Essays at the Intersection of Environment and Development Economics

Essays at the Intersection of Environment and Development Economics

A dissertation presented

by

Elizabeth Ruth Walker

to

The Department of Public Policy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Public Policy

Harvard University

Cambridge, Massachusetts

April 2015

Dissertation Advisor:
Professor Rema Hanna

Author:
Elizabeth Ruth Walker

Essays at the Intersection of Environment and Development Economics

Abstract

The three essays in this dissertation explore how households in Southern Africa interact with and rely upon environmental resources. The first chapter examines the relationship between irrigation dam placement and local infant health outcomes. Irrigation dams can enable farms to harness considerable water resources, and this has been critical to increasing the global food supply. However, irrigation consumes 70 percent of global water resources and returns polluted water back into river systems. I examine the effect of irrigation dams on water pollution and infant health outcomes in South Africa. To remove bias associated with non-random dam placement, I utilize an instrumental variables approach that predicts dam placement using geographic features and time-varying policy changes. I find that each additional dam within a district increases both water pollution and infant mortality. In districts downstream from dams, alternately, dams generate smaller water pollution effects and reduce infant mortality, though magnitudes are much smaller. I argue that this pattern is consistent with pollution-induced health costs that outweigh economic benefits within the districts that receive dams. Downstream, however, pollution generates smaller costs and the economic benefits dominate. Exploring other plausible channels through which irrigation dams may affect infant mortality, I find that while irrigation dams generate substantial effects on district employment and small effects on migration, these factors do not appear to explain the health outcomes observed. Instead, the results suggest that water pollution and reduced water availability may contribute to higher infant mortality near agricultural activity.

The second chapter in this dissertation explores technology adoption of trees with environmental benefits, in the context of a field experiment in Zambia. As context, many technology adoption decisions are made under uncertainty about the costs or benefits of following through

with the technology after take-up. As new information is realized, agents may prefer to abandon a technology that appeared profitable at the time of take-up. Low rates of follow-through are particularly problematic when subsidies are used to increase adoption. This chapter uses a field experiment to generate exogenous variation in the payoffs associated with taking up and following through with a new technology: a tree species that provides fertilizer benefits to adopting farmers. Our empirical results show high rates of abandoning the technology, even after paying a positive price to take it up. The experimental variation offers a novel source of identification for a structural model of intertemporal decision making under uncertainty. Estimation results indicate that the farmers experience idiosyncratic shocks to net payoffs after take-up, which increase take-up but lower average per farmer tree survival. We simulate counterfactual outcomes under different levels of uncertainty and observe that subsidizing take-up of the technology affects the composition of adopters only when the level of uncertainty is relatively low. Thus, uncertainty provides an additional explanation for why many subsidized technologies may not be utilized even when take-up is high.

Finally, the third paper in my dissertation explores the role that mineral wealth has on local economic outcomes. Mineral wealth has been central to the development of the South African economy. However, we have little evidence regarding how it has affected employment and poverty. This paper explores how within-country variation in mineral wealth affects district-level outcomes. Using data on mineral deposits and historical world prices, I construct a plausibly exogenous variable reflecting district-level aggregate mineral wealth, and I use variation in this index to evaluate economic outcomes over three Census rounds (1991, 1996, 2001). In the short run, I show that positive shocks to aggregate mineral wealth generate higher employment, largely driven by increased mining employment. This increase in mining is accompanied by reductions in agricultural employment and slight reductions in manufacturing employment. On average, adult age individuals also report working more hours and earning higher salaries in districts experiencing higher mineral prices. In sum, mineral wealth shocks generate benefits to districts, with less households below the poverty line.

Contents

List of Tables

List of Figures

Acknowledgments

I am incredibly grateful to the people that supported me throughout my PhD.

This dissertation would not have been possible without the guidance and encouragement of my advisors Rema Hanna, Rob Stavins, Rohini Pande and Kelsey Jack. Their mentorship made my PhD an enjoyable, productive and deeply rewarding learning experience. In particular, I am thankful for Rema's incredible patience, technical direction and life advice as I produced my job market paper; for Rob's mentorship on how to teach a great economics class; for Rohini's guidance on how to make research policy-relevant; and for Kelsey's help in connecting research to theory and her practical advice on performing field experiments.

Many colleagues and friends also generously took the time to help me during the dissertation process, including Chris Avery, Bill Clark, Nancy Dickson, Ariel Stern, Rich Sweeney, Sara Lowes, Megan Bailey, Martin Abel, Deanna Ford and the participants in the Harvard Environmental Economics and Development lunches. I am also grateful to Kelsey, Paulina Oliva, Sam Bell, Jonathan Green, and our team in Zambia for their collaboration on our field experiment. Nicole Tateosian and Jason Chapman also were very supportive, and kept me organized and productive.

For financial support, I am grateful to the Vicki Norberg-Bohm Fellowship, the Crump Fellowship, and the Jennifer Perini and Jim Cunningham Dissertation Fellowship. I am grateful for the assistance of the South Africa Department of Water Affairs, particularly Dr. Michael Silberbauer who provided data, insights and support.

For keeping me happy and sane, I thank my running girlfriends for thousands of miles over the last six years. For love and inspiration, I thank my family Dan, Kitty, Maria, Katie and Jay.

To my family

Chapter 1

Irrigation Dams, Water and Infant Mortality: Evidence from South Africa

1.1 Introduction

Water is integral to food production, withdrawing 70 percent of global water resources (FAO, 2005). However, as agriculture intensifies in fast-growing, high food-demand regions of China, India, and Sub-Saharan Africa, growing concerns related to agricultural water use and human health may be warrented (e.g., Tilman et al., 2002; Seckler, 1999). Specifically, agricultural intensification increases yields but relies on water-consuming irrigation techniques and water-polluting agrichemical inputs. Estimating the effects of intensified agriculture on health outcomes is therefore instructive to designing policies that account for these potentially costly health externalities.

Empirically, linking agricultural water pollution to health outcomes is difficult, given that pollution originates from private, non-point sources. Moreover, households near polluted water sources may be systematically different from those further away, and increased pollution may alter behaviors or induce migration (Chay and Greenstone, 2003). As a consequence, most well-identified evidence of the causal effects of water pollution on health focuses on observable improvements to water infrastructure, such as switching from unprotected to protected springs (Kremer, 2011) or from natural to piped water (Jalan and Ravallion, 2003; Galiani et al., 2005).

Only a handful of recent papers incorporate data on surface water quality in order to identify plausibly causal relationships between water pollution and health outcomes (Ebelstein, 2012; Brainerd and Menon, 2014).

This paper builds on the existing literature by investigating a new and important channel: the effect of irrigated agriculture on infant mortality. South Africa is water stressed, and agriculture uses 60 percent of country-wide water resources, compared to 10 percent for urban and domestic use (CSIR, 2010; Oberholster and Ashton, 2008). I use the construction of irrigation dams as a proxy for agricultural water use and estimate the effect of an increase in irrigation dams on infant mortality. Irrigation dams, which henceforth I refer to as "dams," are a useful indicator of intensive water use in this context, where low, erratic rainfall and depressed groundwater levels often make dam construction a primary means for increasing yields (Blignaut et al., 2009). The dams are predominantly small, privately owned, and constructed to support commercial agriculture. Moreover, South Africa has more dams than any other country in Africa, and the number more than doubled between 1980 and 2010, to over 3,000 dams (Dam Safety Office, 2013).[1]

Estimating responses to new dam construction is confounded by non-random dam placement. On the one hand, governments or firms may target locations that are agriculturally more productive, growing faster, or politically better connected than those that do not receive dams. Alternately, dams may be situated in lower functioning districts to spur growth or expand production. As a result, a simple comparison of districts that receive dams with those that do not may be biased by other factors affecting outcomes. To control for endogenous dam placement, I adopt a variation of the approaches in Pande and Duflo (2007) and Strobl and Blanc (2013). I instrument for dam placement by interacting two sources of variation: spatial differences in river gradient suitability for dams and policy changes which affected dam placement over time.[2] I also present and discuss OLS estimates with fixed effects, which demonstrate the same pattern.

The instrumental variables (IV) estimation relies on shifts associated with Apartheid, the

[1] See Figure A.1 and Figure A.2.

[2] Dams are less likely to be constructed on steep rivers, given increased erosion.

political system in place until 1994. Apartheid policies removed millions of black South Africans to marginal "homeland" regions. Dams were rarely constructed within or near homelands until after Apartheid. Thus, the IV strategy interacts a river gradient variable with a step variable reflecting the increasing likelihood that dams were placed in former homelands. This restricts the identification to variation within districts that are politically, economically and geographically more similar. In other words, the IV compares outcomes within former homeland districts that had river gradients more desirable for dams to those with steeper, less desirable river gradients, while controlling for district-specific geographic characteristics.

Restricting the IV analysis to responses *within* former homelands has three empirical advantages. First, evaluating outcomes within former homelands enables plausibly separating the water-induced health effects from those caused by direct chemical exposure associated with working on plots sprayed with agrichemicals. In South Africa, mechanization during the 1950s drastically reduced agricultural labor to a subset of higher-skilled managers (Platsky and Walker, 1985). As of 2001, less than ten percent of the population was involved in agriculture, and in the former homelands less than one percent participated (Census, 2001). Second, the former homelands are more similar to other developing countries in Africa, where many households rely on surface and ground water. This makes the results more applicable to these similar countries. Finally, households within homelands were less able to migrate, which reduces the probability of differential migration driving results.

Given its dual effects on economic productivity and the environment, the effect of dam-induced irrigation on infant mortality is *a priori* ambiguous. Economically, dams affect income through improved labor opportunities and changes to local markets. Even if households are not directly involved in agriculture, there are likely to be increased post-production labor opportunities within the district. Evidence on small irrigation dams in South Africa finds that these dams generate small increases in crop production within districts (Strobl and Blanc, 2013).[3] New labor opportunities could either reduce infant mortality by increasing household

[3]This is in contrast to the research on large dam construction in India and Africa, which suggests that large dams reduce agricultural production within a district while increasing it downstream (Duflo and Pande, 2007; Strobl and Strobl, 2010). Large dams are those at least fifteen meters in height or three million cubic meters in volume. For large dams, the channel through which productivity losses arise is somewhat different. In India, Duflo and Pande suggest that productivity losses arise as a result of community displacement and soil salinity. In

income or increase infant mortality through reductions in time spent caring for the child.

A less considered channel is the concurrent water-related changes that arise as a result of dam-induced irrigation. Irrigation reduces water availability and increases water pollution within districts. Household responses to water-related changes depend on their ex ante water source, with those that rely directly on nearby rivers or groundwater being most vulnerable. Households should unambiguously switch water sources or migrate in response to pollution or low water availability, but this behavior change may be attenuated if pollution is unobservable, if households are unaware of possible adverse effects, or if marginal willingness to pay for improved water sources is low (Kremer et al., 2011; Greenstone and Jack, 2014). Evidence also suggests household characteristics, such as education and income, influence household choices and outcomes in response to water pollution (Jalan and Ravallion, 2003; Kremer et al., 2011).

Despite limited causal evidence, the water channel may be important. Empirical public health research suggests an observable correlation between agricultural water pollution and birth defects in the United States (Croen et al., 2001; Winchester et al., 2009). Building on this, recent causal research from India links agrichemical exposure to increases in poor infant and child health outcomes (Brainerd and Menon, 2014). However, in India, unlike South Africa or the United States, over half of the labor force is involved in agriculture, which means direct agrichemical exposure may contribute to results. Given that dams reduce water availability, research by Field et al. (2011) is also relevant. Using a natural experiment from Bangladesh, the authors estimate that infant and child mortality increase as a result of lower water availability and switching to farther water sources. Finally, within South Africa, countless government reports and news articles document deteriorating water quality and water availability (e.g., DEAT, 2006; Blaine, 2013; UN Water, 2011).[4]

This paper measures the *net* impact of dams on infant mortality, which includes the combined effect of economic, environmental, and other unobservable channels. I find that

South Africa, the data includes dams over five meters in height and the average surface area of constructed dams is 0.465 square km. This is much smaller than the average district area of 3,500 square km. As a result, population displacement by dams is less concerning.

[4]For example, a United Nations report estimates that river pollution, namely nutrient enrichment from agriculture and acidity from industry, has left nearly two-thirds of freshwater species at risk of extinction in South Africa (UN Water, 2011).

each additional dam leads to an average 6 percent increase in infant mortality, off an already high average infant mortality of 4.8 percent. This result is in line with research demonstrating that various types of water pollution and variation in water sources significantly affect infant mortality.[5]

I then examine plausible channels through which dams may affect health. First, I show that newly constructed dams increase water pollution, as observed using chemical indicators. For example, average nitrate concentrations increase almost 50 pecent and other related indicators have smaller but significant increases of 3 percent to 11 percent. In addition, nitrate and total dissolved salt levels are significantly more likely to exceed drinking water standards downstream from new dams. The results corroborate existing case-based research from South Africa showing that small irrigation dams cumulatively reduce water flow and generate concentrated water pollution (e.g., O'Keeffe et al., 1990; O'Connor, 2001; Mantel et al., 2010).

To shed light on other channels, I evaluate employment and migration changes after Apartheid. New dams generate employment, despite falling national employment. While employment rates for men are almost double those for women, dam construction induces slightly greater employment gains for women. This could affect infant health if female caregivers are less attentive to infants. However, stratifying the sample by age demonstrates that employment gains are small for women under 30 and near zero for women over 50, the two age segments most likely to be primary caregivers. Finally, I investigate migration to ensure that selective out-migration is not altering district compositions in a way that could generate the observed infant mortality result. I show that in-migrants are more likely to have piped water, at least a high school education, and an income above the poverty line. These factors are unlikely to increase district infant mortality. Still, out-migrants have similar characteristics. To rule out that this influences the outcome, I show that dam construction does not alter average district characteristics significantly between 1996 and 2001.

This paper contributes to the empirical literature on the determinants of pollution and the

[5] For examples of papers linking water to infant mortality, see Clay et al. (2010) for lead water pollution, Galiani et al. (2005) for transition to piped water sources, Cutler and Miller (2005) for water filtration/chlorination, Field et al. (2011) for switching to less convenient water sources, and Brainerd and Menon (2014) for agrichemical water pollution.

effects of pollution on health. Most estimates of infant health outcomes in response to pollution exposure focus on air pollution, both in the developed world (e.g., Chay and Greenstone, 2003; Currie and Walker, 2011) and developing world (e.g., Borja-Aburto et al., 1997; Jayachandran, 2009; Arceo-Gomez et al., 2012). There are far fewer estimates relating infant health to water pollution given lack of water quality data, the difficulty in obtaining precise estimates from water samples, and the fact that developed countries have long had nearly universal access to clean drinking water sources (as discussed in Olmstead, 2009).[6] This paper contributes new evidence that in places with limited water availability and insufficient water treatment infrastructure, there may be a direct trade-off between agricultural intensification and infant health.

The paper is divided into seven sections. Section 1.2 provides background on irrigated agriculture, water, infant mortality, and the South African context. Section 1.3 describes the data sources. Section 1.4 presents the empirical strategy, Section 1.5 describes results and Section 1.6 explores alternative channels. Section 1.7 concludes.

1.2 Background

In this section, I first describe the link between irrigated agriculture and water quality. Next, I discuss the relationships between water and infant health outcomes. Finally, I provide relevant institutional and policy details.

1.2.1 Irrigated Agriculture and Water Quality

In South Africa, semi-arid environmental conditions make commercial agriculture highly reliant on dams. On average, South Africa receives 500 mm of rain per year, much lower than the global average of 860 mm per year. This rainfall is highly variable and spatially concentrated in the east, falling on only 14 percent of arable land (Davies et al., 1995). Groundwater is also a small portion of water use, given low groundwater levels and recharge rates (Hughes, 2004; CSIR,

[6] Historical estimates, however, imply that the direct human health benefits of improved drinking water quality are sizable: for example, Cutler and Miller (2005) estimate that introducing water disinfection in U.S. cities accounted for three-quarters of the decline in infant mortality during the early twentieth century.

2010). Dams hedge water-related risk and ensure enough water to irrigate crops. While total yields generated from irrigated agriculture are unknown, commercial agriculture generates 95 percent of total output in monetary terms and irrigated crops contribute substantially to this output (FAO, 2005).

Irrigated agriculture alters water systems in three ways that are important to health. First, irrigation dams siphon water from rivers and transport it via canals to irrigate crops. Some portion of the water is consumed by crops or evaporates, and some returns back to river systems.[7] While this regulates water flow in irrigation channels, it reduces water flow and availability to households downstream.[8] Second, irrigation dams increase standing water, which can act as a breeding ground for mosquitos and increase vector-borne infections. Third, water return flows are often contaminated with fertilizer nutrients and pesticides (e.g., Tredoux et al., 2001; Walmsley, 2003). Agrichemical pollution will be exacerbated when total water is reduced.

Changes in water quality can be measured using chemical indicators, which reflect nutrient content and overall quality. Nitrate, phosphate and potassium are the primary components of fertilizer and the indicators most associated with agricultural pollution. Fertilizer run-off also increases total dissolved salts (National Assessment Report, 2002). In South Africa, fertilizers that contain gypsum will lead to elevated sulfate levels as well (Huizenga, 2011).

1.2.2 Water and Infant Mortality

Both water pollution and reduced water access can generate adverse health outcomes.[9] First, existing evidence from the United States demonstrates a correlation between agricultural water pollution and birth defects, the leading cause of infant mortality in this context (e.g., Winchester et al., 2009). Both fertilizer nutrients and pesticides have been linked to developmental defects.

[7]Estimating how much water is consumed by crops rather than returned to the water system is difficult, given its dependence on the rate of evapotranspiration, weather/season, crop types, and other factors. A global average suggests that 70 percent of water removed from rivers for agriculture is consumed by the water system (Cosgrove and Rijsberman, 2000; Shiklomanov, 1999).

[8]For examples, see O'Keeffe et al. (1988), O'Connor (2001), or Mantel et al. (2010).

[9]Although the outcome measure is infant mortality, early life outcomes can be indicative of broader community health and of health impacts that arise later in life. See Almond and Currie (2010) or Currie and Vogl (2012).

For example, high nitrate concentrations may directly induce infant mortality through "blue baby syndrome."[10] Similarly, various pesticides have been associated with adverse fetal and infant health (e.g., Munger, 1992; Garry et al., 1996). Fertilizer nutrients, particularly when combined with slow flowing water, also cause toxic algal growth. Toxic algae pose a serious threat to several water streams in South Africa (e.g., van Ginkel, 2011; Oberholster, 2009). Toxic algae can cause infant death when ingested. Alternately, observable algae may induce households to seek other water sources (WHO, 2009).

Reduced water access prevents households from exercising appropriate hygiene. Water is necessary for hand washing, food preparation, and maintaining toilets. Without water, water-borne and water-washed diseases (e.g., trachoma, scabies, shigella) are more likely. Furthermore, lack of water leads caregivers to devote more time to searching and collecting water and less time to children. For these reasons, lower water usage can increase infant mortality (Pruss et al., 2002). The effect of unreliable water access on infant health is huge: the Millennium Development Goals document that lack of safe drinking water, open defication, and poor hygiene contribute an estimated 88 percent of all deaths for children under the age of five (2007).[11]

South Africa is an important context in which to study water and infant health. At 50 out of 1000 births, infant mortality in South Africa during the 1990s was ten times that found in developed countries. Babies may be vulnerable to infant mortality given mother characteristics and the high HIV/AIDs prevalence. For example, only 66 percent of infants from former homeland districts were delivered in hospitals and only 45 percent were breastfed within the first hour of life. Moreover, relying on breast milk exclusively is rare. While most mothers breastfed their infant at least once, only seven percent exclusive breastfeed in the first six months. This is important because mortality from environmental factors is more likely after

[10]Blue baby syndrome arises because nitrate ingestion leads to decreased oxygen carrying capacity in the blood. Infants under six months are particularly susceptible to this condition, given low levels of the enzymes required to reverse this condition (Knobeloch et al., 2000). Evidence suggests this outcome is rare (Fewtrell, 2004), but does occur. See also Manassaram (2006) or Avery (1999) for reviews of the literature on nitrates and infant health outcomes.

[11]The weights of different health channels are unknown, given the difficulty in isolating them (e.g., as discussed in Hunter, 2010).

children are weened from breast milk (Esrey et al., 1991). Given that the median age of infant death is five months in this sample, exposure to environmental hazards could contribute to mortality.

1.2.3 Institutional and Policy Context

Between the 1950s and the 1970s, several policy shifts forced millions of black South Africans into independent homelands, yielding *de facto* segregation and limited mobility.[12] Water rights were assigned with land rights, which were reserved for whites (Rose, 2005). The ethnic homelands were fragmented resettlement zones on the rural outskirts of productive regions, as shown in Figure 1.1. The homelands negotiated with the government for rights to water and were underserved as a result. Rivers were a primary domestic water source for many households (Funke et al., 2007). When the first democratic government was elected in 1994, the homelands were incorporated into existing provinces and the government made efforts to remove biases against dam placement in former homelands.

[12] The homelands were 13 percent of land area but approximately 40 to 50 percent of the total population as of the 1990s (Anderson, 2005; Thompson, 2001).

Figure 1.1: *Map of Former Homeland Areas*

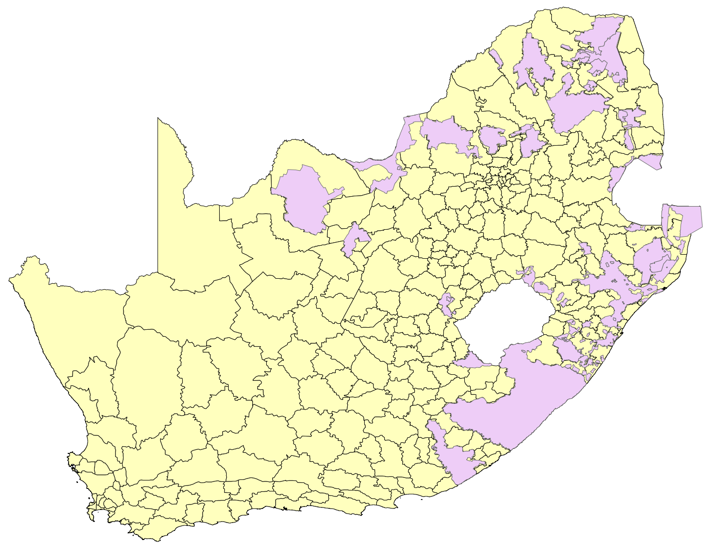

The pink regions reflect former homeland areas under Apartheid. The homeland areas were integrated back into South Africa after the first democratic government was elected in 1994. The analysis relies on magisterial district boundaries (the black lines).

