sign $\frac{\frac{1}{w}\frac{d^2L_2}{d\overline{L}dA_1}}{\frac{dL_2}{dA_1}}$.

Note that the choice of leisure does not enter into the first order conditions for M_k or L_k. Substituting $M_2 = \overline{M} - M_1$ yields the following system of equations

$$\begin{aligned} A_1 F_{1M}(M_1, L_1) - (1+\lambda_M) r &= 0 \\ A_1 F_{1L}(M_1, L_1) - w &= 0 \\ A_2 F_{2M}(\overline{M} - M_1, L_2) - (1+\lambda_M) r &= 0 \\ A_2 F_{2L}(\overline{M} - M_1, L_2) - w &= 0 \end{aligned}$$

Stack the left hand sides into the vector FOC_M. Define the Jacobian $J_M \equiv D_{(M_1, L_1, \lambda_M, L_2)} \text{FOC}_M$. Applying the implicit function theorem yields $D_{(A_1)}(M_1, L_1, \lambda_M, L_2)' = -J_M^{-1} D_{(A_1)} \text{FOC}_M$. Some algebra yields

$$\begin{aligned} J_M &= \begin{pmatrix} A_1 F_{1MM} & A_1 F_{1ML} & -r & 0 \\ A_1 F_{1ML} & A_1 F_{1LL} & 0 & 0 \\ -A_2 F_{2MM} & 0 & -r & A_2 F_{2ML} \\ -A_2 F_{2ML} & 0 & 0 & A_2 F_{2LL} \end{pmatrix} \\ D_{(A_1)} \text{FOC}_M &= (F_{1M}, F_{1L}, 0, 0)' \\ \frac{dM_2}{dA_1} &= k_M A_2 F_{2LL} A_1 (F_{1L} F_{1ML} - F_{1M} F_{1LL}) \\ \frac{dL_2}{dA_1} &= -k_M A_2 F_{2ML} A_1 (F_{1L} F_{1ML} - F_{1M} F_{1LL}) \end{aligned}$$

where k_M is positive.[2] As $F_{2LL} < 0$, $\text{sign}\left(\frac{dM_2}{dA_1}\right) = -\text{sign}\left(F_{1L} F_{1ML} - F_{1M} F_{1LL}\right)$. This is negative whenever productivity growth on plot 1 would cause optimal input allocations, holding fixed the shadow price of inputs, to increase on plot 1. Similarly, $\text{sign}\left(\frac{dL_2}{dA_1}\right) = \text{sign}(F_{2LM}) \text{sign}\left(\frac{dM_2}{dA_1}\right)$. The labor response and input response on the second plot have the same sign whenver labor and inputs are complements on the second plot.

Labor constraint. When only the labor constraint binds, the first order conditions simplify to

$$\begin{aligned} (M_k) \quad & A_k F_{kM} && = r \\ (L_k) \quad & A_k F_{kL} && = (1 - \lambda_L) w \\ (\ell) \quad & \frac{u_\ell}{u_c} && = (1 - \lambda_L) w \end{aligned}$$

[2] $k_M = -\frac{1}{(A_1 F_{1LL}) A_2^2 (F_{2MM} F_{2LL} - F_{2ML}^2) + (A_2 F_{2LL}) A_1^2 (F_{1MM} F_{1LL} - F_{1ML}^2)}$. We make standard assumptions required for unconstrained optimization; second order conditions for unconstrained optimization imply k_M is positive.