More formal policy, laid out in the National Water Act of 1998, nationalized water rights. As a result of the government transition and reforms, millions gained access to improved water sources. Despite increased access to improved water sources, competitive water pricing in the 1990s kept prices high. As of 1995, about 20 percent of households used river water as their primary domestic source, though more used it as a secondary source.

Despite improvements, twenty-seven percent of households reported inadequate water supply during 1995 and 1996. Furthermore, even piped water sources sometimes retain pollution from natural sources. Government reports document non-functioning treatment facilities as a serious problem, in some cases *increasing* pollution levels because of poor practices (e.g., CSIR, 2011; Igbinosa, 2009; Silverbauer, 2009). Agricultural runoff is regulated, but regulations are difficult to enforce.

1.3 Data

In this section, I discuss the data sources utilized for this analysis. Additional details regarding the data are provided in the Appendix.

1.3.1 Dam and Hydrological Data

The dam data was collected from the Dam Safety Office within the Department of Water Affairs. All dams that stand at least 5 meters in height and 50,000 cubic meters in capacity are required to register with the Dam Safety Office. The dataset provides information on the dam owner, designer, catchment area, surface area, wall type, wall height, drainage area, main purpose(s), contractor, wall height, capacity, GPS coordinates, spillway area, catchment area, and completion date for dams. Eighty percent of constructed dams report a purpose of "irrigation" and I restrict the data to these. In 1980, 131 districts (37 percent) had no dams and 83 districts (23 percent) had over five dams. By 2010, only 89 districts (25 percent) had no dams and 147 districts (43 percent) had greater than five.

As shown in Table A.1, less than five percent of dams report multiple purposes. Ninety-five percent of dams in the data are embankment dams, constructed using compacted earth materials to create a water barrier. Water is stored behind the wall in a reservoir and can be transferred nearby using canals. Embankment dams rely on gravity-based flow, thus a gentle gradient ensures enough slope for water flow. However, slope increases velocity, which in turn increases soil erosion and sedimentation. Soil erosion creates cloudy water, suffocates riverbed ecosystems and erodes downstream canals, among other negative consequences. As a result, steep gradients are undesirable for irrigation dams, which I exploit in the identification.

I also utilize water quality and river network data from the South Africa Department of Water Affairs. The water quality data is collected by the National Monitoring Programme and includes 1,672 monitoring stations. Monitoring began as a way to ensure irrigation viability, but today monitors water quality for climate change, domestic use, agriculture, and industry. Most stations had one reading per station per month, which I use to construct station-month-year averages and district-year averages. For the main analysis, I restrict the data to only stations in place as of 1990, in order to avoid biases associated with new stations being placed in response to poor pollution.[13] Although homeland areas were underserved under Apartheid, about 25 percent of stations existed in these districts. Descriptive statistics of chemical indicators are shown in Appendix Table A.2. As the table shows, South African water quality is already poor, with about 10 percent of readings for several indicators above drinking water standards.

To construct the river gradient instrument, I use a complete mapping of the river network from the Department of Water Affairs. About 50 percent of rivers flow perennially, with the rest seasonal/dry (35 percent) or unknown (15 percent). I extract perennial river segments, given that dams are more likely to be placed on these rivers (Strobl and Blanc, 2013).[14] For the subset of river pixels within the district, I match river pixels to land gradient data and calculate the fraction of river pixels which are steep, defined as greater than six percent slope. Figure 1.2 demonstrates that dams are concentrated in areas with gentle to moderate, rather than steep, river gradients.

[13]The results are robust to alternative methods of collapsing the data, including using the means, using only stations with at least seven years of data, and trimming the data (removing top and bottom 1%).

[14]Dams can be constructed on seasonal rivers, but it is less likely given unpredictable flow. I also control for total river length (seasonal and perennial) in the instrumented regressions.

Figure 1.2: *Map of Dam Locations and Average River Gradient by District*

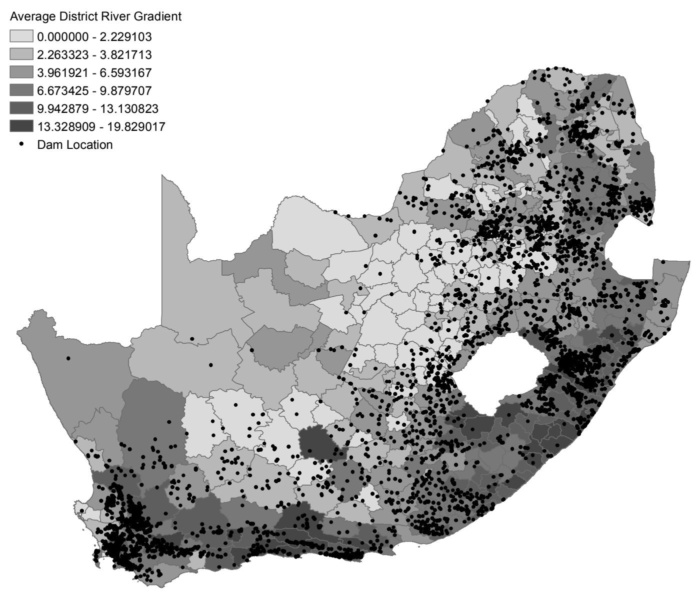

Each dot on the map indicates the location of a dam. Most irrigated grain agriculture is concentrated on the eastern side of the country.

I coded the upstream/downstream relationships among districts based on the direction of flow for each river that crossed a district boundary. Each river segment contains an order code, which increases as a river flows downstream. I rely on this, but verify it using elevation changes along the river. In Figure 1.3, for example, three districts have rivers that flow into the main district, thus there are three upstream districts. In cases where two districts had rivers going in both directions, I marked these district pairs as adjacent and do not include them in the current analysis.

Figure 1.3: *Example of Data*

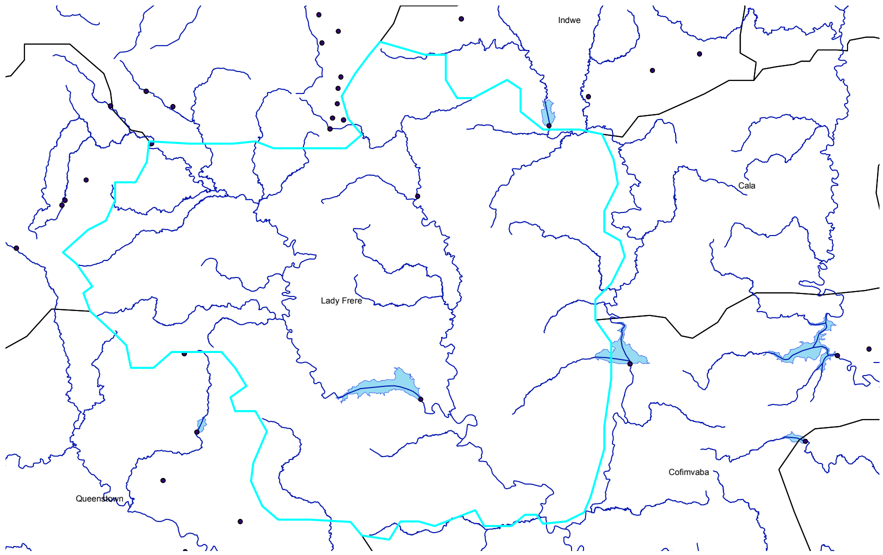

The above figure is a snapshot of the data. The turquoise line is the outline of one district in the sample, Lady Frere. The black lines reflect district boundaries, the black dots indicate dams, and the dark blue lines are the river network. For larger dams, I have the polygon locations, shown in light blue. The surface area of dams in this setting tends to be a small fraction of the total district area.

1.3.2 Survey Data

I pool three rounds of infant birth and death data from the Demographic and Health Surveys. The survey was collected in 1987, 1998, and 2003 and aims to be nationally representative of women between the ages of 15 and 49. I construct a record of child births and deaths based on the detailed fertility history provided by women in the sample. Mothers report the date of birth for each child, whether the child died, and if so, the date of death.[15] In the data, each child birth becomes an observation, and I construct an indicator that equals one if the child died within the first year of life. This provides an estimate of child mortality for a given year. This generated a dataset of 31,807 births and deaths, spread over 22 years.

Household-level data from The October Household Survey (OHS) allows me to further

[15]For the 1987 survey, I only use data on the last two children born, given this data is considered most accurate (Phillips, 1999).

explore behavioral responses to dams. The nationally-representative surveys were conducted to inform post-Apartheid reconstruction and development, and the surveys are similar to the World Bank Living Standards Measurement Surveys. The survey asks questions related to access to water and use of different water sources. It also asked labor-related questions, including employment, hours worked, and salary.

Finally, I utilize the more comprehensive 10 percent population sample of the 1996 and 2001 Census to evaluate migration patterns in response to dams. The limitation with Census data is that I cannot construct a panel given only two rounds of reliable information.

1.3.3 Districts as Unit of Analysis

I link geographic and survey data sources using the magisterial district boundaries that were in place before the end of Apartheid. Magisterial boundary data was obtained from Global Administrative Boundaries. Although boundary re-demarcation occurred several times after Apartheid, I rely on spatial coordinates of subsequent survey enumeration areas to match each survey round with the former boundary areas.[16] There were 354 magisterial districts in South Africa. A magisterial district has an area of about 3,500 square km. Across districts, the population grew from about 82,000 people per district in 1980, to 126,000 people per district in 2000.

To construct former homeland districts, I match homeland maps to magisterial district maps. I define districts that had at least 30 percent former homeland area as a "homeland" district. A map of homeland areas, shown in Figure 1.1, confirms that homelands tended to be concentrated in a subset of all districts. This fact is confirmed in Figure A.3. Because most districts had either no homeland area or a significant portion, the main results hold across different cut-off values. Table 1.1 also demonstrates that homelands tend to be steeper, drier, and more densely populated. Thus, they were somewhat less likely to receive dams, though many are still constructed during the period of my analysis.

[16]For some survey data, the former magisterial district is recorded at the time of surveying. This served as a check on the matching.

Table 1.1: *Geographic Characteristics of All and Former Homeland Districts*

Variable	All Districts Mean (1)	Homelands Mean (2)
Fraction of River Gradient Pixels 0-1.5%	0.294	0.197
	[0.202]	[0.173]
Fraction of River Gradient Pixels 1.5-3%	0.22	0.189
	[0.088]	[0.093]
Fraction of River Gradient Pixels 3-6%	0.189	0.211
	[0.079]	[0.070]
Fraction of River Gradient Pixels > 6%	0.286	0.392
	[0.233]	[0.244]
District Elevation	1.016	817
	[498]	[424]
District Slope	9.14	11.88
	[6.26]	[6.29]
River Length (km)	569.35	418.19
	[936.68]	[457.42]
District Population (1996 Census)	103,696	149,632
	[126,858]	[89,299]
Number of Districts	354	100

Notes:
1. Each observation reflects one district. The first column includes all districts. The second column is only former homeland districts.
2. The means for all districts are reported in column (1), and the means for only former homeland districts are reported in column (2). Standard deviations reported below in brackets.
3. The geographic variables were constructed using ArcGIS, and the population was estimated using the 10 percent sample of the 1996 Census.

1.4 Identification Strategy

1.4.1 Instrumental Variables

The analysis is designed to measure the effects of irrigation dams on health outcomes. If irrigation dam placement is randomly assigned, then the following regression identifies the relationship between dams and outcomes:

$$y_{idt} = \gamma_o + \gamma_1 D_{dt} + \gamma_2 D_{dt}^U + year_t + \delta_d t + \lambda_d + \varepsilon_{idt} \tag{1.1}$$

where y_{idt} is the outcome of interest for observation i in district d and year t. The variable D_{dt} is the cumulative number of operational dams in the district as of year t, and D_{dt}^U is the cumulative number of operating dams in all upstream districts.[17] The remaining right hand side variables control for unobservable heterogeneity: year fixed effects, district-level trends, and district fixed effects.

As discussed above, this strategy is biased, given that dams are possibly placed in locations that are systematically different from other locations along dimensions that are difficult to entirely control for in a regression framework. If factors that affect dam construction are also likely to affect infant mortality, this generates correlation between the true error term and the dam coefficient, resulting in a biased estimate of γ_1, the effect of an additional dam on health outcomes in that district. The direction of the bias depends on which factors drive dam placement. For example, dams are more likely to be placed in areas with growing agriculture, industry, or resources, and newly constructed dams are accompanied by improvements in access to food, health or labor opportunities, then γ_1 will be lower or more negative, as dams reduce infant mortality.

Instead, I instrument for dam construction. First, I construct a variable that denotes the fraction of steep river gradient pixels within a district (greater than six percent slope). To generate time-varying predictions, the steep river gradient variable is interacted with a variable

[17] I also consider the possibility that dams have nonlinear effects on outcomes. The residuals of the regression above are roughly normally distributed (both with and without trends). I also perform the analysis using logs and using quadratic terms, both of which reasonably describe the data. These results available by email.

reflecting dam placement policies. After Apartheid ended and the Water Act was enacted, nationalized water policy favored dam placement broadly (Strobl and Blanc, 2013). To denote this, I create a step function that gives former homeland districts a value of 0 prior to 1994, 1 from 1994 to 1997, and 2 from 1998 to 2010. As a result, the policy variable is an interaction between a location (former homeland area), denoted H_d, and a year, denoted A_t. I interact time-invariant river gradient with the time-varying policy variable, $H_d A_t$. This yields the following first stage:

$$D_{dt} = \beta_o + \beta_1 \left(Steep_d * H_d * A_t\right) + \beta_2 \left(X_d * H_d * A_t\right) + \left(Steep_d * year_t\right) + H_d * A_t + \gamma_p + \varepsilon_{dt} \tag{1.2}$$

where D_{dt} equals the cumulative number of dams in district d and year t, and $Steep_d$ is the fraction of river pixels that have a steep gradient (greater than 6 percent).[18] I also include geographic controls, X_d, including sum of all seasonal and perennial rivers, district elevation categories (0-500 m, 500-1000 m, >1000 m) and district gradient categories (1.5-3%, 3-6%, and >6%). I interact the geographic controls with the policy variable to allow the effect of each control to vary with the policy shifts. The steep river gradient-year interaction term allows for the time-varying shocks, like weather or new technologies, to affect steep river gradients differentially. I also control for district or province fixed effects, γ_p.[19]

The first stage predicts the number of dams in the district and the number of dams upstream from the district.[20] The second stage equation uses these predicted values to estimate health outcomes:

---

[18]Previous versions of this paper used a combination of three river gradient instruments, based on different river gradient slopes (similar to Pande and Duflo, 2007). However, for the small dams in this setting, a single variable - steep river gradient - is sufficient to predict dam placement.

[19]In all cases except infant mortality, I use district fixed effects. Given that infant mortality is a low probability event and the number of observations is limited, I rely on province fixed effects and fraction homeland controls (in addition to all geographic controls).

[20]To predict $\widehat{D_{dt}^U}$ using equation (2), I take the sum of all the upstream predicted dams. For controls, I use the average of upstream districts for variables that enter as averages (elevation, slope, percent homeland) and the sum for variables that enter cumulatively (river length).
$\widehat{D_{dt}^U} = \widehat{D_1^U} + \widehat{D_2^U} + \widehat{D_3^U} + \widehat{D_4^U} = \beta_o + \beta_1 \left(Steep_d^{up1} H_d^{up1} A_t^{up1}\right) + \beta_2 \left(Steep_d^{up2} H_d^{up2} A_t^{up2}\right) + \beta_3 \left(Steep_d^{up3} H_d^{up3} A_t^{up3}\right) + \beta_4 \left(Steep_d^{up4} H_d^{up4} A_t^{up1}\right) + \beta_2 X_d^{upave} + year_t + \gamma_p + H_d^{upave} A_t + \varepsilon_{dt}$

$$y_{idt} = \gamma_o + \gamma_1 \widehat{D_{dt}} + \gamma_2 \widehat{D^U_{dt}} + \gamma_4 Z_d + \gamma_5 Z^U_d + \varepsilon_{idt} \tag{1.3}$$

where $\widehat{D_{dt}}$ is the number of dams predicted in the district, $\widehat{D^U_{dt}}$ is the number of dams predicted upstream, Z_d is the set of district-level controls, and Z^U_d is the set of controls for the upstream district. The outcome variable, y_{idt}, is infant mortality in the primary analysis, and then other outcomes in subsequent analysis (water quality, water source, and employment variables). In each case, the IV regression captures the local average treatment effect of dams. In other words, the estimated coefficient, γ_1, reflects the average effect of each additional dam on districts that receive dams as a result of river gradient desirability.

In the results that follow, I report both IV and OLS coefficients. While the regressions are run on the same observations, the coefficients for the two specifications rely on somewhat different observations for the identification. In the OLS specifications, I use district and year fixed effects for the entire country. Therefore, the coefficient on dams and upstream dams in the OLS represents a reduced-form relationship: the correlation between number of dams and the outcome variable. For the main outcomes, I also show the interaction of dams with homeland district, which identifies the additional effect of dams within homelands. This is more similar to the IV coefficient, which captures the effect of an additional dam within homeland districts over time.

1.4.2 Threats to Identification

The exclusion restriction requires that the instrument only affects outcomes via its effect on the likelihood of irrigation dam construction. This exclusion restriction criteria fails if the instrument is independently correlated with the outcome, after conditioning on covariates.[21] In this case, the covariates are critical. For example, districts may have received preferential access to electricity after the end of Apartheid, and this is correlated with land gradient (Dinkelman, 2011). I control for this using land gradient-policy interactions. Another possible threat to

[21]Note that the river gradient component of the instrument can and is itself correlated with at least one outcome variable, water quality. But, in those regressions I control for district fixed effects. The results are not causal if changes in the policy variable ($P_{dt} = H_d A_t$) differentially affect water quality by river gradient.

the identification is the Land Reform Act enacted to assist in land redistribution toward the poor. However, this is more likely to be targeted based on land gradients (which I control for) and only two percent of land had been redistributed as of 2001 (Twala, 2006; Strobl and Blanc, 2013).

Furthermore, the instrument compares differences in dam construction *within homeland areas* over time. Thus, even if policies across homeland and non-homeland areas were quite different, it will not affect estimates within homeland areas. This does assume that even if dam or water-related policies varied within former homeland areas, the design and enforcement of these policies was not systematically different based on the river gradient. In other words, absent dam construction, changes in infant mortality would not have varied *by river gradient* within homeland areas.

1.5 Results

1.5.1 Balance and First Stage

Table 1.2 presents summary statistics for the mothers in the sample. Column (1) includes all mothers in South Africa and column (2) restricts the sample to mothers within former homelands. Across all districts, mothers are poor, tend to live in rural areas, do not have piped water, and travel about 30 minutes to obtain water.[22] In the third column, I regress each mother characteristic on the steep river gradient variable, with geographic controls. This test provides supporting evidence that the river gradient instrument is uncorrelated with most mother characteristics. The exception is electricity.[23] Given that the instrument relies on variation over time, the IV assumption remains valid so long as outcomes do not vary differentially by river gradient as policy changes.

[22] A caveat to this table is that the responses are reported as of the time of survey, not at the time of the child's birth. This is a limitation of the data.

[23] Given fourteen variables were tested, finding at least one significant variable is expected with a ten percent threshold.

Table 1.2: *Mother Characteristics and Instrument Test*

Variable	All Districts Mean (1)	Homelands Mean (2)	F-stat (3)	N (all) (4)
African	0.809	0.987	0.011	31,809
	[0.393]	[0.111]	(0.196)	
Urban/urban informal	0.463	0.233	0.036	31,809
	[0.499]	[0.423]	(0.85)	
Years education	6.527	6.097	0.937	29,519
	[3.011]	[3.029]	(0.334)	
Doctor delivered child (0/1)	0.305	0.203	0.438	14,839
	[0.461]	[0.402]	(0.509)	
Birth in hospital	0.711	0.656	0.004	14,812
	[0.453]	[0.475]	(0.95)	
Age of mother at child's birth	25.89	25.968	0.111	30,153
	[6.604]	[6.866]	(0.739)	
Ever breastfed	0.85	0.876	0.351	14,733
	[0.357]	[0.330]	(0.554)	
Only breastfed (first year)	0.065	0.050	1.781	1,162
	[0.246]	[0.219]	(0.152)	
Death in months if child died	5.475	5.397	1.960	1,602
	[9.568]	[9.142]	(0.163)	
Electricity	0.826	0.781	4.582	31,809
	[0.379]	[0.413]	(0.033)**	
Average minutes to water	26.632	30.028	0.071	5,783
	[30.109]	[32.217]	(0.790)	
Moved into residence <2 yrs	0.129	0.084	0.278	13,336
	[0.336]	[0.277]	(0.599)	
Piped (house/yard)	0.446	0.211	0.188	31,809
	[0.497]	[0.408]	(0.665)	
Piped (house/yard/public tap)	0.739	0.618	0.777	31,809
	[0.439]	[0.486]	(0.379)	

Notes:
1. Each observation reflects a child's mother in the sample. Mothers may be present more than once if they had multiple children in the sample. The data was pooled across three rounds of survey data. For the variables with N between 13,000 - 15,000, only the 1998 and 2003 rounds were available. For the "only breastfed" variable, only those born in the last year and in the 1998 survey round were available.
2. The table reports the means for all districts in column (1) and the means for only former homeland districts in Column (2). Standard deviations reported below in brackets.
3. Column (3) regresses the mother variable on the steep river gradient variable, with geographic controls. I test whether the steep river gradient variable is significant and I present the F-test, with p-value in parentheses.
4. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
5. The data source is the SA Demographic and Health Surveys (1987, 1998, 2003).

The first stage is presented in Table 1.3 for each subset of years used in the analysis. The instrument is the fraction of steep river gradient pixels interacted with changes in Apartheid policies in Homeland districts. In all specifications, the F-statistic suggests that the instrument is strong.

Table 1.3: *The Effect of River Gradient on Dam Construction (First Stage)*

Dependent Variable:	Number of Dams in District as of Year		
Years:	1980-2002	1980-2010	1995-1999
	(1)	(2)	(3)
Fraction Steep Gradient*Policy	-16.87***	-15.50***	-15.61***
	(3.203)	(3.179)	(4.386)
Geographic Controls*Policy	Yes	Yes	Yes
River Gradient*Year FE	Yes	Yes	Yes
Province FE	Yes	Yes	Yes
F-test for instrument	27.74	23.78	12.67
P-value	0.00	0.00	0.00
Observations	8,142	10,974	1,770
Mean Dependent Variable	7.169	7.653	8.465

Notes:
1. Each observation reflects a district-year combination. The dependent variable is the number of irrigation dams constructed as of that district-year.
2. Each column reports the results of estimating equation (2) for the subset of years for which infant mortality, water quality, and household data are available. The term "Policy" refers to the interacted variable (H*A).
3. The controls including the following geographic variables, each interacted with the policy variable (H*A): total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m). The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and province fixed effects.
4. Robust standard errors clustered at the district level are in parentheses for the panel. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
5. The data source for the dam data is the Department of Water Affairs.

1.5.2 Infant Mortality

Table 1.4 shows the IV and OLS results for the effect of irrigation dams on infant mortality. The estimates capture the combined effect of irrigation dams on infant health, inclusive of all changes arising from markets, the environment and any omitted variables. Columns (5) and (6) present the IV result using estimating equation (3) and show large positive coefficients on irrigation dams in the district. The coefficient (.003) is approximately a six percent increase on the dependent mean infant mortality (0.048).[24] In addition to estimating the IV using two-stage least squares, I present the results using limited information maximum likelihood, given this estimation method is more robust to weak instruments (Stock and Yogo, 2005).

[24]These estimates are consistent with the literature. Brainerd and Menon (2014) estimate that a 10 percent increase in average fertilizer chemical concentrations during the month of conception raised infant mortality likelihoods by 4.6 percent. Galiani et al. (2005) estimate reductions in under-5 child death of 5-8 percent associated with privatization. Field and Glennester (2011), alternatively, found that when households switched from closer arsenic-contaminated wells to further away wells or surface water sources, infant and child mortality increased 27 percent.

Table 1.4: *The Effect of Dam Construction on Infant Mortality*

Dependent Variable:	Infant Death					
Type of Regression:	OLS	OLS	OLS	Logit (w/FE)	IV (2SLS)	IV (LIML)
	(1)	(2)	(3)	(4)	(5)	(6)
Dams in District	0.00218*	0.00158*	0.000805	0.0685***	0.00309**	0.00313**
	(0.00113)	(0.000947)	(0.00101)	(0.0210)	(0.00134)	(0.00137)
Dams Upstream	0.000288	0.000249	0.00132***	0.00705	-0.000813*	-0.000826*
	(0.000334)	(0.000291)	(0.000440)	(0.00682)	(0.000488)	(0.000496)
Dams in District*			0.0104***			
Homeland			(0.00288)			
Dams Upstream*			-0.00168**			
Homeland			(0.000665)			
District FE	Yes	Yes	Yes	Yes		
Year FE	Yes	Yes	Yes	Yes		
Province-Year Trends		Yes				
Geo. Controls*Policy					Yes	Yes
River Gradient*Year FE					Yes	Yes
Province FE					Yes	Yes
First Stage F-Statistic					51.94	51.94
Observations	31,809	31,809	31,809	29,366	31,809	31,809
Mean Dep. Variable	0.048	0.048	0.048	0.052	0.048	0.048

Notes:
1. Each observation reflects a single birth from 1980-2002. The outcome variable equals zero if the child is alive at one year, and one if the child has died within the first year.
2. Columns (1) through (4) report the regression of infant mortality on number of dams in the district and number of dams upstream, with controls as reported. Column (3) interacts the dam coefficients with a binary variable indicating former homeland. Column (4) reports the results of a logit regression using conditional fixed effects.
3. Columns (5) and (6) report the results using the instrument. For these specifications, the following control variables are included, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m), and fraction homeland category. For upstream dams, I control for the same variables: average upstream slope, average upstream elevation, total upstream river length, and total upstream district area. The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and province fixed effects.
4. The F-Statistic is the Cragg-Donald Wald F-Statistic.
5. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
6. The infant mortality data is from the South Africa Demographic and Health Surveys (1987, 1998, 2003).

The OLS estimates, which rely on fixed effects, provide corroborating evidence that dams increase infant mortality. These estimates control for district and year fixed effects in column (1), and district and year fixed effects with a province time trend in column (2). While the trends absorb some variation, the effect size remains positive and significant at the 10 percent level. In column (3), I re-run the OLS specification interacting the dam variable with an indicator for being a former homeland district. I find that the interaction is positive, suggesting that dams are correlated with relatively higher infant mortality increases within former homelands. This is unsurprising, given that households within homelands were more disadvantaged than the average household and mothers in former homelands were more likely to rely on river water sources.[25] This result is similar to the IV result, which looks at the effect size within former homelands. The OLS somewhat overestimates the effect of dams on infant mortality, which could suggest that dams are placed in areas that are experiencing lower growth or declining health.

Downstream, dams generate net reductions in infant mortality, at least in the IV. However, estimated coefficients are an order of magnitude smaller than those within the district and the estimates are not statistically different from zero in the OLS. While net impacts are small, the IV results suggest that districts downstream from dams experience net health benefits greater than costs. The small effects of upstream dams are as predicted, given the small size of dams.

In Table 1.5, I check that the increase in infant mortality is robust to the inclusion of mother controls. The controls are characteristics reported by the mother at the time of surveying, not at the time of the child's birth. Because some of the mother controls may be endogenous, the regression coefficients for the characteristics cannot be easily interpreted. However, the controls should be correlated with the characteristic at the time of child birth. Thus, this table provides some confidence that the estimated effects are robust to the inclusion of household characteristics. Unfortunately, the mother characteristics are not asked in every survey round; for that reason, columns (2), (3) and (5) are smaller samples.

[25] Because infant mortality is a binary variable, I also present conditional logit with district and year fixed effects in column (4). The average marginal effect from the conditional logit regression is much larger at the mean, given it is bounded (0/1). However, it is useful to check that it remains positive.

Table 1.5: *Robustness Checks using Additional Controls*

Dependent Variable: Infant Mortality							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Instrumental Variables							
Dams in District	0.00272**	0.00224	0.00239*	0.00303**	0.00207	0.00274**	0.00300**
	(0.00121)	(0.000985)	(0.00144)	(0.00133)	(0.00139)	(0.00131)	(0.00129)
Dams Upstream	-0.000772**	-0.000579	-0.000540	-0.000795	-0.000404	-0.000678	-0.000931*
	(0.000350)	(0.000776)	(0.000589)	(0.000488)	(0.000583)	(0.000485)	(0.000492)
African	0.0400***						
	(0.00553)						
Employed		-0.00583					
		(0.00519)					
Female HH			0.00243				
			(0.00426)				
Urban				-0.00378			
				(0.00466)			
Electricity					-0.00867		
					(0.00574)		
Piped						-0.0222***	
						(0.00399)	
Education							-0.00291***
							(0.000612)
Panel B: OLS							
Dams in District	0.00195*	0.00196	0.00201	0.00221*	0.00211	0.00219*	0.00218**
	(0.00108)	(0.00198)	(0.00198)	(0.00113)	(0.00201)	(0.00113)	(0.00108)
Dams Upstream	0.000208	-0.000401	-0.000450	0.000261	-0.000483	0.000174	0.000288
	(0.000316)	(0.00112)	(0.00111)	(0.000328)	(0.00110)	(0.000334)	(0.000335)
African	0.0354***						
	(0.00479)						
Employed		-0.00458					
		(0.00503)					
Female HH			0.000508				
			(0.00429)				
Urban				-9.81e-07			
				(0.00406)			
Electricity					-0.00719		
					(0.00491)		
Piped						-0.0191***	
						(0.00432)	
Education							-0.00312***
							(0.000585)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	31,726	13,082	13,087	31,809	13,030	31,809	29,227

Notes:
1. Each observation reflects a single birth from 1980-2002.
2. The IV regression includes the following control variables, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m), and fraction homeland category. For upstream dams, I control for the same variables: average upstream slope, average upstream elevation, total upstream river length, and total upstream district area. The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and province fixed effects.
3. The OLS regressions include for district and year fixed effects.
4. Employment, female household head and electricity questions were not available for the 1987 round, so the regression has a smaller sample.
5. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.

This version of the paper assumes no health spillovers to adjacent districts. This assumption is generally reasonable, given that districts are large and operate as roughly separate labor markets (as discussed in Dinkelman, 2011). In addition, pollution is unlikely to travel further than the district. However, it is possible that dams affect adjacent districts through one or several channels. For example, water for irrigation is transferred in some instances and could affect results. In this case, the result would provide a conservative estimate of the effect of dams on health.

1.5.3 Other Health Outcomes

Evidence of water-related illnesses in response to dams provide support for the health-water channel. To evaluate this, I rely on mother responses to other health-related questions in the DHS data. In each survey round (1987, 1998, 2003), mothers are asked about the health of their children in the last two weeks.[26] In Table 1.6, I present evidence that children in districts which receive more dams experience more fevers, coughs and diarrhea in the last two weeks. All three outcomes could indicate greater water-related illness. Fevers indicate any infection, including malaria, which is predicted to increase in response to stagnant water (EPA, 2014; WHO, 2014). Similarly, coughing can indicate respiratory illness. Finally, infants consuming polluted water or in households with low water access are more likely to experience diarrhea, though the effect is imprecise and not statistically distinguishable from zero.

[26] I restrict the sample to children under five years of age, given this is all that was surveyed in 2003 round. In the 1987 survey, the responses are further restricted to the last two children born into the household.

Table 1.6: *Other Health Outcomes*

Dependent Variable:	Fever in the last two weeks	Cough in the last two weeks	Diarrhea in last two weeks	Diarrhea in last two weeks (≤2 yr old)
	(1)	(2)	(3)	(4)
Panel A. Instrumental Variables				
Dams in District	0.0518***	0.0606***	0.0124*	0.0134
	(0.0131)	(0.0228)	(0.00691)	(0.0110)
Dams Upstream	0.00808	0.0230	-0.00437	-0.0289***
	(0.0151)	(0.0194)	(0.00989)	(0.00991)
Panel B. OLS				
Dams in District	0.00611**	0.00600**	-0.00175	-0.00423
	(0.00255)	(0.00287)	(0.00181)	(0.00351)
Dams Upstream	0.000176	0.000435	0.00117	0.00209
	(0.000837)	(0.00108)	(0.000958)	(0.00151)
Controls	Yes	Yes	Yes	Yes
Observations	32,315	29,315	29,311	9,619
Mean Dep. Variable	0.113	0.125	0.066	0.133

Notes:
1. Each observation is a living child age 5 or younger, surveyed in one of three rounds of DHS data (1987, 1998, and 2003). The analysis is restricted to "currently living" at the survey time, but the results look the same if deceased children are included.
2. The question in the survey asked whether the respondent had experienced one of the above conditions during the past two weeks.
3. The IV regressions include the following control variables, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m), and fraction homeland category. For upstream dams, I control for the same variables: average upstream slope, average upstream elevation, total upstream river length, and total upstream district area. The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and district fixed effects.
4. The OLS regressions include district and year fixed effects.
5. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
6. Infant mortality data is from the South Africa Demographic and Health Surveys (1987, 1998, 2003).

1.6 Channels

I next evaluate several possible channels through which dams may affect infant mortality, including changes in water quality, water sources, employment and migration.

1.6.1 Water Quality

I first explore water quality responses to dams. Table 1.7 presents the results of the IV regression in Panel A and the OLS regression in Panel B. In Panel A, each additional dam in a district generates significant increases in indicators most often associated with agricultural water pollution. In the IV, nitrates in particular increase by around forty-five percent, consistent with high nitrate-based fertilizer concentrations in the water.[27] Given a median increase of five dams per district, some district-year averages increase over two hundred percent. Still, district-year averages do not generally exceed drinking water standards.

[27] Interestingly, phosphate, a usual component of fertilizers, is more heavily regulated in South Africa and does not appear to jump significantly in the district-year averages.

Table 1.7: *District-Level Water Quality Changes*

Dependent Variable:	Sulfate	Nitrates	Potassium	TDS	Chloride	Phosphate
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Instrumental Variables						
Dams in District	10.0***	0.322**	0.338***	35.6***	6.23*	0.00776
	(3.50)	(0.134)	(0.110)	(9.53)	(3.35)	(0.0119)
Dams Upstream	2.26	0.0609	0.0493	-8.60	-0.799	0.0108
	(1.63)	(0.0623)	(0.0464)	(5.57)	(1.512)	(0.00655)
Panel B. OLS						
Dams in District	0.922**	0.0331*	0.0406***	4.26***	1.92**	0.000261
	(0.422)	(0.0200)	(0.0141)	(1.55)	(0.758)	(0.00314)
Dams Upstream	0.887*	0.00926	0.0172*	1.80**	0.131	0.000548
	(0.475)	(0.0102)	(0.00930)	(0.836)	(0.371)	(0.00147)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,192	5,138	5,693	5,509	5,685	5,198
Mean Dependent Variable	86.0	0.707	5.34	575	173	0.184
Joint test on outcomes (OLS):	chi2 =	8.51				
	p-value	0.00				

Notes:
1. The table presents district-year level observations, with geographic controls and district FE. The IV and OLS regressions are weighted by the 1996 district population estimates. Unweighted results are similar in size and significance.
2. Panels A and B regress the district-year mean values of each chemical indicator on number of dams and dams upstream. The mean relies on only stations in place as of 1990. Only districts that have health data are included, but results are robust to using all districts.
3. The results are weighted by district populations in 1996.
4. The IV regressions include the following control variables, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m), and fraction homeland category. For upstream dams, I control for the same variables: average upstream slope, average upstream elevation, total upstream river length, and total upstream district area. The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and district fixed effects.
5. The OLS regressions include district and year fixed effects.
6. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.

In Table 1.8, I test an additional specification to evaluate whether dam construction increases the likelihood of obtaining a water quality reading over the standard.[28] Specifically, I construct threshold indicators for each chemical. The indicators equal one when a water quality reading is over the standard. I then match each water quality monitoring station to its nearest dam, and I create an indicator for whether the monitoring station is upstream or downstream from the dam.[29] Using this, I implement a difference-in-differences regression estimating the effect of dams upstream, after the dam is constructed, on the likelihood that a downstream water quality reading exceeds the threshold. The regression is as follows:

$$Threshold_{smt} = \gamma_o + \gamma_1 L_s C_t + \lambda_s + \delta_m + year_t + \varepsilon_{st} \tag{1.4}$$

where $Threshold_{smt}$ is an indicator variable equal to one if a given water quality observation is above the threshold and zero if the indicator is below the threshold; L_s is an indicator equal to one if the water quality station is located downstream from the dam and zero if the station is upstream; and C_t is an indicator that equals one starting in the year that the closest dam becomes operational and is zero prior to the dam becoming operational. The main effect of L_s is absorbed by the station fixed effect, and the main effect of C_t is absorbed by the year fixed effect. The coefficient of interest is on the interaction term, γ_1. This measures the effect of being located downstream from a dam, after the dam has become operational. As Table 1.8 demonstrates, stations located downstream from dams are more likely to generate readings about the standard for nitrates and total dissolved salts (TDS).

[28]Estimating the regression at the station-month-year level allows for greater precision. It also allows evaluation of how often individual monthly readings go above the threshold, rather than district-year averages (which are much less likely to exceed the threshold). From a health perspective, acute nitrate exposure can lead to infant mortality.

[29]The average distance between a dam and its nearest station is 15 km. Half the readings are within 6 km, while the farthest "closest dam" reading is 685 km away.

Table 1.8: *Probability of Exceeding WQ Standards Downstream from Dams*

Water Quality Indicator:	Sulfates	Nitrates	Potassium	TDS	Chloride	Phosphate
Downstream Dam*	-0.000183	0.00661**	0.00231	0.0180**	0.00678	0.00267
Dam Constructed	(0.00121)	(0.00328)	(0.00186)	(0.0080)	(0.00673)	(0.00571)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	195,178	159,959	168,544	145,151	165,746	176,143
Mean Dep. Variable	.00800	.00308	.0206	.0955	.0199	.0444
Standard (mg/l)	10	100	200	1000	12000	1.0

Notes:
1. Each observation reflects a station-month-year reading for the given chemical indicator. The dependent variable is a binary variable indicating whether or not the reading is over the drinking water standard, for the given chemical indicator. The data includes all water quality monitoring stations that are on rivers with dams. Each monitoring station is matched to the nearest dam.
2. Each column reports the results of estimating equation (4). The regression tests whether being located downstream from a dam, after the dam is constructed, leads to an increase in the likelihood of obtaining a water quality indicator reading over the standard. The coefficient of interest is therefore an interaction between location (downstream=1) and time (dam constructed=1). The regression includes station, month, and year fixed effects.
3. Robust standard errors clustered at the station level. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.

Note that this analysis abstracts from some of the spatial and seasonal distribution of water quality within districts. Dam-enabled agriculture leads to elevated nitrate levels upstream and downstream, and nitrate levels are higher during dry summer months.[30]

1.6.2 Water Sources

An additional concern is that households alter their choice of primary domestic water source in response to dam-induced changes in water availability. This could happen if dams reduce river or groundwater levels. Similarly, income effects from dams or reduced food prices could enable households to switch to more expensive but cleaner sources. A third alternative is that dams enable construction of piped water. However, less than two percent of dams report being used for domestic water supply (See Appendix Table A.1).

Households should be less likely to experience infant mortality if they move toward piped sources. In Table 1.9, I present imprecise but unsurprising estimates of shifts in primary water source. Forty percent of households use piped water in their homes. As dams are constructed, households are more likely to use piped water or water directly from dams, and they are less likely to rely on rivers or public taps. This could be explained by any of the above described mechanisms. Regardless of mechanism, this pattern is unlikely to generate the observed increase in infant mortality.

[30] Rainfall has ambiguous effects on water pollution. On the one hand, greater rainfall can increase runoff. On the other hand, it can dilute pollution. In most parts of South Africa, the dilution effect dominates.

Table 1.9: *Household Outcomes Related to Changes in Water Source*

	Main Source of Domestic Water				Other Water-Related Questions		
Dependent Variable:	Piped in household	Public Tap	River/Spring	Dam	Dist. > 200m to Water*	Pay for Any Water*	Enough Water**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Instrumental Variables							
Dams in District	0.0374	-0.0724	-0.0820**	0.0612	0.257*	0.0854**	-0.00319
	(0.0478)	(0.0707)	(0.00341)	(0.000909)	(0.147)	(0.0424)	(0.00229)
Dams Upstream	-0.0465	0.154	0.0503	-0.104	-0.0703	-0.0466	-.000822
	(0.0763)	(0.121)	(0.0726)	(0.0817)	(0.0809)	(0.0478)	(0.00105)
Panel B. OLS							
Dams in District	0.00543	0.00210	-0.00535	-0.000332	0.0144	0.00262	-0.0361
	(0.00399)	(0.00403)	(0.00341)	(0.000909)	(0.0375)	(0.00504)	(0.0261)
Dams Upstream	0.00595*	-0.000574	-0.000104	-0.000996	-0.0209	-0.00142	0.0115
	(0.00328)	(0.00192)	(0.00160)	(0.000957)	(0.0156)	(0.00388)	(0.0179)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	119,116	119,116	119,116	119,116	13,363	118,953	45,477
Mean Dep. Variable	0.4072	0.1655	0.0870	0.0159	0.506	0.5064	0.866

Notes:
1. Each observation is a household surveyed in the October Household Survey, pooled across the years 1995 to 1999.
2. The IV regressions include the following control variables, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m), and fraction homeland category. For upstream dams, I control for the same variables: average upstream slope, average upstream elevation, total upstream river length, and total upstream district area. The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and district fixed effects.
3. The OLS regressions include district and year fixed effects.
4. Each observation is weighted using the provided household level weights.
5. For the distance to water question, only those that rely on natural water sources are included. For the enough water question, only two years of data (1995-1996) are available.
6. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
7. The data is the October Household Survey (1995-1999).

Households also answer other water-related questions, shown in the last three columns of Table 1.9. Households report less access to adequate water, though again the results are imprecise. Among households using river sources, households report walking further, in line with closer sources becoming unavailable or becoming observably polluted. In addition, households are more likely to pay for water, suggesting they are better off or have new access to more costly piped water.

1.6.3 Labor Market Effects

As described above, dams are predicted to increase employment. In Table 1.10, I evaluate the effect of dams on adult responses to labor-related questions. New dams increase the likelihood of current employment, employment in the last year and total hours worked. They also decrease the likelihood of reporting an income below the poverty line.[31]

[31] Without data on the entire population and its composition, I cannot distinguish whether the employment result is driven by increased demand or by falling supply. Fortunately, this isn't needed to show that irrigation dams increase employment, which is unlikely to directly increase infant mortality.

Table 1.10: *Employment Outcomes in Response to Dams*

Dependent Variable:	Currently Working (1)	Working in the Last Year (2)	Total Hours in Last 7 Days (3)	Salary Below Poverty (4)
Panel A. Instrumental Variables				
Dams in District	0.0425**	0.0296	2.475*	-0.0604***
	(0.0191)	(0.0269)	(1.305)	(0.0222)
Dams Upstream	0.0160	0.0469	-0.179	0.00581
	(0.0245)	(0.0331)	(0.941)	(0.0343)
Panel B. OLS				
Dams in District	0.00962***	0.0102**	0.521***	-0.00891***
	(0.00351)	(0.00411)	(0.139)	(0.00265)
Dams Upstream	0.00507***	0.00601***	0.210***	-0.00472***
	(0.00163)	(0.00130)	(0.0756)	(0.00173)
Controls	Yes	Yes	Yes	Yes
Observations	348,618	348,618	346,969	348,618
Mean Dep. Variable	0.337	0.356	15.0	0.814

Notes:
1. Each observation is an individual 15 years or older, within the households surveyed for the OHS survey. It includes workers, the involuntary unemployed, and those who are not economically active.
2. The IV regressions include the following control variables, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6%), district elevation category (0-500 m, 500-1000 m, >1000 m), and fraction homeland category. For upstream dams, I control for the same variables: average upstream slope, average upstream elevation, total upstream river length, and total upstream district area. The regression also controls for the main effect of the policy variable (H*A), steep gradient*year fixed effects, and district fixed effects.
3. For OLS, the controls include district and year fixed effects.
4. Currently employed includes informal and formal work within the last seven days, including if the individual is employed but did not attend work over the last seven days (e.g., due to illness). Work in the last year extends this to the last year. Salary below poverty line indicates income below six hundred rand per month.
5. Each observation is weighted using the person weights provided.
6. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
7. The data source is the October Household Survey (1995-1999).

At face value, an increase in employment should unambiguously increase household income and reduce the likelihood of infant mortality. However, employment effects could indicate that parents were spending less time with their children and neglecting infant care. To investigate this, I separately evaluate the regression by gender and age in Appendix Tables A.5 and A.6. First, while employment gains are slightly higher for women in the IV, the overall employment of women remains low. Furthermore, the segment of women who are most likely to care for children are not driving the employment gains. Specifically, women between 30 and 50 are twice as likely to be employed and have double the increase in hours worked, as compared to women under 30. Likewise, the effects are near zero for women over fifty.[32] An additional check on the results is shown in Table 1.5, where I regress infant mortality on dams while controlling for any household employment. The infant mortality result remains positive though smaller and statistically indistinguishable from zero, likely a consequence of the smaller sample.