Substituting $\ell = \overline{L} - L^O - L_1 - L_2$ and $L^O = \overline{L^O}$, and some rearranging yields

$$A_1 F_{1M}(M_1, L_1) - r = 0$$
$$A_1 F_{1L}(M_1, L_1) - (1 + \lambda_L) w = 0$$
$$A_2 F_{2M}(M_2, L_2) - r = 0$$
$$A_2 F_{2L}(M_2, L_2) - (1 + \lambda_L) w = 0$$
$$u_\ell \left(\sum_{k \in \{1,2\}} A_k F_k(M_k, L_k) + r(\overline{M} - M_1 - M_2) + w\overline{L^O}, \overline{L} - L^O - L_1 - L_2 \right) -$$
$$(1 + \lambda_L) w u_c \left(\sum_{k \in \{1,2\}} A_k F_k(M_k, L_k) + r(\overline{M} - M_1 - M_2) + w\overline{L^O}, \overline{L} - L^O - L_1 - L_2 \right) = 0$$

Stack the left hand sides into the vector FOC_L.

Additionally, it will be convenient to define the following derivatives of on farm labor demand on plot k, LD_k, with respect to the shadow wage w^* and productivity A_k, on farm input demand on plot k, MD_k, with respect to productivity A_k, and on farm labor supply, LS, with respect to the shadow wage w^* and consumption (through shifts to wealth) c. Let

$$\text{LD}_{kw^*} = \frac{A_k F_{kMM}}{A_k^2 (F_{kMM} F_{kLL} - F_{kML}^2)}$$
$$\text{LD}_{kA_k} = \frac{A_k F_{kM} F_{kML} - A_k F_{kL} F_{kMM}}{A_k^2 (F_{kMM} F_{kLL} - F_{kML}^2)}$$
$$\text{MD}_{kA_k} = \frac{A_k F_{kL} F_{kML} - A_k F_{kM} F_{kLL}}{A_k^2 (F_{kMM} F_{kLL} - F_{kML}^2)}$$
$$\text{LS}_{w^*} = -\frac{u_c}{u_{\ell\ell} - (1 + \lambda_L) w u_{c\ell}}$$
$$\text{LS}_c = -\frac{u_{c\ell} - (1 + \lambda_L) w u_{cc}}{u_{\ell\ell} - (1 + \lambda_L) w u_{c\ell}}$$

We make standard assumptions required for unconstrained optimization; these imply LD_{kw^*} is negative (labor demand decreasing in shadow wage), and LS_{w^*} is positive (labor supply increasing in shadow wage). We further assume LD_{kA_k} and MD_{kA_k} are positive (labor demand and input demand are increasing in productivity); an additional sufficient assumption for this is that F is homogeneous. We further assume LS_c is negative (labor supply is decreasing in wealth); an additional sufficient assumption for this is that u is additively separable in c and ℓ.

Next, define the Jacobian $J_L \equiv D_{(M_1,L_1,M_2,L_2,\lambda_L)}\text{FOC}_L$. Some algebra yields

$$J_L = \begin{pmatrix} A_1F_{1MM} & A_1F_{1ML} & 0 & 0 & 0 \\ A_1F_{1ML} & A_1F_{1LL} & 0 & 0 & -w \\ 0 & 0 & A_2F_{2MM} & A_2F_{2ML} & 0 \\ 0 & 0 & A_2F_{2ML} & A_2F_{2LL} & -w \\ \frac{d\text{FOC}_{L,\ell}}{dM_1} & \frac{d\text{FOC}_{L,\ell}}{dL_1} & \frac{d\text{FOC}_{L,\ell}}{dM_2} & \frac{d\text{FOC}_{L,\ell}}{dL_2} & -wu_c \end{pmatrix}$$

$$\frac{d\text{FOC}_{L,\ell}}{dM_1} = A_1F_{1M}(u_{c\ell} - (1+\lambda_L)wu_{cc})$$

$$\frac{d\text{FOC}_{L,\ell}}{dL_1} = A_1F_{1L}(u_{c\ell} - (1+\lambda_L)wu_{cc}) - (u_{\ell\ell} - (1+\lambda_L)wu_{c\ell})$$

$$\frac{d\text{FOC}_{L,\ell}}{dM_2} = A_2F_{2M}(u_{c\ell} - (1+\lambda_L)wu_{cc})$$

$$\frac{d\text{FOC}_{L,\ell}}{dL_2} = A_2F_{2L}(u_{c\ell} - (1+\lambda_L)wu_{cc}) - (u_{\ell\ell} - (1+\lambda_L)wu_{c\ell})$$