1.6.4 Population and Migration Effects

A final concern to the interpretation of the results is that a changing composition of district households is causing the measured infant mortality increase. To evaluate this, I first look at overall population changes as a function of the number of dams received. In Appendix Table A.4, I show that each additional dam constructed in a district generates negative population growth. Between 1996 and 2001, the average district population grew by 17 percent, while districts that received dams grew by 6 percent less for each additional dam. In other words, despite national growth, households shift away from areas with greater irrigated agriculture.[33]

Given greater population declines in districts with more dams, I perform further analysis of district in-migrants and out-migrants. In the 2001 Census, individuals report whether they moved into their current district within the past five years. If the answer is yes, individuals report the name of their prior residence. Using this information, I construct a rough measure

[32]Dinkelman (2011) showed that women who are younger than 30 years or older than 50 years are far more likely to be primary caregivers.

[33]This is consistent with evidence of rural out-migration after the end of Apartheid , but in contrast to the effects of rural electrification documented by Dinkelman (2011). This is not surprising: rural electrification benefits households, while irrigation dams may not.

of the number of district in-migrants and out-migrants, which I report based on the type of district in Table 1.11. The fractions cannot be directly compared, given that the out-migrant fraction is based on the 2001 population less new migrants, and the in-migrant fraction is based on the 2001 population. However, total flows out are greater than total flows in. Furthermore, the table provides a few facts. First, former homelands experience lower migration rates in both directions than non-homelands. In particular, few individuals move into former homeland districts, consistent with the fact that homelands are less desirable and land ownership within homelands was generally determined by tribal leaders. However, the rates of in-migration and out-migration are not very different between districts that do not receive any dams between 1996-2001 and districts that do receive any dams.

Table 1.11: *Average District Migration Statistics (1996 to 2001)*

	No Irrigation Dams Mean (1)	Any Irrigation Dams Mean (2)	Difference (p-value) (3)
Panel A. Former Homeland Districts			
Fraction that outmigrate	0.090	0.088	0.00185
	[.003750]	[.00631]	(0.825)
Fraction of new migrants	0.053	0.0638	-0.010768
	[0.0030]	[.0069]	(0.1326)
Number of observations	81	19	
Panel B. All Districts			
Fraction that outmigrate	0.119	0.106	0.0124
	[0.0048]	[0.0040]	(0.153)
Fraction of new migrants	0.108	0.113	-0.0054
	[0.004]	[.007]	(0.505)
Number of observations	266	88	

Notes:
1. Each observation is a district. Column (1) reports the mean fraction that out-migrate from districts that received no dams from 1996-2001 and the mean fraction of new migrants into districts that received no dams. Column (2) calculates the same fractions, but for districts that do receive dams. Standard deviations are reported in brackets.
2. For both in-migration and out-migration, fractions are based on the movement of the household head.
3. Panel A restricts the estimates to former homeland districts, while Panel B includes all districts.
4. The data source is the 2001 Census.

Migration patterns may affect the interpretation of the results if net migration masks out-growth by some "types" and in-growth by other "types." To address this, in Table 1.12, I examine the characteristics of individuals that leave districts which receive dams and the composition of individuals that enter districts which receive dams. Panel A restricts the analysis to former homeland areas, while Panel B includes all districts. I show that the composition of those individuals migrating into and out of districts within homeland areas are better off along characteristics that are correlated with infant mortality: less poor, higher educated, and more likely to have piped water access. This suggests that in-migration is not making the average individual more vulnerable to infant mortality.

Table 1.12: *Profile of In-migrants and Out-migrants*

	Non-migrants Mean	In-migrants Mean	Diff. to (1)	Out-migrants Mean	Diff. to (1)
Panel A: Former homeland districts that receive dams from 1996-2001					
	(1)	(2)	(3)	(4)	(5)
African race	0.9832	0.9122	0.0710	0.9569	-0.0263
	[0.1286]	[0.2831]	(0.000)***	[0.2032]	(0.000)***
Employed	0.2015	0.19467	0.0068	0.2498	0.0482715
	[0.4011]	[0.3960]	(0.2929)	[.4329]	(0.000)***
HS education	0.0465	0.1076	-0.0611	0.1051	0.0585771
	[0.2105]	[0.3099]	(0.000)***	[0.3066]	(0.000)***
Piped water	0.2834	0.4732	-0.1898	0.5964	0.3130
	[0.4507]	[0.4993]	(0.000)***	[0.49066]	(0.000)***
Below poverty line	0.9870	0.9573	0.0297	0.9662	-0.020756
	[0.1133]	[0.2022]	(0.000)***	[0.1806]	(0.000)***
Observations	60,185	4,053		6,930	
Panel B: All districts that receive dams from 1996-2001					
	Non-migrants Mean	In-migrants Mean	Diff. to (1)	Out-migrants Mean	Diff. to (1)
	(1)	(2)	(3)	(4)	(5)
African race	0.7317	0.6317	0.1000	0.6297	-0.1020
	[.4431]	[.4823]	(0.000)***	[0.4829]	(0.000)***
Employed	0.1749	0.1637	0.0112	0.1661	-0.0088
	[.3799]	[.3700]	(0.000)***	[0.3722]	(0.000)***
HS education	0.0863	0.1908	0.1001	0.1716	0.0853
	[.2808]	[.3930]	(0.000)***	[0.3770]	(0.000)***
Piped water	0.607	0.7572	-0.1502	0.7522	0.1452
	[0.4884]	[0.4288]	(0.000)***	[0.4318]	(0.000)***
Below poverty line	0.9222	0.8421	0.0801	0.8566	-0.0656
	[0.2678]	[0.3646]	(0.000)***	[.3505]	(0.000)***
Observations	203,029	30,868		34,213	

Notes:
1. The table presents characteristics for in-migrants and out-migrants to districts that receive any dams between 1996 and 2001. The data source is the 2001 Census.
2. Columns (1), (2) and (4) present the means, with standard deviations in brackets. Column (3) reports the difference between the mean for non-migrants and the mean for in-migrants. Column (5) reports the difference between the mean for non-migrants and the mean for out-migrants. P-values are in parentheses.
3. Panel A restricts the sample to only those moving into or out of former homeland districts. Panel B includes all districts.

However, out-migration may bias estimates, given that the individuals who leave districts also have characteristics less correlated with infant mortality. I perform two tasks to address this. First, in Table A.4, I show that the average education level in districts which receive dams falls only two percent more than in the average district and that the fraction with piped water goes up slightly, though imprecisely. In addition, average district employment increases, as also shown in the household survey data. Thus, districts that receive dams are not observably worse off. Second, while data limitations prevent measuring infant mortality outcomes for only non-migrants, I can identify those mothers who have lived in their residence for more than two years. In Table A.3, I show that the infant mortality result for these mothers remains positive.

1.7 Conclusion

Irrigated agriculture has multiplied the global food supply, but agricultural production benefits may be coupled with health consequences. In particular, while it is well understood that intensive agriculture uses significant water and returns polluted (nutrient-enriched) water to waterways, the health effects of these water-related changes on nearby households are not well documented (Vitousek et al., 2009).

In this context, irrigation dams provide a useful indicator of increased agricultural activity, particularly greater water use. This paper measures the effect of irrigation dams on infant mortality within former homeland districts of South Africa, providing some of the first evidence that irrigation dams generate adverse infant health outcomes. Specifically, I demonstrate that each additional irrigation dam led to an increase in infant mortality of 6 percent, off of a baseline of 4.8 percent. I then investigate changes in water quality, one plausible channel through which the health effects arise. I find that each additional dam within a district increased concentrations of several chemical indicators associated with agricultural pollution. For example, dams increase nitrate concentrations by forty-five percent, on average, and also increase the likelihood of obtaining a nitrate reading over the drinking water standard by 80 percent, though these outcomes are still rare.

Because the analysis measures the combined effect of dams on infant mortality, I also evaluate changes in domestic water sources, employment and migration. I do not find clear evidence that households switch water sources. Moreover, I find increased employment, which should decrease infant mortality. While I find small declines in district populations, I don't find evidence that migration altered district compositions in a way that would increase infant mortality.

The current data does not allow for separate identification of different water-related mechanisms through which irrigation dams affect infant health. For example, anecdotal reports suggest that both increased water pollution and reduced water availability occur, but these two channels cannot be separately identified with this data. Rather, the water quality result could indicate increased fertilizer runoff or, alternatively, the results could reflect more concentrated pollutants as a result of less total water. Slow flowing and stagnant water may also cause malaria or other vector borne diseases, explaining some of the health effect. Future versions of this paper will incorporate new data on water flow in order to better clarify these underlying health mechanisms. Despite this limitation, econometric approaches like the one utilized in this paper are valuable to policy because they incorporate behavioral responses to pollution, a useful complement to dose-response estimates of human responses to direct exposure.

While using the construction of irrigation dams as an identification strategy is specific to this context, the findings suggest that coordination among agriculture, water and health policy-makers may be important in managing water supplies and in helping households to identify safe drinking water sources. Compared to most of Sub-Saharan Africa, South Africa has practiced more intensive agricultural techniques for decades, a result of the historical political and land ownership structure. Today, significant agricultural water use in South Africa contributes to its' current water stress (WRG, 2012). Similar irrigation techniques are promoted worldwide, especially in developing countries where food and water demand are highest. This paper demonstrates the need for countries with expanding agricultural production to consider the water-related consequences. Complementary infrastructure or implementation of more sophisticated water management strategies may be of greater value as agriculture intensifies. Without effective water management in place, water resources may become polluted and depleted, a classic "tragedy of the commons." Furthermore, those without the resources to find alternatives to polluted and reduced water streams will be the hardest hit.

Chapter 2

Technology Adoption Under Uncertainty[1]

2.1 Introduction

Many technology adoption decisions consist of at least two parts, which occur at different points in time: a take-up decision and a use or follow-through decision. While subsidies are often used to increase take-up, critics of subsidies for technology adoption, in development, health and environmental policy, worry that subsidizing the initial take-up decision lowers subsequent follow-through. For example, are subsidized health treatments as likely to be taken according to directions? Will subsidies for home energy audits lower weatherization rates among those receiving audits? Are agricultural technologies used as often if their take-up is subsidized?

Numerous studies in health and development have examined the reasons why follow-through might be affected by the initial cost of the technology, including screening or selection effects, learning, and psychological channels such as sunk cost effects or procrastination (Ashraf et al., 2010; Cohen and Dupas, 2010; Berry et al., 2012; Ashraf et al., 2013, Tarozzi et al., 2013, Dupas, 2014, Fischer et al., 2014). With the exception of learning, little attention has been paid

[1]Co-authored paper with Samuel Bell, Kelsey Jack, Paulina Oliva, and Chris Severen

to the role of uncertainty in the initial take-up decision. Specifically, at the time of take-up, many of the benefits and costs associated with the follow-through decision may be unknown. New information arrives in the form of learning about the technology (Oster and Thornton, 2012; Dupas, 2014; Carter et al., 2014; Fischer et al., 2014) or in the form of transient shocks to the opportunity cost of follow-through. If the new information is bad news about the profitability of the technology, then adopters may prefer to abandon the technology as long as it is cheap to do so.

We study technology adoption under uncertainty in three steps. First, we develop a simplified theoretical model that shows that uncertainty about the profitability of the technology lowers follow-through and undermines the degree to which a higher initial price of the technology screens out individuals who are unlikely to follow-through. Second, we provide reduced form evidence that uncertainty explains low follow-through in a field experiment on technology adoption in Zambia. Third, our experimental data identifies a structural model, which (a) shows that most of our reduced form results can be replicated by a data generating process that is fully consistent with uncertainty, and (b) allows us to perform counterfactual simulations of take-up and follow-through outcomes at different levels of uncertainty.

To develop intuition, we begin with a stylized model of intertemporal adoption under uncertainty in the presence of subsidies, where adoption consists of a take-up and a follow-through decision, both of which are binary.[2] The net costs of adoption include any direct costs or opportunity costs, net of benefits, associated with following-through with the technology.[3] The model allows for some share of these net costs of adoption to be observed at the time the take-up decision is made, and the remaining share to be observed at the time of the follow-through decision. The net costs that are unobserved at the time of take-up represent the uncertainty that the adopter faces. The theoretical model generates clear predictions about

[2]Abstracting from the intensive margin of both decisions simplifies the model and intuition. Much of the existing research focuses on adoption decisions where follow-through, or usage, has an intensive margin. This is true for durable goods such as bed nets, light bulbs, or stoves (Ashraf et al., 2010; Cohen and Dupas, 2010; Berry et al., 2012; Dupas, 2014). In our empirical application, we introduce an intensive margin to the follow-through decision (how many trees to cultivate) to to explore the effect of uncertainty on both the extensive and intensive margins

[3]The only cost associated with take-up is assumed to be the price of the technology, which may be subsidized.

the relationship between uncertainty and adoption outcomes. First, as uncertainty increases, take-up becomes more attractive as long as the follow-through decision can be postponed until after additional information is revealed. The expected profit of take-up therefore includes an option value associated with the decision to abandon the technology if the new information makes follow-through unprofitable. Second, as the share of information that is unobserved at take-up increases, the decision to take-up becomes less predictive of follow-through and high take-up rates may be accompanied by low follow-through. Third, uncertainty undermines the screening effect of the take-up price: if adopters know little about their net cost of follow-through when they take-up, then a higher take-up price will not be effective at screening out those who realize a high cost at follow-through.

Next, we generate empirical evidence for the model using a multi-period field experiment in rural Zambia. Farmers decide whether to take-up and care for a tree species that generates soil fertility benefits over the long term but carries short-run costs. In addition to the private soil fertility benefits, adoption of the trees generates sizable public benefits in the form of carbon sequestration and reduction of soil erosion.[4] The take-up decision consists of the purchase of a 50-tree seedling package at the start of the agricultural cycle. The follow-through decision consists of the number of seedlings that the farmer chooses to plant and care for (which we together refer to as tree cultivation), and occurs over the course of the subsequent year. We measure follow-through as tree survival at the end of the one-year period, and assume that farmers can guarantee tree survival for some level of costly effort.[5] Farmers who take-up the technology are subject to numerous shocks to their opportunity cost of follow-through, such as family illness, pests, weather and price shocks. Thus, the adoption decision we study is one characterized by uncertainty in some share of the net costs of follow-through at the time of take-up. Note that farmers can abandon the technology without penalty in states of the

[4]The present value of sequestration associated with half a hectare of surviving trees is around USD 353. As we discuss in Section 2.3, the back-of-the-envelope cost benefit analysis is negative from the individual perspective but positive from a social standpoint, offering a potential justification for subsidies.

[5]The choice of minimum effort that guarantees survival is optimal under convexity of the survival risk function as a function of effort. Therefore, the only source of uncertainty in tree survival according to our baseline model is the endogenous choice of follow-through by farmers, which depends on uncertain costs of effort. This assumption is examined in greater detail in Appendix B.3.

world where scarce resources, such as labor, are more profitably allocated to other productive activities. Just as with other technologies, the costless exit in the face of uncertainty generates an option value tied to the take-up decision.[6]

We introduce exogenous variation into this adoption decision at two different points in time. First, we randomize the cost of take-up through a subsidy on the purchase price of the 50-seedling package. This creates variation in take-up rates that helps characterize the distribution of expected net costs from the program across farmers. Second, we generate exogenous variation in the benefits associated with follow-through using a randomized threshold performance reward: a payment that is conditioned on the survival of 35 trees one year after take-up. The tree cultivation choices farmers make in response to the reward help us characterize the distribution of follow-through costs after potential shocks have been realized. Under the assumption that shocks are independent across farmers, the difference in the variance of net costs between the two points in time can be attributed to uncertainty.[7] Note that, rather than artificially varying the allocation of shocks across our study population, the reward creates exogenous variation in the variance of possible outcomes faced by farmers in our sample, and therefore in the distribution of uncertainty. By offering a positive payoff, the performance reward mitigates some of the downside risk much in the way that varying the terms of an insurance contract has a state-contingent effect on the distribution of outcomes.[8] Together, the different sources of variation at two points in time identify a structural model of intertemporal decisions that can distinguish between static and dynamic explanations for the outcomes we observe.

Under fewer assumptions than are needed for the structural model, the research design produces a number of stylized facts and reduced form results that are consistent with the

[6]The option value is akin to the "quasi-option" value described by Arrow and Fischer (1974): the value of the information revealed by delaying investment, or – in our case – the delayed decision to abandon the project. In this sense, the contract in our study is a put option (Pindyck, 1993). Like Stange (2012), we calculate the option value as the difference between a binding participation decision and one with free exit.

[7]The cross-farmer independence assumption rules out common shocks. In a model variant, discussed in Section 2.5, we also relax the independence assumption by allowing for an unexpected common shock to all farmers.

[8]A number of other authors study how affecting uncertainty via an insurance contract affects technology adoption and other outcomes (e.g., Bryan, 2014; Einav, 2013; Karlan, 2014).

predictions of our theoretical model in the presence of uncertainty. First, farmer choices are sensitive to economic incentives: they are more likely to take-up under higher take-up subsidies and higher rewards, and they are more likely to cultivate 35 or more trees under higher rewards. This simple observation rules out that take-up and follow-through decisions are divorced from economic incentives. Second, our conceptual model predicts that in the absence of uncertainty, farmers would only pay a positive price for the technology if they intended to follow-through with a strictly positive number of trees. Contrary to this, we find that many farmers (around 35 percent) who paid a positive price for take-up have zero trees one year later.[9] Third, our conceptual model predicts that under high levels of uncertainty, the relationship between take-up subsidies and follow-through is weak: individuals have little incentive to self-select based on the likelihood of follow-through, and a strong incentive to take-up if uncertainty is high.[10] Consistent with this, we find no significant reduced form relationship between the take-up price and follow-through.[11]

Although these facts are suggestive of uncertainty in farmers' decisions, we turn to our structural model to answer a number of remaining questions. First, are there sources of static heterogeneity (i.e. without uncertainty) that can explain the absence of a screening effect of the take-up subsidy? Second, how large is the uncertainty implied by farmers' decisions compared to the cross-sectional variation in expected net costs of adoption? And third, how much would uncertainty have to be reduced to generate meaningful correlations between take-up subsidies and follow-through?

First, we note that, in theory, static heterogeneity in private benefits can also explain an

[9]We rule out a number of alternative motives for taking-up and cultivating zero trees, such as side-selling and experimenter demand effects, using a combination of the study design and our empirical results. These are further described in Section 2.4.

[10]Liquidity constraints could also affect self-selection and follow-through. However, strong liquidity constraints would result in low take-up, and the take-up rates we observe are high across all levels of the take-up subsidy. Furthermore, farmers receive a show-up fee sufficient to cover take-up costs. Finally, we include an additional test for self-selection that does not depend on liquidity: variation in the timing of the reward announcement. We discuss these tests for liquidity constraints and other confounds, such as sunk cost effects, in Section 2.4.

[11]These results are in contrast with Jack (2013), who provides evidence that farmers self-select based on future costs into a tree planting incentive contract in Malawi. The differences may be driven by the context or the nature of the contract. Beaman et al. (2014) also find evidence of self-selection based on future returns. In their study, farmers with relatively high returns to credit select into an agricultural loan.

absence of or even negative screening effects of prices whenever the follow-through decision has an intensive margin. Heterogeneity in private profits from adoption potentially has multiple dimensions to the heterogeneity, which can be positively or negatively correlated, or uncorrelated. More specifically, there may be heterogeneity in both the level of the profit (for example, if there are fixed costs to adoption) and in the optimal usage rate (i.e. in the interior solution to the maximization problem); and these two types of heterogeneity may or may not be correlated. Only a positive correlation between the level of private profit and privately optimal usage rates would generate higher follow-through rates among those who participate at higher cost (i.e. a positive screening effect of prices). But zero or negative correlation between profit and optimal usage rates could result in a null or negative screening effect of prices. [12] To investigate the role of static heterogeneity in delivering similarly low follow-through outcomes across all levels of subsidy, our structural model allows for heterogeneity in the privately optimal number of trees as well as in the net cost of follow-through, and allows these two random (but known to the farmer) components of private profit to be freely correlated. The distribution of privately optimal number of trees is separately identified from the variance of the shocks to net costs, hence our model can distinguish between static and dynamic explanations for the lack of screening effects of prices. Our structural estimation delivers both a positive variance for the shocks to net cost, which offers a dynamic explanation for low tree survival, as well as a positive correlation between the component of net costs that is always observable to the farmer and the privately optimal number of trees (or a negative correlation between the level of private profit and the privately optimal number of trees). Both of these sources of heterogeneity reduce the ability of take-up prices to screen adopters that are likely to have high rates of follow-through.

Second, despite finding a substantial level of static heterogeneity (the standard deviation in the known component of net costs is 1.5 as large as the standard deviation of shocks), the

[12] Note that in many cases the source of heterogeneity in individual adoption choices is left unspecified. However, some source of heterogeneity, either static or transient, is necessary for take-up and follow-through decisions to differ across experiment participants within each branch of treatment. Most of the literature implicitly assumes that there is a positive correlation between the level of the profits and the privately optimal usage rate, which is necessary for "positive selection" i.e. a screening effect from the price (e.g., Ashraf et al., 2010; Cohen and Dupas, 2010). A notable exception is Suri (2011), who offers evidence of a negative correlation between optimal rates of usage and fixed costs of adoption in the case of fertilizer.

the scale of the shocks to profit appear to be economically meaningful and have a substantial effect on outcomes. Specifically, among farmers whose expected number of trees is above the reward threshold at the time of take-up, around 15 percent do not reach the threshold because their ex post optimal number of trees is lower than their ex ante expected number, once new information is accounted for. This is in addition to the 36 percent of farmers who choose to take-up cultivate zero trees even when the take-up cost is high, as shown in the reduced form results. Though we cannot precisely distinguish between learning and shocks in our empirical application, we describe supporting evidence in Section 2.7 that – in our case – shocks to the net cost of follow-through appear to be more important than learning.[13]

Third, to shed light on the amount of uncertainty that can eliminate screening effects, we implement counterfactual simulations that vary the magnitude of the uncertainty at the time of take-up, modeled as an increase in the standard deviation of the unknown cost component. Consistent with our theoretical model, uncertainty increases take-up and lowers follow-through conditional on take-up. Only at low levels of uncertainty do we see that higher prices for take-up have a positive effect on follow-through, by screening out farmers with a low probability of follow-through. If uncertainty is high enough, then subsidies will have little effect on either selection or take-up since the option value is high enough to generate high take-up rates at all prices.

Methodologically, our econometric framework is an example of sequential identification of subjective and objective opportunity cost components in a dynamic discrete choice model (Heckman and Navarro 2007, 2005). As described in Heckman and Navarro (2007), we can account for selection into treatment (in our case take-up of the tree planting program as well as non-corner solution tree survival outcomes) when identifying the distribution of the unobserved opportunity cost determinants. We do so by introducing two layers of random variation in economic incentives, one of which produces a probability of take-up equal to one for a sub-population and a second of which produces interior solution in tree cultivation outcomes with probability one in the limit. The use of experimental variation in treatments at

[13]Time inconsistent preferences (see e.g., Mahajan, 2011) are also consistent with our models, as discussed in Section 2.7.

two different points in time offers an alternative to a panel data structure (used for example, in (EIA), since statistically independent samples are exposed to each of the different treatment combinations. To our knowledge, this is the first paper to introduce experimental variation in order to satisfy the exclusion restrictions needed for sequential identification.[14] Our approach resembles a selective trial as described by Chassang et al. (2012) in that the take-up decision reveals private information about farmer type. However, we use random variation in the take-up subsidy as opposed to incentive compatible mechanisms to reveal agent's information about her own returns. In addition, we allow for the information set available to the agent to change between the time the take-up decision is made and the time the follow-through decision is made. Along with this methodological contribution, our paper pulls together two previously unrelated strands of literature: first, the literature on subsidies for technology adoption in development (e.g., Cohen and Dupas 2010; Ashraf et al. 2010; Berry et al. 2012; Oster and Thornton 2012; Dupas 2014), and second, the literature on investment under uncertainty (e.g., Pindyck 1993; Dixit and Pindyck 1994).[15]

The rest of the paper proceeds as follows. We begin with a simple theoretical model to generate intuition for the adoption problem. We turn in Section 2.3 to a description of the empirical context and experimental design, and show summary statistics and reduced form results in Section 2.4. We present the empirical model and its identification in Section 2.5 and show estimation results and simulations in Section 2.6. Section 2.7 discusses the interpretation of the results, including learning and procrastination, and Section 2.8 concludes.

[14] Combining field experiments with structural modeling is an increasingly popular approach to extending the generalizability of experimental findings and testing among alternative models of behavior.

[15] A small number of other papers span these literatures. For example, Bryan et al. (2014) study a household's decision to migrate under uncertainty about the payoffs from doing so. They find substantial increases in migration and profits in response to a small subsidy for migration. Levine et al. (2012) also find evidence consistent with adoption under uncertainty. In their study, money-back guarantees and free trials increase adoption of an improved cookstove in Uganda. A much larger literature in development economics examines risk and uncertainty as part of the technology adoption process (see Foster and Rosenzweig (2010) for a summary), though little attention is paid to the implications for adoption subsidies.

2.2 A Simple Model of Intertemporal Technology Adoption

Consider a two period model, where each agent choses to purchase a single unit of a technology (take-up) in the first period (time 0) at a cost of c, and to follow-through with implementation of the technology in the second period (time 1). There is a net cost of complying, $F_0 + F_1$, with only F_0 known to the agent at the time the technology is purchased. The remaining component of the cost, F_1, is random with known distribution to the agent at time 0,[16] and its realization is revealed to the agent at time 1.[17] $F_0 + F_1$ can be interpreted as the "net cost" of following-through for the agent and can be positive (if costs outweighs benefits) or negative (if benefits outweigh costs). Note that although F_0 is known at time 0, both F_0 and F_1 are incurred at time 1. In addition to the private (or intrinsic) net costs faced by the agent, there could be external economic incentives to either take-up the technology (a subsidy, A) or to follow-through (a reward, R). We assume these, as well as the unsubsidized cost c, are constant across agents.

We assume linear utility – or risk neutrality – throughout the paper, including the empirical analysis. Although risk aversion is an important component of intertemporal decisions with costs or benefits that represent substantial shares of household income, our theoretical and empirical frameworks are best applied to technology adoption decisions that will cause relatively small changes to income. Importantly, our theoretical and empirical frameworks model decisions as a function of the profits associated with adoption relative to the best alternative use of household resources. Thus, a positive shock to the opportunity cost of adoption could correspond to an increase or a decrease in overall household income. Incorporating risk aversion into our theoretical model would require us to make modeling assumptions about the nature of the opportunity cost of adoption. For example, an increase in profitability of a competing economic activity and a labor shortage due to health could both represent an

[16]Note that because the distribution of F_1 is mean zero by construction since its distribution is known to the agent and it enters the profit function additively, thus any non-zero mean is known to the agent at time 0 and incorporated into F_0.

[17]The information revelation process can be one of learning about the technology itself or one of shock realization. The current model is forward-looking and relevant for both. Ex post, they diverge. As we move to our field setting, we will attempt to distinguish these two interpretations of the nature of the information.

increase in the opportunity cost of adoption, but would have opposite effects on total income and thus on marginal utility of income. Hence, we prefer to leave the source of the opportunity cost unspecified, which makes our framework generalizable to any source of uncertainty to the extent that adoption decisions do not result in large changes in income.[18]

Following backward induction, the agent's decision at time 1 is whether to follow-through or not, and this decision is made to maximize her utility at each point in time. With a linear utility assumption, the net benefit derived by the agent in period 1 is given by $R - F_0 - F_1$ if she choses to follow-through and 0 if she choses not to follow-through.

At time 0, the agent's decision is whether or not to take-up by purchasing the technology. The agent chooses to take-up if

$$c - A - \delta \mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) > 0 \tag{2.1}$$

where δ is the one-period discount factor and the expectation in (2.1) is taken with respect to the conditional density of F_1 given F_0.

To simplify the exposition, assume $F_0 \perp F_1$, and that the random net cost component takes one of two values

$$F_1 = \{f_L, f_H\}, \text{ with } f_L < f_H, \text{ and } \mathbb{E}_{F_1|F_0}(F_1) = \mathrm{g}_1(f_L) f_L + \mathrm{g}_1(f_H) f_H$$

where $\mathrm{g}_1(.)$ is the probability mass function of F_1. The independence assumption on F_1 and F_0 implies that $\mathrm{g}_1(.)$ is constant across agents. Assume also that F_0 is continuously distributed across agents with cumulative distribution function $\mathrm{G}_0(.)$.[19]

The distributional assumption on F_1 allows us to classify individuals into three types: those who never follow-through (hereforth *non-adopters*), those who follow-through only if the low net cost shock is realized (hereforth *contingent adopters*) and those who always follow through,

[18] In Section 2.5, we document that the size of the benefits derived from our program is small relative to total household income; and thus the risk neutrality assumption is unlikely to affect our conclusions.

[19] This model simplifies our empirical setting in two key ways: first, it assumes a binary follow-through decision and second, it assumes a discrete distribution on F_1. As we show when we present our empirical model, the propositions derived from this model are not an artifact of the distributional assumption on F_1 nor of the binary decision that characterizes follow-through in this simple model. The independence assumption between F_0 and F_1 can be thought of as an assumption on how uncertainty enters profits in the model: in an additive way.

regardless of the realization of F_1 (hereforth *always adopters*). These three types of agents can be characterized by their value of F_0 (what they know about their net costs at the time of take-up): non-adopters have the highest value of F_0 (between $R - f_L$ and ∞); contingent adopters have a mid-size value of F_0 (between $R - f_H$ and $R - f_L$); always adopters have the smallest value of F_0 (between $-\infty$ and $R - f_H$). Figure 2.1 graphically shows the proportions for each type of agent using areas under a symbolic bell-shaped distribution for F_0, separated by gray dashed lines. Figure 2.1 also illustrates two thresholds (along the support of F_0) for take-up in black dashed lines. The first take-up threshold (labeled $R - \mathbb{E}(F_1) - \frac{c-A}{\delta}$) is only binding if it falls to the left of the threshold that defines always adopters ($R - f_H$). When this first take-up threshold binds, only always adopters take-up. The second take-up threshold (labeled $R - f_L - \frac{(c-A)}{\delta g_1(f_L)}$) is perhaps more interesting. When binding, it determines the full share of types that takes-up, which includes all always adopters and a subset of contingent adopters. We further use this figure to explain intuitively the results outlined by each of our propositions below. The formal proofs of these propositions can be found in Appendix B.1.

Proposition 1 *Under low levels of uncertainty, follow-through conditional on take-up increases as a function of the total (potentially subsidized) take-up cost.*

To see this, note that as take-up cost increases (represented by $c - A$ in Figure 2.1), the second take-up threshold moves to the left, bringing down the overall share of contingent adopters among the set of individuals who take-up. Since contingent adopters follow-through with probability less than one ($g_1(f_L)$), this in turn increases the share of individuals who follow-through among those who take-up. Proposition 3 shows how this relationship weakens as uncertainty increases.[20]

[20]If the take-up cost, $c - A$, increases enough that the first take-up threshold becomes binding, follow-through conditional on take-up reaches 100 percent and stays constant for further increases in the take-up cost.

Figure 2.1: *Take-up and Follow-through Thresholds as a Function of Agent Type*

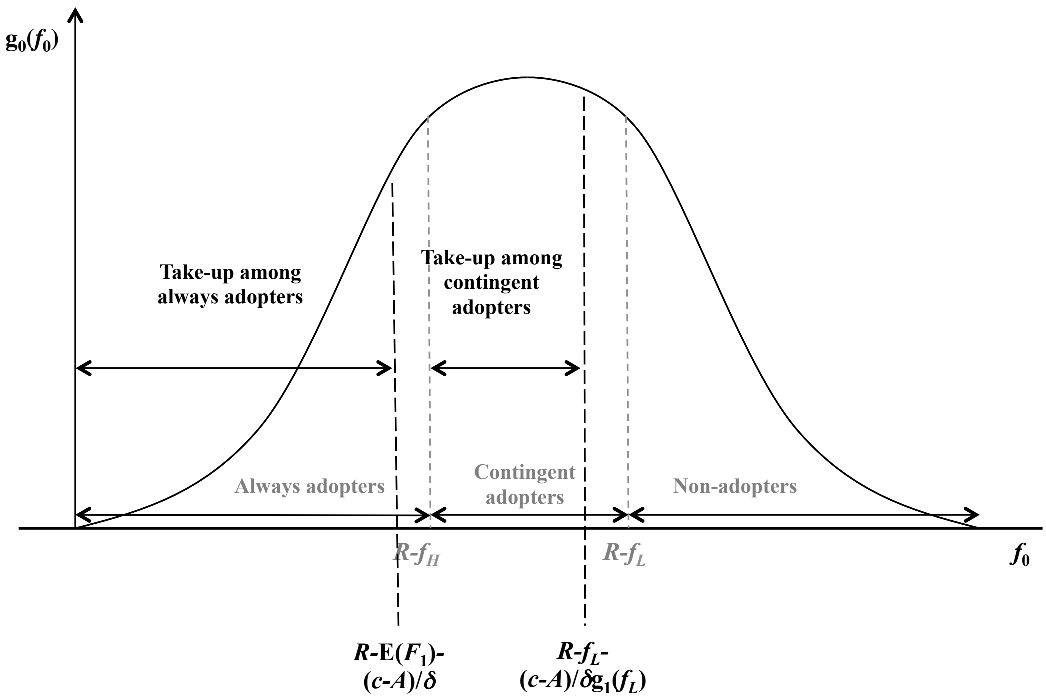

The figure shows the shares of always adopters, contingent adopters and non-adopters over a symbolic probability density function of F_i. The grey thresholds ($R - F_H$ and $R - F_L$) correspond to the follow-through thresholds, while the black thresholds correspond to the take-up thresholds.

Proposition 2 *An increase in uncertainty reduces follow-through conditional on take-up.*

Given the assumed distribution for F_1, an increase in uncertainty can be represented by an increase in the distance between f_L and f_H. This, in turn, causes the share of contingent adopters to increase (as the two grey dashed lines move further apart). Note that as uncertainty increases, the position of the second take-up threshold does not change relative to the threshold that determines the upper bound for contingent adopters. Thus, contingent adopters become a larger share of those who take up, reducing average follow-through.

Corollary 2.1 *Under no uncertainty, everyone who takes-up follows-through.*

This is easy to see from figure 2.1, as under no uncertainty (where $f_L = f_H$) there would be only always adopters and never adopters.

Proposition 3 *An increase in uncertainty weakens the relationship between take-up cost and conditional follow-through shown in Proposition 1.*

To see this, visualize the take aways of Propositions 1 and 2 simultaneously in the graph. Under a high level of uncertainty (recall uncertainty expands the number of contingent adopters while keeping the relative position of the second take-up threshold and the $R - f_L$ line constant), the take-up cost is less effective at deterring contingent adopters from taking-up. More precisely, the share of contingent adopters that are excluded by an increase in the take-up cost becomes a smaller proportion of all those who take-up when uncertainty increases.

Proposition 4 *The option value associated with take-up is increasing in uncertainty, which results in higher take-up at all take-up cost levels.*

We define option value as the difference between value of a dynamic contract (with a two-period decision framework, as outlined above) and a static contract (where take-up and follow-through decisions are made simultaneously at time 0):[21]

$$OV(F_0) = \mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) - \max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right) \tag{2.2}$$

Under this hypothetical static contract, a farmer would be required to stick to a decision on follow-through made jointly with the take-up decision. The value of the static contract (second term to the right of the equal sign in 2.2) does not vary with uncertainty: the symmetric distribution of shocks means that the ex post expected profit does not increase (or decrease) with the variance of the shock. Thus, this result also implies that there is a positive relationship between the expected profit of the contract and uncertainty stemming solely from the response of the option value to uncertainty.

In the appendix, we show how this quantity increases with uncertainty and has an asymmetric relationship with the upper and lower bounds of the F_1 distribution. Intuitively, the option value is the value of "reoptimizing" after the take-up decision is made and the new information (the realization of F_1) emerges. As uncertainty ($f_H - f_L$) increases, the gains from

[21] This corresponds most closely to the "value of waiting" in Dixit and Pindyck's (1994) framework, when the decision is over whether to abandon a project.

follow-through conditional on a low cost shock (f_L) increase. Because agents can choose not to follow through, f_H is less important in determining the take-up decision. Thus, an increase in uncertainty increases the option value of the contract, resulting in higher take-up.

2.2.1 Common Shocks, Transitory Shocks and Learning

So far, we have left open the question of whether F_1 should be interpreted as a permanent or a transitory shock. That is, if the F_1 component of the returns to the technology is permanent, the agent will learn about it in period 1 and will incorporate it into future choices about the technology. If F_1 is transitory, the same decisions in the future will have the same level of uncertainty and the same option value associated with take-up as when the technology is introduced for the first time. Note that the interpretation does not matter for the exposition so far, since we restrict our attention to the first instance a new technology is introduced. However, incorporating the private and social value of future take-up and follow-through decisions could change conclusions regarding the optimal take-up level. We cannot completely disentangle these two interpretations of the model in our context, though we use survey data to provide suggestive evidence on the extent of learning (see Section 2.7).

In addition, we have not specified whether F_1 represents common or idiosyncratic shocks. This distinction does not play a role in the analysis so far, since farmers can form expectations on either type of shock as long as they have knowledge of its distribution. However, this distinction will play a more important role in the empirical framework. Thus, it is illustrative to understand how common and idiosyncratic shocks will affect cross-sectional variation in outcomes using this simple theoretical model. Assume all farmers face the same distribution of F_1. If F_1 represents a common shock, all farmers will have the same realization (either f_H or f_L) of the shock ex-post. Hence, all contingent adopters will either follow-through (if f_L is realized) or not (if f_H is realized), i.e. a common shock does not translate into cross-sectional variation in follow-through choices among contingent compliers, while idiosyncratic shocks do. Note, however, that the cross-sectional variation in the take-up decision is unchanged by the nature of the shock.

2.3 Context and Experimental Design

We bring the propositions from our conceptual model to a two part technology adoption problem, characterized by uncertainty in the costs and benefits of following through with the technology. In the context of an ongoing project to encourage the adoption of agroforestry trees (*Faidherbia albida*), we introduce exogenous variation in the payoffs to farmers at the time of their take-up and follow-through decisions. Our ultimate goal is to use experimental variation to uncover the existing levels of static heterogeneity and uncertainty in the population of farmers, which we model as random parameters. This design relies on orthogonal assignment to treatments that affects farmers at two different points in time to reveal how farmers' cost-benefit calculation, and therefore the information they have, changes between the first (take-up) and second (follow-through) decisions. This section describes the context and the experimental design in detail and Section 2.4 shows some reduced form results that provide simple empirical facts consistent that are consistent with the presence of uncertainty. We discuss identification assumptions and estimation of the structural model in more detail in Section 2.5.

The study was implemented in coordination with Dunavant Cotton Ltd., a large cotton growing company with over 60,000 outgrower farmers in Zambia, and with an NGO, Shared Value Africa. The project, based in Chipata, Zambia, targeted approximately 1,300 farmers growing cotton under contract with Dunavant, alongside other subsistence crops. The project is part of the NGO partner's portfolio of carbon market development projects in Zambia.

2.3.1 The Technology

Faidherbia albida is an agroforestry species endemic to Zambia that fixes nitrogen, a limiting nutrient in agricultural production, in its roots and leaves. Optimal spacing of *Faidherbia* is around 100 trees per hectare, or at intervals of 10 meters. The relatively wide spacing, together with the fact that the tree sheds its leaves at the onset of the cropping season, means that planting *Faidherbia* does not displace other crop production (Akinnifesi, 2010). Agronomic studies suggest significant yield gains from *Faidherbia*. However, these private benefits take 7-10 years to reach their full value, and may be insufficient to justify the up-front investment costs. Indeed, a back of the envelope estimate suggests that, at the high discount rates typically observed in developing country settings, the adoption of *Faidherbia* is not a privately profitable investment for most farmers. Average benefits for adoption on half a hectare of land are around USD 2 and average costs around USD 19.[22] Negative net benefits, on average, are consistent with the low adoption rates we observe at baseline: less than 10 percent of the study households reported any *Faidherbia* on their land.[23]

Subsidies may therefore be necessary to increase take-up rates, and are justified by positive environmental externalities and market failures that contribute to high private discount rates. Environmental benefits include erosion control, wind breaks, and carbon sequestration. Based on allometric equations from Brown (1997), adapted to the growth curves for *Faidherbia*, we estimate that over the course of 30 years, a tree sequesters around 4 tons of carbon dioxide equivalent. Discounting the annual sequestration at 15 percent leads to a present value of around 0.48 tons. At a marginal social cost of carbon of 21 USD/ton, the present value of sequestration associated with half a hectare of surviving trees is around USD 353.

Both the private and the public benefits associated with adoption require that farmers continue to invest in the technology, even after their initial take-up decision is made. To obtain the private and public benefits, farmers must plant, water, weed and otherwise care for the trees, activities that are costly in the short run. In addition, the opportunity cost of these investments may be highly uncertain and depend on shocks to household labor supply, weather, pests and prices, all of which are realized after the initial take-up decision is made. Therefore, the technology maps clearly onto our conceptual framework, with farmers deciding

[22] Agronomic studies suggest yield gains from planting *Faidherbia* range from 100 to 400 percent, relative to production without fertilizer (Saka et al., 1994; Barnes and Fagg, 2003), corresponding to income gains of up to USD 400 per hectare per year. For our back of the envelope calculation, we take mid-range values from the literature and assume that yields on half a hectare of land increase by 525 kilograms, starting 10 years after trees are planted. At current maize prices, this equals additional income of 136.50 USD per year from years 10 to 20. At a discount rate of 0.67, which is based on survey data and the literature, the present value of this investment is USD 2.02. While this discount rate is high, it is in line with observed interest rates and elicited individual discount rates in rural developing country settings (Conning and Udry, 2007; Cardenas and Carpenter, 2008). It corresponds to a discount factor of 0.6, which we employ in the structural model. With respect to the private costs, surveys implemented in conjunction with our study indicate that farmers invest around 38 hours planting and caring for the trees. At an agricultural wage rate of around USD 2.5 per day, these time investments correspond to USD 19 of labor costs incurred in the first year.

[23] Informal land tenure in the project area presents an additional barrier to adoption. By focusing on landholders engaged in contract farming arrangements, the project targets households with relatively secure tenure.

whether or not to follow-through by choosing whether to continue to invest in keeping trees alive after the initial take-up decision.

2.3.2 Experimental Design and Randomization

Our experimental design serves two purposes. First, the treatments allow for reduced form tests for the presence of uncertainty about the costs and benefits of follow-through at the time of the farmer's take-up decision stemming from our conceptual model in Section 2.2. Second, the experimental variation identifies a structural model of intertemporal technology adoption under uncertainty. The study design varied two major margins of the farmer's decision to adopt *Faidherbia albida*. All farmers were given a take it or leave it offer of a fixed number of seedlings (50, or enough to cover half a hectare) to be planted and managed by the farmer and his or her household. The cost of the seedlings (the take-up decision) was subsidized in some cases. The program offered a threshold payment conditional on follow-through (tree survival) after one year (time 1). Farmers received the reward if they kept 70 percent (35) of the trees alive through the first dry season (for 1 year). The threshold reward, as opposed to a per-tree incentive, allows us to draw a sharper distinction between internal incentives and external incentives to cultivate the trees, which aids identification of the structural model. [24] We collected data on the binary take-up decision, the number of trees that survived after a year (whether or not this number was above the 35-tree threshold) and whether or not the farmer earned the reward associated with the 35-tree threshold.