Applying the implicit function theorem yields $D_{(A_1)}(M_1, L_1, M_2, L_2, \lambda_L)' = -J_L^{-1}D_{(A_1)}\text{FOC}_L$. Some further algebra, and substitution, yields

$$D_{(A_1)}\text{FOC}_L = (F_{1M}, F_{1L}, 0, 0, (u_{c\ell} - (1+\lambda_L)wu_{cc})F_1)'$$

$$\frac{dL_2}{dA_1} = \text{LD}_{2w^*}\frac{\text{LD}_{1A_1} - \text{LS}_c(F_{1M}\text{MD}_{1A_1} + F_{1L}\text{LD}_{1A_1} + F_1)}{\text{LS}_{w^*} - (\text{LD}_{1w^*} + \text{LD}_{2w^*}) - \text{LS}_c(\text{LD}_{1A_1} + \text{LD}_{2A_2})}$$

$$\frac{dL_2}{d\bar{L}} = \text{LD}_{2w^*}\frac{1}{\text{LS}_{w^*} - (\text{LD}_{1w^*} + \text{LD}_{2w^*}) - \text{LS}_c(\text{LD}_{1A_1} + \text{LD}_{2A_2})}$$

$$\frac{dL_2}{d\bar{M}} = \text{LD}_{2w^*}\frac{r\text{LS}_c}{\text{LS}_{w^*} - (\text{LD}_{1w^*} + \text{LD}_{2w^*}) - \text{LS}_c(\text{LD}_{1A_1} + \text{LD}_{2A_2})}$$

$\frac{dL_2}{dA_1} < 0$; for interpretation, note that this expression is the derivative of labor demand on plot 2 with respect to the shadow wage, times the effect of the shock to A_1 on the shadow wage. The numerator of the latter is the effect the shock on negative residual labor supply through direct effects (LD_{1A_1}) and wealth effects, including through adjustments of labor and inputs ($-\text{LS}_c(F_{1M}\text{MD}_{1A_1} + F_{1L}\text{LD}_{1A_1} + F_1)$). The denominator of the latter is the derivative of residual labor supply with respect to the shadow wage, adjusted for wealth effects ($\text{LS}_{w^*} - (\text{LD}_{1w^*} + \text{LD}_{2w^*}) - \text{LS}_c(\text{LD}_{1A_1} + \text{LD}_{2A_2})$).

The signs of $\frac{d^2L_2}{d\bar{L}dA_1}$ and $\frac{d^2L_2}{d\bar{M}dA_1}$ are ambiguous. However, unlike the cases of input market failures or insurance market failures, here these second derivatives may have opposite signs. To see one example of this, consider a case where on farm labor and input demands are approximately linear in the shadow wage and productivity, and on farm labor supply is approximately linear in consumption, but exhibits meaningful curvature with respect to the shadow wage. In this case, $\text{sign}(\frac{d^2L_2}{d\bar{L}dA_1}) = \text{sign}\left(\frac{d}{d\bar{L}}\text{LS}_{w^*}\right)$ and $\text{sign}(\frac{d^2L_2}{d\bar{L}dA_1}) = \text{sign}\left(\frac{d}{d\bar{M}}\text{LS}_{w^*}\right)$. To focus on one case, larger households are less responsive to the A_1 shock ($\frac{d^2 L_2}{d\bar{L} dA_1} > 0$) if and only if they are on a more elastic portion of their labor supply curve ($\frac{d}{d\bar{L}}\text{LS}_{w^*} > 0$). That larger households, with more labor available for agriculture, or poorer households, who likely have fewer productive opportunities outside agriculture, would be on a more elastic portion of their labor supply curve is consistent with proposed models of household labor supply dating back to Lewis (1954). This motivates the prediction we focus on: that larger households should be less responsive to the A_1 shock, and richer households should be more responsive to the A_1 shock.