Because the threshold reward pays out after a year after which no further incentives for tree survival are offered, farmers' decisions on take-up and follow-through are based on their perceptions about costs and benefits during the first year only. The first year is the most relevant on the cost side, since the costs associated with planting and caring for the trees are highest during the first year when the trees are vulnerable and require substantial attention

[24] In addition to its advantages in our empirical design, the threshold reward was easy to enforce, consistent with other related contracts offered by Dunavant, and easy to explain to the farmers. The threshold was chosen based on *Faidherbia* survival rates in other programs in Eastern Zambia. Specifically, for a sample of around 3000 farmers tracked by another NGO, which offered no performance incentive, a 70 percent survival rate was achieved by around 20 percent of farmers.

in the form of watering, weeding and protection from pests. After they survive the first dry season, costs decrease substantially. The follow-through decision we observe is more accurately described as the cumulative outcome from numerous follow-through decisions made over the course of the year after take-up. New information may reveal itself starting immediately after the take-up decision is made, or gradually, as family members fall ill, crops fail, or input and output prices change.[25] Each time new information arrives that affects the opportunity cost of caring for the trees, farmers may reoptimize on the number of trees they continue to cultivate (if any). Therefore, at the time that we measure follow-through, we observe the cumulative effect of all of these decisions. On the benefit side, we expect to see little change in information within the first year since the private benefits take considerably longer than the costs to materialize. Of course, farmers may still face uncertainty about the costs and benefits of keeping trees alive even after follow-through is measured.[26]

Figure 2.2 summarizes the experimental design. First, the size of the input subsidy (A) was varied at the farmer group level in increments of 4,000 ZMK from a subsidy of 12,000 ZMK (fully subsidized) to zero. At zero subsidy, farmers paid 12,000 ZMK (approximately USD 2.60) for inputs, which is the cost recovery price for the implementing organization. Groups were randomly assigned to one of four input subsidy treatments using the min max T approach (Bruhn and McKenzie, 2009), balanced on Dunavant shed, farmer group size and day of the training. Second, the size of the threshold performance reward (R) was varied at the individual level, in increments of 1,000 ZMK, ranging from zero to 150,000 ZMK or approximately 30 USD.[27] Variation in the reward was introduced using a random draw at the time of the take-up decision. One-fifth of all draws were for zero ZMK with the remaining four-fifths distributed

[25]The take-up decision is made at the beginning of the planting season. This is the natural timing of take-up decisions for other crops and technologies. Therefore, our design allows for an amount of time between take-up and follow-through that is similar to many other agricultural technologies.

[26]As mentioned in the previous section, the fertilizer benefits of the technology do not appear until years 7-10. The only potential for learning during the course of the experiment is therefore about the costs associated with keeping trees alive, which is an activity similar to what most farmers do each year with other crops.

[27]At the time of the study, the exchange rate was just under 5000 ZMK = 1 USD. In piloting, the distribution of payments extended to 200,000 but was scaled back prior to implementation. The scratch cards with values between 150,000 and 200,000 were removed from the prepared cards by hand, but six of them were missed. For the main analysis, we top-code payments at 150,000.

uniformly over the range.

Figure 2.2: *Experimental Design*

Timing	Treatment	Randomization
Time 0	Take-up subsidy A = 0, 4000, 8000, 12000	Group level
	Follow-through threshold reward announced R = (0 (1000) 150000)	Individual level
Time 1 (1 year later)	Follow-through threshold reward administered	Individual level

Treatments are administered at two points in time: time 0 (November 2011) and time 1 (November 2012). The take-up subsidy is applied to the bundle of 50 trees that the farmer can choose to adopt or not. The follow-through threshold reward pays out conditional on the survival of 35 or more of those trees one year later. The take-up subsidy is varied at the farmer group level; the follow-through threshold reward is varied at the individual level. See text for a description of implementation.

Third, the timing of the reward draw was varied at the individual level to occur either before or after (a surprise reward treatment) the farmer's take-up decision. The timing was pre-determined based on the registration number. A roughly equal number of farmers were assigned to each timing treatment (52.5% were assigned to the surprise reward treatment). Recall from Section 2.2 that R has a direct effect on both the take-up decision and the follow-through decision, while A affects follow-through only via its effect on the types of agents who choose to take-up. We use the variation in the timing of when the reward was announced to separate the effect of R on take-up and on follow-through, in a similar spirit to Karlan and Zinman (2009).[28] This third source of variation also provides an additional test for liquidity constraints as a driver of selection outcomes.

[28]We do not manipulate or measure beliefs about potential financial benefits from joining in the surprise reward treatment, and cannot therefore assume that farmers in the surprise reward treatment assumed $R = 0$ at the time of take-up.

2.3.3 Data and Implementation

The field experiment was implemented between November 2011 and December 2012 with 125 farmer groups and around 1300 farmers. Implementation of the study relied on Dunavant's outgrower infrastructure, which is organized around sheds, each of which serves several dozen farmer groups. The farmer groups consist of 10-15 farmers and a lead farmer, who is trained by Dunavant each year and in turn trains his or her own farmers on a variety of agricultural practices. The study included the following stages: farmer training, program enrollment, baseline survey, end line survey, tree monitoring, reward payment. We describe each of these stages in turn.

Farmer Training

Lead farmers attended a training on *Faidherbia* at their Dunavant shed, which provided them with the information and materials necessary to train their own farmer groups.[29] The group trainings were scheduled by the project team for approximately four weeks after the lead farmer training and were attended by study enumerators. Either the household head or an immediate family member participated in the group training, as per Dunavant's usual practices. At the group training, farmers were provided with instructions on planting and caring for the trees, were given information about the private fertilizer benefits and public environmental benefits of the trees, and were informed about the eligibility requirements for the program.[30] All farmers who attended the group training received a show up fee of 12,000 ZMK and lunch.

---

[29] As described above, this is standard practice. At the lead farmer training, study enumerators collected basic information from the lead farmers for input into the randomization. Thus, any selection into the *Faidherbia* training by lead farmers is orthogonal to treatment. In addition to training their own farmers, lead farmers were responsible for constructing a nursery in their community, and were paid for the labor and materials. Backup nurseries were also constructed at the Dunavant sheds.

[30] Eligibility required that land must have been un-forested for 20 years, must be owned by the farmer, and must not be under flood irrigation. Study enumerators confirmed that all lead farmers covered the necessary training material and summarized the main points at the end of the training to ensure that all farmers had accurate information.

Program Take-Up

After the technical training, study enumerators explained the details of adoption decision, including the cost of the inputs (12,000 ZMK minus the subsidy), which was announced to the group as a whole. The un-subsidized price (12,000 ZMK) is approximately one day's agricultural wages. To implement the individual-level randomization of the rewards and allow participants to make their take-up decision in private, farmers were called upon individually by study enumerators. The farmer was told that she could earn a reward conditional on reaching a threshold of 70 percent tree survival after one year. The farmer then drew a scratch off card from a bucket, which revealed the individual reward value, after which the take-up decision was recorded. The surprise reward treatment proceeded similarly but the take-up decision was recorded before the farmer was informed about the threshold reward. Conditional on agreeing to join the program, the farmer was told about the rewards and given the opportunity to draw a scratch card. After the value was revealed, farmers could change their minds about take-up. No one did.

Survey Data

Following the take-up decision, all farmers were given a baseline survey that lasted for approximately one hour. Of the 1314 farmers for whom a participation decision was recorded, 1292 are in the baseline survey sample. After the survey, participating farmers signed a contract indicating their agreement with the program terms, paid the take-up cost and collected their seedlings.[31]

Two additional data collection instruments were administered between the baseline and endline surveys. First, one-fifth of the farmers were randomly sampled for ongoing visits to collect information on activities and inputs associated with the contract and with other crops. These visits, which we refer to as effort monitoring, consisted of a very short survey

[31]Some farmers chose to leave their seedlings with the lead farmer for later collection. In a small number of cases, the lead farmer's nursery had not generated enough seedlings to support the group. In those cases, seedlings from a backup nursery were transported to the site.

(~20 minutes), during which a project monitor asked the farmer about agricultural activities since the last visit.[32] We control for the effort monitoring subsample in our analysis but do not analyze its effects on program outcomes. Second, all farmers were visited at the end of the planting season for a very brief survey on seedling pick up, transplanting and care, primarily for the purpose of obtaining qualitative information about the timing of planting and tree care activities.

All farmers in the study sample were given an endline survey, regardless of their decision to participate in the contract. Of the 1292 farmers in the baseline survey, 1232 (>95 percent) are also in the end line survey.

Tree Survival Monitoring and Reward Payments

One year after program enrollment, tree survival outcomes were monitored and rewards delivered, conditional on reaching the 35-tree threshold. Approximately one week after the endline survey, farmers with contracts were visited for field monitoring, during which the farmer and a study enumerator examined each tree, and recorded whether it was sick, healthy or dead. All surviving trees counted toward the tree survival threshold. In case of disputes, a third party was called in from the village.

Finally, within a couple of days of the monitoring visit, the reward payment team visited each farmer group and paid out rewards to farmers with 35 or more surviving trees. Keeping the payments separate from the monitoring was intended to improve monitors' objectivity.[33] Of the farmers eligible for monitoring, 9 were not located. We assume zero tree survival for these individuals in our main analyses.

[32]No information was provided to the farmers about their performance and monitors were instructed not to prompt specific activities or answer technical questions.

[33]As a check for collusion between the monitors and farmers, we test whether individual monitors are associated with a higher probability that a farmer passes the tree survival threshold. No single monitor indicator is significantly correlated with reaching the threshold, nor are the monitor indicators jointly predictive. This implies that either all monitors were cheating to the same degree or that no monitors were cheating. Given differences in career concerns with the study implementers (some had higher paid jobs as survey supervisors when not engaged in monitoring), cheating by all monitors is unlikely.

2.4 Summary Statistics and Reduced Form Results

Sample Characteristics

Appendix Table B.5 shows summary statistics from the baseline survey. Around 70 percent of participants are heads of household and 13 percent of households are female headed. Respondents have, on average, just over 5 years of education and live in households with just over 5 members. Households have around 3 hectares of land spread across just under 3 fields, which are an average of around 20 minutes away from their dwelling. Around 10 percent of households state that soil fertility is one of the major challenges that their household faces. Households have worked with Dunavant Cotton for an average of over 4 years and over 40 percent interact regularly with their lead farmer. Very few are affiliated with other local organizations that promote *Faidherbia albida* (CFU and COMACO). Almost 70 percent of respondents report familiarity with the technology but only around 10 percent had adopted prior to the program.

Balance and Attrition

We test the randomization outcomes by comparing observable characteristics across treatment groups. Appendix Table B.5 tests balance for the take-up subsidy, threshold performance reward, and surprise reward treatment. We observe some imbalance in the assignment of the take-up subsidy treatment with slightly larger households with more non-agricultural assets receiving lower input subsidy assignments on average. We also observe that older respondents with larger households and better self-reported soil fertility are marginally more likely to be assigned to the surprise reward treatment. The table tests balance for 17 variables and consists of 51 separate regressions. Five significant coefficients is therefore consistent with significance threshold of 10 percent.

We also examine whether non-random attrition may undermine internal validity.[34] We

[34]Selection into treatment is also a threat to the experiment's internal validity. By design, this is unlikely: group level participation subsidy treatments were revealed only after individuals arrived for training, and individual-level reward treatments were assigned in a one-on-one interaction with study enumerators.

examine attrition at each stage of data collection in Appendix Table B.6. Take-up rates in the baseline survey are over 98 percent, while the end line included 94 percent of all study farmers and over 95 percent of baseline respondents. We see some evidence that farmers who received lower take-up subsidies were marginally less likely ($p < 0.10$) to participate in the surveys. Otherwise, attrition is balanced across treatments. For the tree survival monitoring, over 95 percent of the 1092 households that took up the program were located.

Finally, spillovers across treatments pose a threat to the experimental design. Because the take-up subsidy treatment was assigned at the group level, spillovers are relatively unlikely. During regular effort monitoring visits, no farmers asked about differences in take-up costs. The value of the threshold reward, on the other hand, varied at the individual level within-group. By revealing the reward value privately to each farmer before the take-up decision, we mitigate the potential that take-up is affected by rewards received by others.[35] However relative reward values may still affect performance since farmers can share information after they leave the training. We observe some suggestive evidence that a higher average reward among other farmers in the group has a positive spillover on own survival probabilities, though the magnitude of the effect is small relative to the direct effect of the reward.[36]

[35]This is confirmed empirically by regressing the probability of take-up on the average random reward draw that preceded a farmer's own draw. The outcomes of preceding draws have no effect on the probability of take-up.

[36]Randomization of rewards at the individual level also gives rise to concerns about reselling seedlings to those with higher rewards or transplanting young trees just before monitoring. We use several pieces of information to investigate these concerns. First, we would expect that take-up would reflect the potential to re-sell if farmers were aware of the individual-level variation in rewards. In other words, farmers that were aware of the arbitrage opportunities generated by the variation in the reward treatment should be more likely to take-up, and increasingly so as the size of the subsidy falls (i.e. would respond less to the subsidy). However, a regression of take-up on the interaction of the surprise reward treatment and the level of the input subsidy shows no significant interaction or clear pattern of coefficients. Second, to investigate the potential for transplanting, we take advantage of planting data collected for all farmers shortly after the end of the rainy season. We construct a measure of the difference between the planting and the monitoring tree counts, which is positive for around 100 farmers. The positive value indicates either very delayed planting or transplanting. Restricting attention to those with a positive value, the coefficient from regressing this measure of extra trees on the size of the reward is insignificant, and becomes negative (and insignificant) when group fixed effects are included. Third, we examine the within-group spillovers associated with the effect of the reward on tree survival outcomes. To the extent that transfers of any kind are happening within group, we expect a steeper slope on the reward within-group than on average. We observe a slightly smaller and statistically indistinguishable coefficient on the reward when group fixed effects are included, relative to the coefficient without fixed effects. Appendix Table B.4 shows spillover effects of the rewards.

2.4.1 Reduced Form Results

Table 2.1 displays means and standard deviations for several program outcomes: take-up, follow-through (tree survival ≥ 35), zero surviving trees and the number of trees conditional on positive survival rates. These statistics are broken down by treatment and show clear patterns in responses to the incentives offered in the experiment: take-up responds positively to the price subsidy and with the threshold reward, and follow-through responds positively to the threshold reward.

Table 2.1: *Summary Statistics*

	(1) Take-up	(2) 35-tree threshold	(3) # trees | # trees> 0	(4) Zero trees
		Panel A: full sample		
mean	0.83	0.25	27.42	0.36
sd	0.38	0.44	14.31	0.48
		Panel B: by take up subsidy treatment		
A = 0				
mean	0.71	0.26	27.60	0.37
sd	0.46	0.44	14.31	0.48
A = 4000				
mean	0.76	0.29	28.86	0.36
sd	0.43	0.45	13.67	0.48
A = 8000				
mean	0.86	0.27	29.30	0.38
sd	0.35	0.44	14.19	0.49
A = 12000				
mean	0.97	0.22	24.93	0.33
sd	0.17	0.41	14.52	0.47
		Panel C: by reward treatment		
R = 0				
mean	0.90	0.13	22.00	0.49
sd	0.31	0.34	14.70	0.50
R = (0,70000]				
mean	0.90	0.21	25.45	0.40
sd	0.30	0.41	14.62	0.49
R = (70000,150000]				
mean	0.93	0.32	29.53	0.30
sd	0.25	0.47	13.67	0.46

Notes: Means and standard deviations of take-up (column 1) and follow-through (columns 2-4) outcomes, by experimental treatment. Column 1 includes all farmers (N=1314). Columns 2-4 are conditional on take-up (N=1092). Column 2 reports the number of farmers who reached the performance reward threshold.

Uncertainty

We use the means and standard deviations presented in Table 2.1 to provide evidence for the presence of uncertainty, consistent with our conceptual model. Regression-based results are shown in Table 2.2. First, notice that take-up rates are increasing across values of the take-up subsidy. Take-up rates are high, on average, even in the zero subsidy condition, where over 70 percent of farmers take-up. This could be due to high known payoffs from follow-through, on average, or to high expected values driven by option value (see Proposition 4).

Table 2.2: *Comparison of Structural and Reduced Form Estimates*

	(1) Take up	(2) 35-tree threshold	(3) # trees | # trees>0	(4) 1.(zero trees)		(5) Take up	(6) 35-tree threshold	(7) # trees | # trees>0	(8) 1.(zero trees)
					Panel A. Observed Data				
Input Subsidy	0.022***	-0.004	-0.229	-0.003	Reward	0.001*	0.001***	0.044***	-0.001***
	(0.005)	(0.004)	(0.200)	(0.005)		(0.000)	(0.000)	(0.013)	(0.000)
Observations	1,314	1,092	701	1,092		624	1,092	701	1,112
R-squared	0.071	0.002	0.005	0.001		0.006	0.018	0.022	0.013
					Panel B. No Mean Shift				
Input Subsidy	0.021***	-0.003	0.061	0.010***	Reward	0.001*	0.003***	0.091***	-0.001***
	(0.002)	(0.003)	(0.130)	(0.003)		(0.000)	(0.000)	(0.012)	(0.000)
Observations	1,314	1,104	625	1,104		624	1,104	625	1,104
R-squared	0.069	0.001	0.000	0.009		0.005	0.105	0.082	0.011
					Panel C. Allowing for Mean Shift				
Input Subsidy	0.020***	-0.004	0.022	0.009***	Reward	0.001*	0.003***	0.094***	-0.001***
	(0.002)	(0.003)	(0.131)	(0.003)		(0.000)	(0.000)	(0.012)	(0.000)
Observations	1,314	1,112	603	1,112		624	1,112	603	1,112
R-squared	0.066	0.001	0.000	0.007		0.005	0.104	0.088	0.013

Notes: This table shows coefficients from regressions of each of four indicator variables (take up, threshold of 35 trees reached, tree survival larger than zero, and no tree survival) on each of our randomized treatments (input subsidy and threshold reward). Panel A shows these regression outcomes for the true data. Panel B and C show the fit of the strcutural model by simulating all four outocomes using the model estimates and examining the how much the linear relationships between outcomes and treatments resemble those in Panel A. Panel B uses baseline model estimates (Panel A of Table 2), while Panel C uses estimates from the mean shift model (Panel B of Table 2).

Second, we observe that follow-through rates vary considerably within subsidy and reward treatments and are low, on average, with only 25 percent of farmers reaching the 35-tree threshold (column 2). This holds even in the zero subsidy condition, ruling out that farmers were certain about high payoffs associated with cultivating a large number of trees at the time of take-up. Low follow-through conditional on take-up is consistent with Proposition 2.[37]

Third, a large number of farmers abandon the technology altogether (have a survival of zero trees), even conditional on taking up with zero subsidy (37 percent, column 4). This rules out that farmers were certain about positive payoffs from a small number of trees at the time of take-up, and is consistent with Corollary 2.1 of our conceptual model.

Finally, we see no reduced form effect of the subsidy treatment on the likelihood of reaching the 35-tree threshold or of abandoning the technology (zero trees). More formally, we implement a two-sample t-test for equal means between the highest and lowest subsidy condition. For continuous tree survival, the probability of reaching the threshold ($\geq$ 35 trees) and zero trees, the p-values are 0.63, 0.25 and 0.32, respectively. The linear regression test of the effect of the take-up subsidy on tree survival outcomes (shown in Table 2.2) is also statistically insignificant. This is consistent with Proposition 3, which states that the selection effect of subsidies will be diminished by high levels of uncertainty in the net benefits of follow-through.

We also examine whether outcomes can be explained by observables. Appendix Table B.7 shows that, overall, observables explains relatively little of the variation in outcomes: the R-squared from a regression of outcomes on observables is 0.0296, 0.0297 and 0.0314 for take-up, reaching the 35-tree threshold and tree survival, respectively. Adding the treatment variables improves the explanatory power substantially (even numbered columns). The low explanatory power of observables further motivates our use of a structural model to estimate the heterogeneity across farmers at both take-up and follow-through.

[37] Behavioral explanations such as over-optimism or procrastination might also be consistent with high take-up and low-follow through, even at positive take-up prices. We discuss behavioral explanations consistent with the reduced form results, as well as the interpretation of the type of new information, in Section 2.7.

Liquidity Constraints

Our design allows us to address concerns that liquidity constraints complicate the relationship between the take-up decision and farmers' expected payoffs at the time of take-up. We do this in two ways. First, as noted above, all farmers received a training show up fee sufficient to cover the cost of take-up in even the lowest subsidy treatment. Thus, cash-on-hand is unlikely to interfere with take-up. Second, the variation in the timing of the reward for follow-through provides a separate test for self-selection, which does not depend on immediate concerns about liquidity. Specifically, because some farmers were aware of the rewards at the time of the take-up decision and some were not, self selection based on expected payoffs should also incorporate the value of the reward. Importantly, the reward is paid after a year, and thus the response to it at the time of take-up should not be contaminated by immediate liquidity constraints. We see no difference in the response of follow-through to the value of the reward based on the timing of the reward announcement (see Appendix Table B.8). This provides further support for the conclusion that the new information after take-up plays a substantial role in the follow-through decision.

Sunk Cost, Information Signaling and Crowding Out Effects

Previous studies of the effect of subsidies on follow-through, via screening on private benefits, have worried about psychological effects associated with the initial price paid for the technology. First, sunk cost effects would cause higher follow-through among adopters who pay for more take-up, because adopters would consider their expenditure at take-up when making their follow-through decision (Ashraf et al. 2010; Berry et al. 2012; Cohen and Dupas 2010). Second, farmers could extract information about the quality of the technology from the NGO's decision to subsidize (Ashraf et al. 2013). If higher subsidies accompany better technologies, then farmers in a higher subsidy condition might have a higher follow-through (tree survival). This works in the opposite direction as the sunk cost effect. Third, if paying farmers to take up the technology crowds out their intrinsic motivation for the technology, we might see higher subsidies leading to lower follow-through (2014). This crowding out effect would work in the same direction as the sunk cost effect. Because we observe no effect of the exogenous variation in take-up subsidies on follow-through, we are able to rule out all three of these explanations.[38] Note that were we to observe a screening effect of the take-up price, we would not be able to distinguish it from these other channels using our design.

2.5 Model, Identification and Estimation

The reduced form results and stylized facts in Section 2.4 provide evidence that is consistence with uncertainty in the opportunity costs of follow-through. However, they do not rule out that, in addition to uncertainty, other sources of heterogeneity in costs may explain part of the apparent disconnect between high take-up rates and low follow-through. This could be the case if, for instance, there is a negative correlation between the privately optimal number of trees and the private net profit derived from them.[39] In addition, the reduced form results do not offer any specific facts about the magnitude of the uncertainty that farmers face and how does it compare to a set up where selection effects are present.

We modify our simple theoretical model described in Section 2.2 to adapt it to our empirical setting and explicitly estimate the distribution of random parameters governing a quasi-profit function (a "reduced form" profit function of sorts). This expanded model allows us to answer the questions above.

[38]That we do not see differences in follow-through across the reward timing conditions provides additional evidence against crowding out of intrinsic motivation associated with selecting in to the program based on expectation of a reward.

[39]A correlation (positive or negative) between the privately optimal number of trees and the level of profit can emerge from the joint distribution of the primitive parameters that govern a profit function (e.g. marginal costs, fixed costs, marginal benefits, etc.). This is somewhat similar to Suri (2011), who finds that low adoption rates of hybrid maize among farmers who seem to have high returns from adoption can be traced to a positive correlation between costs and benefits from adoption using a random coefficients model.

2.5.1 Farmer Net Benefits

We begin with a general characterization of farmer profits at $t = 1$, after new all relevant information is revealed. Then we specify a time line for information and introduce the technology take-up decision under uncertainty (at $t = 0$).

General Profit Function Consider the decision of a farmer to plant and care for trees for their private benefits. We preserve the assumption of risk neutrality made in Section 2.2.[40] Thus, we can represent farmer's utility from trees in terms of a profit function. Consider a general quadratic profit function for a number of trees N given by

$$\Pi(N) = \left[\sum_{t=7}^{\bar{T}} \frac{1}{(1+r)^t} \left(\alpha_0 N - \alpha_1 N^2\right)\right] - \gamma_0 N - \gamma_1 N^2 - \gamma_2 \times \mathbf{1}(N > 0) \tag{2.3}$$

where α_0 and α_1 are private benefit parameters, r is the annual discount factor, $\bar{T}$ is the maximum number of years a tree lives, γ_0 and γ_1 are cost of implementation parameters that are proportional to the number of trees and γ_2 is a fixed cost of implementation. Recall that private benefits in the form of soil fertility start around the 7th year. Under some parameter values,[41]equation (2.3) describes a convex function in the number of trees cultivated. This is consistent with our empirical observation that a number of farmers find it optimal to cultivate a number of trees between zero and 50 (the number of seedlings they receive) in the absence of an external incentive (i.e. when the reward for tree survival is equal to zero). In other words,

[40] As is discussed in Section 2.2, this assumption will have little consequence for our results to the extent that the changes in income produced by our program are small relative to total income. This is indeed the case, as the highest reward from our program is roughly 3.5 percent of average annual income. Although our baseline survey did not collect comprehensive measures of income, we know that the average annual income from the household's main crops has a mean of ZMK 2.4 million (s.d. ZMK 3.4 million), which is 16 times the maximum reward derived from our program and 34 times the average reward. This measure of income is likely to underestimate total income substantially, since it ignores income from minor crops, consumption value from subsistence crops and non-agricultural income. Fink et al. (2014) collected detailed income data from a larger sample from the same District, but without the requirement that they work with Dunavant, find that mean annual household income for smallholder farmers in the District is 4.3 ZMK million (s.d. 5 million).

[41] Equation (2.3) can be rewritten as

$$\Pi(N) = (\tau\alpha_0 - \gamma_0)\, N - (\tau\alpha_1 + \gamma_1)\, N^2 - \gamma_2 \times \mathbf{1}(N > 0)$$

where $\tau = \sum_{t=7}^{\bar{T}} \frac{1}{(1+r)^t}$. This function is concave in the number of trees as long as $\alpha_0 > 0$, $\alpha_1 > 0$, $\gamma_0 > 0$, $\gamma_1 > 0$, $\gamma_2 > 0$, and $\tau\alpha_0 - \gamma_0 \geq 0$.

there the privately optimal number of trees appears to correspond to an interior solution to the private profit maximization problem for many farmers.

The solution to the profit maximization problem defined by (2.3) is:

$$N^* = \begin{cases} \frac{\tau\alpha_0-\gamma_0}{2(\tau\alpha_1+\gamma_1)} & \text{if } \frac{(\tau\alpha_0-\gamma_0)^2}{4(\tau\alpha_1+\gamma_1)} - \gamma_2 > 0 \\ 0 & \text{if } \frac{(\tau\alpha_0-\gamma_0)^2}{4(\tau\alpha_1+\gamma_1)} - \gamma_2 \leq 0 \end{cases}$$

where $\tau = \sum_{t=7}^{\bar{T}} \frac{1}{(1+r)^t}$. Thus, in this version of the model the follow-through decision corresponds to the choice of the optimal number of trees.

Note that, in principle, all parameters in (2.3), $\alpha_0, \alpha_1, \tau, \gamma_0, \gamma_1$, and γ_2, may vary across farmers. Our experimental variation, however, does not allow us to separately identify all potential sources of heterogeneity across farmers. We therefore turn to a quasi-profit function that uses our experimental variation to characterize farmer heterogeneity along two important dimensions of the farmer's profit maximization problem, $\max_N \Pi(N)$: the interior solution and the profit level evaluated at the optimal number of trees.

General Profit Function The same interior and corner solutions conditions delivered by (2.3) are generated by the following quasi-profit function indexed by two random parameters, T_i and F_i:

$$\Pi(N) = N - \frac{1}{2T_i}N^2 - F_i \times \mathbf{1}(N > 0) \tag{2.4}$$

where $T_i = \frac{\tau\alpha_0-\gamma_0}{2(\tau\alpha_1+\gamma_1)}$ and $F_i = \gamma_2 + \left(\frac{(\tau\alpha_0-\gamma_0)-(\tau\alpha_0-\gamma_0)^2}{4(\tau\alpha_1+\gamma_1)}\right)$. This alternative profit function allows for heterogeneity across farmers in the interior solution (which is equal to T_i) as well as parameter F_i, which is a scaling parameter that ensures profits in the quasi-profit function coincide with profits in the generic quadratic function when evaluated at the interior solution. The advantage of (2.4) over (2.3) is that all random parameters are identified out of the variation induced by our experiment. Because both functions have the same value at the optimal solution, (2.4) can be used to evaluate welfare under the more general profit function (2.3).

Introducing a reward for follow-through

The farmer's objective function when faced with an external reward for follow-through (tree survival above the threshold $\bar{N}$) is

$$\Pi(N|F_i, T_i, R_i) = \left(N - \frac{1}{2T_i}N^2\right) + \mathbf{1}(N \geq \bar{N})R_i - \mathbf{1}(N > 0)F_i \tag{2.5}$$

where R_i denotes the randomized reward. As we discuss below, the random variation in R_i will allow us to identify the conditional distribution of F_i given T_i.

Through out our estimation, we assume that survival of trees is deterministic conditional on farmers' costly effort. As we explain in explain in greater detail in the appendix, this assumption is consistent with a model where probability of survival is a convex continuous function of effort, e, up to $\tilde{e}$, where it attains one. Farmers would respond to such probability profile by investing the minimum effort that guarantees survival, $\tilde{e}$, in all trees they choose to plant.[42] We believe this is a reasonable assumption to make in our context, given that we observe a small bunching of tree survival at 35, which is consistent with perfect targeting of the tree threshold that guarantees receiving the reward (see Appendix figure B.3.1).

Alternatively, one could assume that effort choices still affect the probability of survival, but no level of effort can guarantee survival of the tree. In Appendix B.3 we use Monte Carlo simulations to explore how this alternative assumption would affect farmers decisions and observed outcomes. Given the bunching at 35 trees, we assume that the maximum probability of survival is relatively close to one, as smaller probabilities would eliminate the discontinuity in the distribution. The results from the simulations suggest that introducing stochastic survival conditional on effort does not change farmers' behavior relative to the incentives we provide in any meaningful way. Specifically, the responses of take-up and follow-through decisions to the take-up subsidy and threshold rewards remain approximately of the same magnitude.

2.5.2 Dynamics and Take-up Decision

As in the conceptual model, we assume the farmer makes adoption-related decisions in two periods: $t = 0, 1$. The random parameter F_i described in the previous section is divided into

[42] Except, perhaps, on one of them, as is explained in Appendix B.3.

two orthogonal and additive components: F_{0i} and F_{1i}, where F_{0i} is known at all periods and F_{1i} is known at $t = 1$ but not at $t = 0$. In addition, we assume that T_i (which corresponds to the interior solution to the maximization problem) is known to the farmer at all times. Given that the random parameter F_i governs the level of profits (evaluated at the optimal number of trees), this amounts to assuming that there is uncertainty about the net returns to tree cultivation, while the optimal scale of the technology is known to the farmer at all points in time. The advantage of assuming this particular structure of information is that it allows us to nest a model without uncertainty (i.e., $\text{Var}(F_{1i}) = 0$) that could also generate high take-up and low follow-through within our more general model.

Thus, summarizing, there are two sources of static heterogeneity: the scale of the profit component, F_{0i}, and the interior solution, T_i. Through a negative correlation in these two random parameters, the model allows for farmers who take-up under a full price to have higher or lower tree survival outcomes than they would under a subsidized price. This is important, since part of the aim of this paper is to investigate the importance of uncertainty in explaining low follow-through decisions after take-up vs. other alternative explanations that rely on static heterogeneity.

At $t = 0$, the farmer decides whether or not to pay to take-up the technology. At this point in time, the farmer has partial information about her net benefits from the contract. We assume that the farmer knows the the distribution of F_{1i} conditional on F_{0i} and T_i at the time she makes her decision to take-up. More precisely, farmers know that $F_{1i} \sim \text{normal}(0, \sigma_F^2)$, such that $F_i \sim \text{normal}(F_{0i}, \sigma_F^2)$. Thus, the farmer chooses to take-up if

$$\delta \mathbb{E}\left[\max_N \Pi(N|T_i, F_{0i}, F_{1i}, T_i, R_i) \middle| F_{0i}, T_i\right] - c + A_i \geq 0 \tag{2.6}$$

where c is the cost of the seedlings, A_i is the randomly determined subsidy, and δ is the discount factor, which we assume is equal to 0.6.[43]The expectation in (2.6) is taken over the

[43]Like Stange (2012), we note that in the context of stochastic dynamic structural models the discount factor is not separately identified from the scale parameter of future period shocks. We collected survey data on time preferences, however, the survey tool used to elicit the discount rate is very coarse and responses are most informative as a relative ranking, so we do not use them directly in our structural estimation. They do, however, inform our choice of 0.6 for the discount factor in the structural estimation. See Footnote 22 for further discussion of this parameter value.

distribution of F_{1i} conditional on (F_{0i}, T_i).

2.5.3 Identification and Estimation

Identification of the structural model consists of uniquely identifying the joint distribution of unobservables T_i, F_{0i} and F_{1i}. Identification relies on the observed variation in tree survival and take-up decisions across different treatments and a distributional assumption on the random vector (F_{0i}, F_{1i}, T_i).

The identification of the joint distribution of $F_{0i} + F_{1i}$ (or F_i) and T_i can be non-parametrically identified in the subset of the support such that $\bar{N} < T_i < 50$ and $F_i < \frac{1}{2}T_i$. To see this, consider the follow-through decision of the subset of the sample for which

$$\lim_{a \to \mathcal{A}_1} \Pr(\mathbb{E}\left[\max_N \Pi(N|T_i, F_{0i}, F_{1i}, T_i, R_i)\middle| F_{0i}, T_i\right] \geq c - a) = 1,$$

such that there is no selection on take-up. Within this subset of the sample, we can use the variation in R to identify the joint distribution of F_i, T_i as well as the marginal distribution of T_i.

For this group, the probability of cultivating $N^* = n > \bar{N}$ trees when $R = r$ can be written as

$$\Pr(N^* = n; R = r) = \Pr\left(F_i < r + \frac{1}{2}n\middle| T_i = n\right)\Pr(T_i = n) \tag{2.7}$$

Because the left hand side of (2.7) is empirically observable, increments of r trace out the joint distribution of F_i given T_i.

Since non-parametric identification of the joint distribution of F_i and T_i occurs only in the subset of the support such that $\bar{N} < T_i < 50$, additional parametric assumptions are required to fully characterize these distributions. In the empirical estimation we assume that $\ln T_i$ and F_i are jointly normally distributed, thus the marginal distribution of T is log-normal(μ_T, σ_T), while F_i is normally distributed with mean μ_F and variance $\sigma_F^2 = \sigma_{F_0}^2 + \sigma_{F_1}^2$ with the correlation between F_i and T_i given by ρ.

In addition to the additivity and orthogonality assumptions on F_{0i} and F_{1i}, we need to assume that the distribution of F_{1i} is *iid* across farmers in order to separately identify the distribution of F_{0i} and F_{1i}. This new assumption rules out common shocks to farmers. Under these assumptions, the decision to take-up provides independent identification of the known component of F_i, F_{0i}, through the variation in R and A. This allows us to capture the role of information and uncertainty in the take-up decision and to determine the share of private profit that comes from the option to costlessly exit at $t = 1$.

Identification of the distribution of F_{0i} is obtained from the decision to participate, which is characterized by the inequality in (2.6). The left side of (2.6) is a known function of the random variable F_{0i}. Denote this function $h(F_{0i}; r_i)$, so we can rewrite (2.6) as

$$h(F_{0i}; R_i = r_i) \geq c - a_j \tag{2.8}$$

The right side of (2.8) can take one of four known values, $c - a_j$ for $j = 1, ..., 4$, and is randomly determined by the research design. The left hand side of (2.8) is known up to F_{0i} and varies across individuals according to the known cost determinant, r_i. Provided that $h(F_{0i}; r_i)$ is invertible,[44] we can identify the distribution of F_{0i}, from the random variation in a_j and r_i:

$$\Pr\left(F_{0i} \leq h^{-1}(c - a_j, r_i) | \mu_F, \sigma_F^2, \mu_T, \sigma_T, \rho\right) = \Pr(Part_i | A_i = a_j, R_i = r_i) \tag{2.9}$$

Note that in (2.9), parameters μ_F, σ_F^2, μ_T, σ_T, and ρ can be treated as known since they are identified from tree survival as described above. However, the random variation in R_i and A_i allows us to identify μ_F in (2.9) from the participation decision. We use this feature of our research design to explore whether farmers have correct beliefs on $\mathbb{E}(F_i)$ at $t = 0$, and in a model variant we allow for a mean shift in μ_F, which would be consistent with an unexpected common shock to all farmers. This variant of the model partially relaxes the assumption of independence of shocks across farmers. However, it is somewhat limited in that its unexpected nature means it does not affect the option value of the contract. A complete relaxation of this assumption, which would allow for free patterns of known correlation in shocks to farmers' costs, is not possible given the limited time horizon in our data.

Estimation

We estimate the model using simulated maximum likelihood. The log-likelihood function is over observations of the number of planted trees, $N = 0, ..., 50$, and the participation decision,

[44]It can be shown that there exists some $\bar{f}$ such that $h(F_{0i}; r_i)$ is strictly monotonically decreasing on $(-\infty, \bar{f})$.

$P = 0, 1$. The sample includes the 1314 farmers who made a take-up decision. Because there are no trees planted whenever the individual chooses not to participate, the support of this bivariate vector is given by the 52 (P, N) pairs: $(0, 0), (1, 0), (1, 1), (1, 2), ..., (1, 50)$.

We use numerical methods to minimize the negative of the simulated log-likelihood. For each likelihood evaluation, we use 500 draws of (T_i, F_{0i}, F_{1i}). Also within each likelihood evaluation and for each draw of (T_i, F_{0i}, F_{1i}), the expectation in the right hand side of equation (2.6) is numerically computed using 100 draws of (T_i, F_{1i}) conditional on the draw of F_{0i}.[45] Standard errors for the estimated parameters are obtained as the inverse of the inner product of simulated scores.[46]

2.6 Structural Estimates and Simulation Results

In this section, we describe the structural estimates and carry out a number of counterfactual simulations.

2.6.1 Structural Estimates and Model Fit

Table 2.3 shows the point estimates for the main parameters described in Section 2.5.3.[47] Panel A shows the estimates of our baseline model, which assumes that farmers' expectations about the F are correct and consistent over time. Panel B shows the results of allowing for

[45]Simulated methods often result in stepwise objective functions which work poorly with gradient-based numerical optimization algorithms. To facilitate the numerical optimization, we "smooth" the objective function by computing the multilogit formula for each decision over participation and the number of trees. We assume a relatively small variance parameter of the logistic error term: 0.5. However, we experiment with different values for this parameter. We find that smoothing does not significantly affect the point estimates and does improve substantially the curvature of our objective function.

[46]See Appendix B.2 for a more detailed description of our standard error calculation.

[47]There are two remaining parameters that are omitted from Table 2.3 but discussed in Appendix B.2 for the sake of brevity: these are the surprise treatment parameter, α_S, and the monitoring treatment parameter, α_m. Recall that farmers in the surprise reward treatment made a take-up decision before learning their reward. We model this aspect of the design by assuming individuals expect a threshold reward of 0 when their participation decision is made, but incorporate the reward value they draw in their follow-through decision. Because our reduced form results show that individuals in the surprise treatment had higher participation rates than those individuals who drew a reward of zero ZMK, we allowed the surprise reward treatment to have an independent effect on the participation decision. The structural estimates suggest that the boost to participation is equivalent to offering them between 92 and 54 ZMK (in the model with a mean shifter). In all models, we allow the regular visits to collect data on program implementation that were administered to one-fifth of farmers to independently affect the tree survival decision (but not the participation decision). The estimated parameter is -238.40 (s.e. 73.887) in Panel A and -229.53 (s.e. 74.444) in Panel B.

an unexpected common shock to all farmers at $t = 1$ (a mean shifter). As discussed in Section 2.5.3, this alternative model partially relaxes the assumption of independence in shocks across farmers. Because point estimates are somewhat hard to interpret (e.g. the μ_T and σ_T parameters do not correspond to the mean and standard deviation of the log-normally distributed parameter T), we convert the estimated parameters into more easily interpretable outcomes using simulation.[48] The estimated joint distribution of T and F shown in Panel A is such that the mean ex-post privately optimal number of trees is 8.46 (s.d. 14.64), with about 59 percent of farmers choosing to plant no trees.[49]

Table 2.3: *Structural Parameter Estimates*

Parameters in T					Parameters in F			
μ_T	σ_T	ϱ	μ_F	σ_{F0}	σ_{F1}	α_s	α_m	μ_{Fs}
				Panel A. No Mean Shift				
3.539	1.401	0.818	107.58	307.87	211.42	-91.79	-238.40	-
(0.057)	(0.066)	(0.066)	(11.822)	(93.278)	(49.953)	(16.222)	(73.887)	-
				Panel B. Allowing for Mean Shift				
3.579	1.392	0.835	74.48	290.06	193.05	-54.42	-229.53	53.29
(0.071)	(0.075)	(0.073)	(15.47)	(84.622)	(45.427)	(20.47)	(74.444)	(26.761)

Notes: Parameters fitted by simulated maximum likelihood using 1500 draws of the random vector *(F_{0i}, F_{1i}, T_i)*, with smoothing (lambda is 0.5) and tolerance (1e-15). The baseline model (Panel A) restricts the mean of F_i to be the same in both time periods. The mean shift model (Panel B) allows the mean of F_i to differ between the two periods, and the parameter F_shift to capture this difference. The log-likelihood value for the baseline model is 11142.064, while it is 11138.996 for the mean shift model.

This joint distribution also implies that the average ex-post private profit from the optimal number of trees is 108.39 thousand ZMK . However, ex-post private profits vary widely across farmers: their s.d. is 185.47. Importantly, the model estimates that about 39 percent of the variance in ex-post profits is due to new information that emerges after the take-up decision is made.

[48]The point estimates in Table 2.3 can be used to simulate farmer's draws of F and T. These draws can then be used to compute optimal tree cultivating decisions that account for interior and corner solutions to the optimization problem. The optimal solutions can then be plugged back into the profit function to compute maximized profit. The statistics presented below correspond to means and variances from 10,000 simulated draws.

[49]These statistics assume that they can plant a maximum of 50 trees. Although we allow for the distribution of T to be unbounded, we chose to present statistics of the bounded distribution because we fit the econometric model using only this range of outcomes. According to our estimates, about 6 percent of the farmers would choose a private optimum of 50 or more trees.

The variance of shocks is partially identified out of the difference in the variances of expected profits at the time the take-up and ex-post profits at the time the follow-through decision is made and partially identified out of its non-linear effect on the mean level of the expected profits at the time of take-up.[50] The presence of common shocks will generate a tug-of-war between these two sources of identification, as the expected value of the profits will pull the variance of shocks parameter, σ_{F1}, up while the ex-post variance in profits will not reflect the variance of common shocks and thus will pull, σ_{F1}, down. Our mean shifter model helps us explore the importance of the bias in σ_{F1} generated by the presence of common shocks, as it allows for the mean level of profits to be different at the time of take-up and at the time the follow-through decision is made.

The corresponding results from the mean shifter model are presented in Panel B of Table 2.3. The results are consistent with the presence of common shocks, but also suggest the share of variance attributed to them is relatively small: the mean shifter parameter is positive, which reflects the higher mean level of expected profits; while the idiosyncratic shock variance parameter, σ_{F1}, falls moderately from 211.42 to 193.05, which is consistent with a positive bias in our baseline model stemming from the presence of common shocks.[51] The remaining results are qualitatively similar. The mean number of privately optimal trees is 7.67 (s.d. 14.00) and the mean ex post private profit is 90.46 (s.d. 163.73). The share of the variance in ex-post profits attributed to idiosyncratic shocks remains high: 40 percent.

We now turn to the interpretation of the parameters that govern the distribution of known (at the time of take-up) sources of heterogeneity across farmers. We estimate a high positive correlation between F and T, $\rho = 0.81$ ($\rho = 0.83$ in the mean shifter model).[52] Because F enters negatively in the profit function, this translates into a negative correlation between the privately optimal number of trees and the level of profits. This negative correlation generates

[50] Recall that a higher variance in shocks results in a larger option value for the contract.

[51] Unfortunately, we cannot calculate the share of the variance attributed to common shocks using model estimates as we are not explicitly modeling random common shocks with a known distribution at the time of take-up.

[52] F_1 is assumed to be orthogonal to F_0 and T. Thus, the correlation between F and T stems solely from the correlation between F_0 (the known component of F) and T.

low, but strictly positive, follow-through rates (as in low numbers of trees planted) among farmers whose expected value from the contract is high and thus take-up under a high price of the contract. In other words, the static heterogeneity identified by the model contributes to undermining the effect of high prices at take-up on positive (high follow-through) self-selection of farmers.

Because F and T are reduced form parameters (see Subsection 2.5.1), the positive correlation between them could stem from two sources: a positive correlation between farmers' fixed costs and farmers interior solution to the maximization problem (a component of the reduced form random parameter, F, is the fixed cost of the generic quadratic profit function), or a high mean and high variance in the linear term of net marginal returns (the term $\tau\alpha_0 - \gamma_0$ in the reduced form expression for T and F), which enters linearly in T and non-linearly in and F, and in high ranges may generate a negative relationship between them.[53]

Next, we examine model fit by comparing the reduced form treatment effects using simulated outcomes from the estimated model and the observed data. Most of the magnitudes and signs between the treatments and outcomes are well matched by our model estimates. Panel A of Table 2.2 shows results with the observed data, while Panels B and C show the corresponding simulations using the estimates from Table 2.3. Columns 1-4 estimate the effect of a thousand ZMK increase in the subsidy on take-up and follow-through outcomes. Columns 5-8 repeat the regressions with the reward (in '000 ZMK) on the righthand side. The effect of the subsidy and the reward on take-up (columns 1 and 5) show very similar coefficients for the observed and simulated datasets. The estimates for whether the farmer reached the 35-tree threshold (columns 2 and 4, conditional on take-up) are also similar across the observed and simulated data, with slightly higher estimated effects of the reward on reaching the threshold in the simulated data.

There are, however, some discrepancies between what our model predicts and the observed data. Columns 3 and 7 show the effect of the take-up subsidy and reward on the number of surviving trees, excluding zeros (which are shown in columns 4 and 8). Interestingly, the sign

[53] The reduced form expressions for T and F as a function of the generic quadratic profit function parameters correspond to the first and second brackets in (2.4).

of the coefficient on the subsidy is different between the observed and simulated data, though standard errors are relatively large. The effect of the reward on the number of trees (column 7) is larger in the simulated data, consistent with the effects on reaching the 35-tree threshold. Finally, we see evidence that the take-up subsidy selects for farmers more likely to keep zero trees alive (column 4) in the simulated but not in the observed data, indicating some selection effect on abandoning the technology altogether in the simulated data only. The effect of the reward on zero tree survival is the same in the observed and simulated data.

Using simulations, we explored whether some of the discrepancies between our estimated model and the data stem from the assumption of deterministic survival of the trees conditional on effort. To do this, we take the parameter estimates as given (we use the estimates from the mean shifter model) and incorporate a binomial and a beta binomial distribution to the survival outcomes of trees under a couple of different parameter values.[54] We find little to no improvement when stochasticity is introduced into the tree survival outcomes: The simulations that incorporate stochastic survival of trees show no improvements on matching the observed relationships between the take-up subsidy and the positive number of trees or the take-up subsidy and choosing zero trees. The beta binomial distribution shows little improvement at matching the relationship between the take-up subsidy and the positive number of trees as well as the choice of zero trees; although it does so at the expense of worsening the fit for the relationship between the reward and take-up.The results from these simulations are further discussed in Appendix B.3.

2.6.2 The Effect of Uncertainty on Farmer Profits and Program Outcomes

We use estimates from the structural model to perform counterfactual simulations of farmer profits and program outcomes (take-up and tree survival) under different levels of uncertainty. For these analyses, we use the results from the model that allows a more flexible

[54]We chose not to reestimate the structural parameters under the alternative assumption for tree survival, given that our simulations suggested that stochastic survival of trees was unlikely to improve the fit of the model, and that estimation of the model is very time consuming.

specification of F_i with the mean-shifter parameter, such that both the mean and the standard deviation of F_i may vary over time (Panel B in Table 2.3). This version of the model offers the best approximation of our observed take-up and tree survival results. In the simulated results below, we set the mean shifter for $t = 1$, μ_{F_s}, equal to zero, so that we can equate the expected benefits with the true average discounted benefits from the program in the welfare evaluations.[55] Results from using our baseline model estimates instead are qualitatively similar.

Farmer Profits

We begin by examining the effect of uncertainty on the average per-farmer expected private profit (right hand side of inequality 2.6) implied by the empirical model.

In order to do so, we simulate the value of the expected profit for each farmer at different values of σ_{F1}. The relationship between the mean expected profit and σ_{F1} corresponds to the solid black curves in Figure 2.3. Panel A shows the relationship for a reward of zero and Panel B for the largest reward offered (150 thousand Kwacha). Both are shown at a full subsidy, so that take-up is 100 percent. The first takeaway from this figure is that uncertainty increases the expected profit, for both low and high rewards. This result is analogous to the theoretical result discussed in Proposition 4: the option value from the contract increases with uncertainty and thus drives a positive relationship between the expected profit and uncertainty. The option value, as defined in 2.2, is shown by the dashed lines in Figure 2.3 for different values of the reward. The option value is always non-negative (as farmers could only do better with more information, but never worse), and is also the only component of the expected private profit that varies with uncertainty.

[55] This treatment of the mean-shifter parameter is consistent with the common-shock interpretation of this parameter.

Figure 2.3: *Farmer Expected Profit as a Function of Uncertainty*

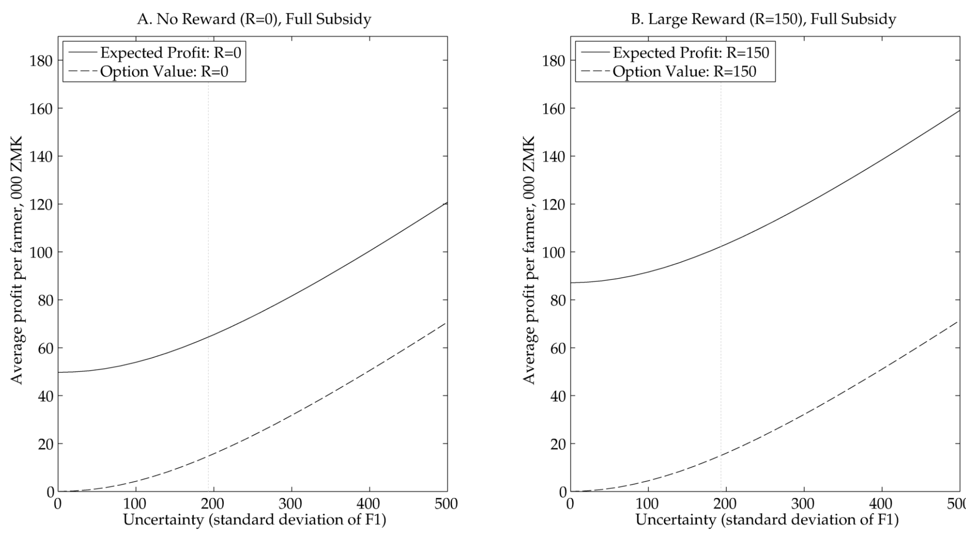

This figure shows a counterfactual simulation of farmers' mean expected profit where we vary the level of uncertainty (the standard deviation of F_1). For the simulations, we use the estimated parameters from Panel B of 2.3, except for σ_{F_1}, which we vary along the horizontal axis. The mean per-farmer profit is shown in a solid black line for low (Panel A, $R = 0$) and high (Panel B, $R = 150$) reward values. Profits are expressed in '000 ZMK. The dashed lines show the mean option value for the two different reward levels. We define the option value as the difference between expected profits under the possibility of choosing the number of trees to cultivate after F_1 is realized and under the static contract (where the choice on the number of trees to cultivate has to be made ex-ante).

Recall that this result stems from the asymmetric response of the expected profit to positive and negative shocks: If the realization of the random component of the profit from follow-through is high (low F_1), then the farmer responds by securing that profit through positive tree tree survival. However, if the realization of the random component of the profit drives it below zero, the farmer will respond by not cultivating any trees at all (effectively bounding the profit realizations at zero). This optimizing behavior turns the high variance of the shocks into an asset of the contract, which in turn results in higher take-up.

The positive relationship between expected private profit and uncertainty has an important implication for take-up decisions: more farmers are ex ante attracted to the contract under higher uncertainty, despite the fact that its ex post expected value remains unchanged. Note that a high enough level of uncertainty may result in an expected profit that exceeds the take-up cost even under low subsidies. Hence, the ability of subsidies to "tease out" those who will be more likely to comply with the tree survival goal decreases with uncertainty. This can be observed more directly when looking at take-up and follow-through outcomes as a function of uncertainty, which we turn to next.

Program Outcomes

Figure 2.4 plots average take-up at low and high subsidies (dashed and solid lines) and low (Panel A) and high (Panel B) rewards as a function of uncertainty. Panels C and D show the share of individuals who reach the 35-tree goal conditional on take-up for the same combinations of subsidies and rewards. Figure 2.5 shows the effects of uncertainty on the average number of trees for different values of the subsidy (A) and reward (R). Panels A and B show the average tree survival, unconditional on take-up (non-participants have zero surviving trees). Panels C and D show average tree survival conditional on take up. Because take-up is 100 percent with a full subsidy (A=12), the solid lines in Panels C and D coincide with the solid lines in the Panels A and B, respectively.

Figure 2.4: *Take-up and Threshold Outcomes as a Function of Uncertainty*

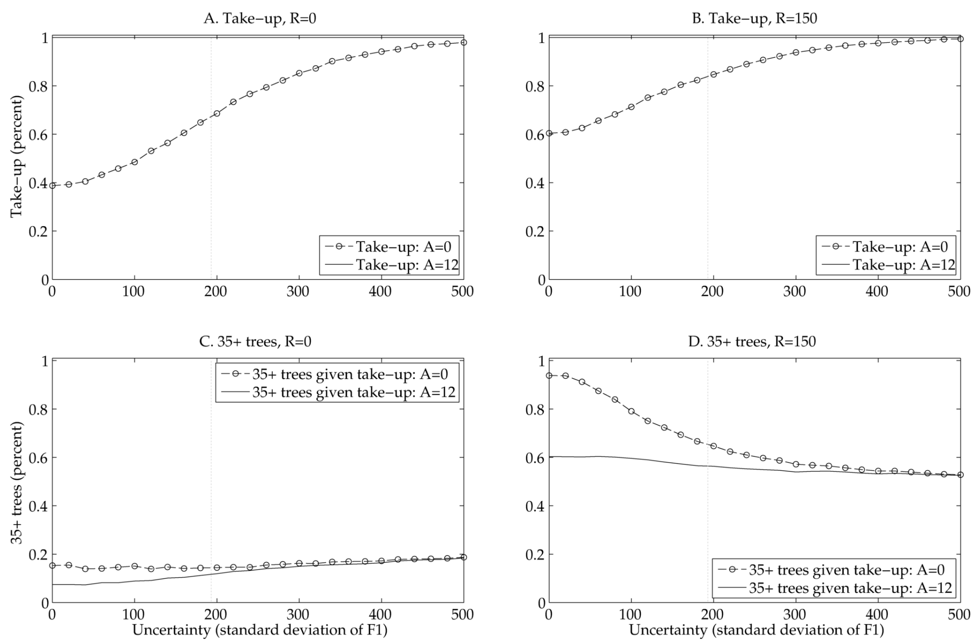

This figure shows a counterfactual simulation of farmers' average take-up and 35-tree reward threshold outcome as we vary the level of uncertainty (the standard deviation of F_1). For the simulations, we use the estimated parameters from Panel B of Table 2.3, except for σ_{F_1}, which we vary along the horizontal axis. Take-up and threshold (trees $\geq$ 35) outcomes are shown for different combinations of the reward value and take-up subsidy.

Figure 2.5: *Tree Survival as a Function of Uncertainty*

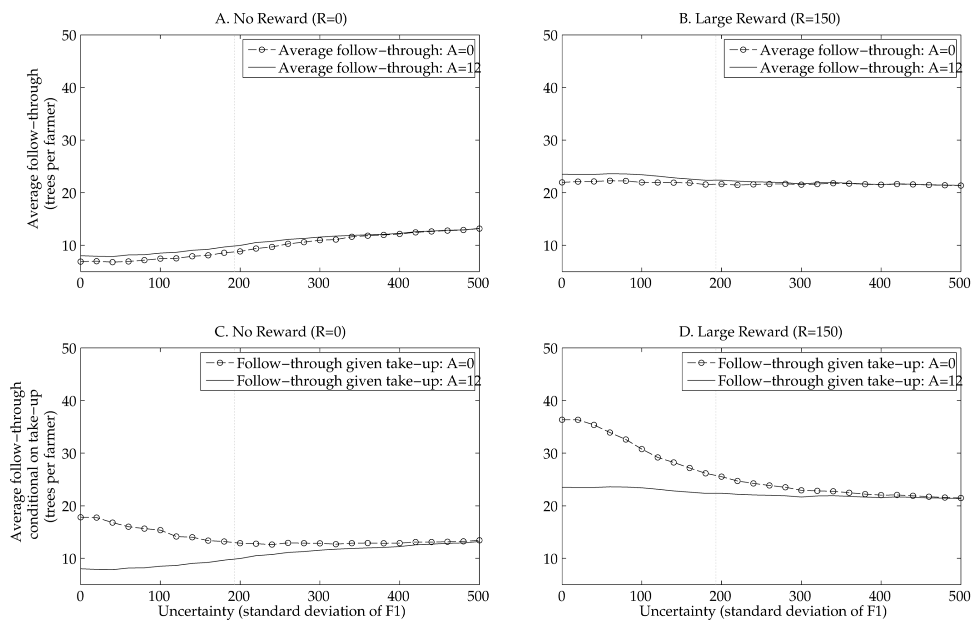

This figure shows a counterfactual simulation of farmers' mean number of surviving trees as we vary the level of uncertainty (the standard deviation of F_1). For the simulations, we use the estimated parameters from Panel B of Table 2.3, except for σ_{F_1}, which we vary along the horizontal axis. The mean number of surviving trees is shown for different combinations of the external threshold reward and the subsidy for inputs. The top panels show per-farmer tree survival for all farmers (those who didn't take up have zero trees); the bottom panels show tree survival conditional on take-up.

First note that the average number of trees conditional on take-up is boosted by a high take-up cost (low A) when uncertainty is low (the boost corresponds to the difference between the two lines in Panels C and D). This boost is what we refer to as the selection effect of price. This result from the model, the existence of a selection effect for low levels of uncertainty, is analogous to Propositions 1 in our conceptual model (Section 2.2). For the level of uncertainty identified from our data, a value of σ_{F1} of 195, the boost from high take-up cost to tree survival is modest: less than 5 trees; and it continues falling as uncertainty further increases. The reduction in the selection effect as uncertainty increases is analogous to Proposition 3 in our conceptual model.

Second, uncertainty lowers follow-through conditional on take-up whenever the price for take-up is high. This is illustrated by the downward slope of the dashed line with circles in Panels C and D in figure 2.5. This effect is driven by the selection effect of the positive price fading away as uncertainty increases: as can be seen from Panels A and B in figure 2.4, take-up converges to 100 percent as uncertainty increases, making the pool of takers ever more similar to the overall pool. This is broadly consistent with Proposition 2 from our conceptual model, although the analogous result in this richer model is qualified by an additional competing effect: the variance of shocks can have a positive or negative effect in the number of farmers that choose zero trees depending on whether the sign of this corner solution threshold is to the right or to the left of the mean of F. We call this effect *the corner solution effect.* Whenever the reward for the 35 tree goal is at 150 ZMK, the threshold is positive. Whenever the reward for the 35 tree goal is zero, the threshold is negative. Thus, when the reward is zero, we see that the unconditional average number of trees increases with uncertainty (both lines in Panel A and the solid line in Panel C increase with σ_{F_1}). Whereas the opposite is true when the reward for trees is 150 ZMK (both lines in Panel A and the solid line in Panel D fall with σ_{F_1}).[56]

[56] This effect holds whenever the distribution of shocks is continuous and symmetric around the mean. To see this, denote the standard normal cdf as $\Phi(.)$ and the threshold along the support of F above which the private profit associated with the interior solution is negative as $\tilde{F}$. The probability of choosing to plant zero trees takes the form $1 - \Phi\left(\frac{\tilde{F}-\mu_F}{\sigma_{F_0}+\sigma_{F_1}}\right)$. Note that the derivative of this probability with respect to σ_{F_1} has the same sign as the numerator of the argument of $\Phi(.)$. When $R = 0$, $\tilde{F} = \frac{1}{2}T_i$. Given that μ_F is above 100 and the mode of T_i is 8.9, the numerator tends to be negative. However, when $R = 150$ and $N^* \geq 35$, $\tilde{F} = N^* - \frac{1}{2T_i}N^{*2} + 150$; and, thus, the numerator tends to be positive.

Thus, in this richer model, the effect of uncertainty on follow-through conditional on take-up carries both effects: the selection effect, which lowers follow-through; and the corner solution effect, which has an ambiguous effect on follow-through. In our context, the selection effect dominates the corner solution effect, and thus follow-through falls with uncertainty conditional on take-up.

The fact that average follow-through conditional on take-up falls with uncertainty has important implications for technologies whose benefits kick in above a certain "usage" level: under low levels of uncertainty about implementation costs, a high cost of take-up will help select those individuals who are likely to engage in a more intensive usage of the technology.

The third observation is that boost to take-up effect from increasing the subsidy may trump some of the selection effects. This is most clearly seen by comparing unconditional tree survival (Panels A and B of Figure 2.5) with tree survival conditional on take-up, which excludes the take-up effect (Panels C and D). The boost associated with the selection effect observed at low levels of uncertainty in Panels C and D is more than offset by the take-up effect: many more farmers take-up when subsidies are high (see Panels A and B in Figure 2.4). The two counteracting effects lead to similar average tree survival across subsidy levels, unconditional on take-up (Panels A and B). Hence, for technologies whose benefits kick in with total follow-through (whether or not follow-through is spread about few or many adopters), subsidies may increase the benefits of adoption. Note, however, that uncertainty removes the take-up advantage of high subsidies, as take-up increases with uncertainty for all subsidy levels (e.g. doubling uncertainty relative to the level observed in our setting increases take-up from 85 to 97.5 percent at the average take-up subsidy in our setting).

The last observation from both Figures 2.5 and 2.4 is that under high uncertainty, a reward conditional on follow-through is more effective at inducing higher levels of follow-through than a subsidy. This can be appreciated when comparing the lines with higher R with the lines with lower . We also note that the combination of low take-up subsidy and high threshold reward works best to maximize the share of individuals whose follow-through falls above the 35-tree threshold conditional on take-up, but this combination of incentives works best under low uncertainty (Panel D of figure 2.4). Even with this combination of incentives, uncertainty can bring down the number of farmers that reach the 35-tree threshold. Taking the level of uncertainty we identify from our data as a baseline, a standard deviation in the unknown part of the cost (σ_{F_1}) that is twice as large would bring down the share of farmers that attain 35 trees or more from 65.5 to 54.7 percent. Conversely, halving the standard deviation would increase the share of individuals who reach 35 trees or more to 80 percent (Panel D of figure 2.4).[57]

2.7 Discussion and Interpretation

Our model and results are consistent with several interpretations of intertemporal adoption decisions, and the nature of the uncertainty that farmers face. Our default interpretation is one of idiosyncratic and common shocks to the opportunity cost of follow-through with the technology, with idiosyncratic shocks playing a relatively more important role. However, results would look similar if instead farmers learn about the technology between the time of the take-up and follow-through decisions. A third explanation for the intertemporal decisions we observe, which would violate our identifying assumptions, is that farmers have time inconsistent preferences. We describe supplementary evidence on each of these explanations in this section.

2.7.1 Common vs. Idiosyncratic Shocks

In our surveys and in other developing country micro data, farmers frequently refer to shocks as a primary determinant of agricultural outcomes. The first version of our model treats the ex post variance in profits across farmers as the sum of the variances of the known and unknown components of the private profit, and therefore allows for idiosyncratic but not common shocks. This restriction on the distribution of the information revealed at time 1 is relaxed in our second version of the model, where we allow for an unexpected mean shift on the distribution of private profit at time 1. The data indicate that the mean shift, which could

[57] We can also perform these simulations with the randomized treatment levels we used in our experiment. Relative to the level of uncertainty we identify from our data, which results in 49.6 percent of farmers reaching the 35-tree goal, doubling uncertainty lowers this share of farmers to 39.96 percent, while halving it increases follow-through to 67.12 percent, conditional on take-up. In the absence of uncertainty ($\sigma_{F1} = 0$), 88.7 percent of farmers who take-up reach the follow-through threshold.

be interpreted as an unexpected common transitory shock or as homogenous learning (see below), is consistent with a common increase in costs after take-up. The estimated mean shift is small in magnitude relative to the variance (see Table 2.3). We interpret this estimate as evidence of common shocks but of limited importance.

The variation in our data does not allow us to identify other types of common shocks: namely, shocks whose distribution is known at time 0, or that are correlated across subsets of farmers (as opposed to all). Therefore, to assess the validity of our identifying assumptions, we consider two additional sources of information about the importance of common and idiosyncratic shocks: household self-reports from the baseline and endline data and the existing literature on agricultural productivity in rural Sub-Saharan Africa.

First, two of the most frequently discussed common shocks are weather and prices. In our setting, the primary source of price risk is from cotton, which experienced a very negative shock in the year of study. At baseline, 97 percent of respondents who forecasted a minimum cotton price for the coming year that exceeded the realized price. Close to 80 percent of the households in our study grew cotton in the contract year, and were therefore affected by the price shock. To the extent that the negative shock to cotton prices was very unlikely from the farmer's stand point, the mean shifter model is likely to capture its effects on farmers' decisions and to reduce the bias in our estimates of the variance of idiosyncratic shocks. Unlike cotton prices, rainfall was fairly typical during the year of the contract, according to local agronomists. Note, however, that to the extent that there are common shocks (such as crop prices and rainfall) whose variance is known and thus are incorporated in the take-up decision, our estimates for the variance of idiosyncratic shocks will continue to be biased upwards even in the mean shifter model.

According to our baseline data, idiosyncratic shocks seem to be an important concern for farmers. When participants were asked about the greatest challenges facing their households, respondents describe issues unrelated to correlated shocks. Two-thirds of respondents list health problems as their greatest challenge. At endline, almost 50 percent of households report losing cattle or livestock to death or theft during the past year and 10 percent of households reporting the death or marriage of a working age member.

Second, a growing literature documents large negative effects of household illness on labor supply and agricultural productivity in the region of study (e.g., Fink, 2012) and on consumption more broadly (for example, Gertler, 2002). These findings are consistent with a larger literature that documents, in most cases, a disproportionate share of income risk from idiosyncratic factors in rural developing country settings (summarized in Dercon, 2002). Thus, we conclude that while common shocks are potentially an important source of uncertainty in agriculture, the fact that our model focuses on idiosyncratic shocks appears reasonable in this context.

2.7.2 Learning

As discussed throughout the paper, the theoretical and empirical models are similar, from an ex ante perspective, if the information that arrives between $t = 0$ and $t = 1$ is in the form of learning or in the form of a shock to opportunity cost, provided that both types of new information are independent across farmers. In addition, to the extent that non-idiosyncratic learning shocks are common across all farmers and unexpected, the mean shifter model could account for learning. Interpreted this way, the positive mean shift estimate is consistent with farmers systematically underestimating the costs of follow-through at the time of take-up. That said, the estimated magnitude of the mean shift parameter is small, suggesting relatively little systematic updating across all farmers. Learning shocks that are independent across farmers and whose distribution is known by the farmers ex ante would show up in the variance of F_1. Thus, we cannot distinguish between permanent (learning) and transitory shocks to the opportunity costs of planting the trees.

The potential for learning about the value of the technology during the first year of tree production is somewhat limited given that tree survival benefits are not realized until several years after the end of the study, and planting and caring for trees resembles activities that farmers undertake regularly. Thus, if learning occurs, it is likely related to the opportunity costs of cultivating the trees, rather than the benefits. We use survey data to explore the extent to which knowledge of the technology changes over the course of the year-long contract.

Our most reliable survey measure of knowledge about the cost of cultivating the trees is the number of risks to tree survival that the farmer is able to list. Out of four major risks to the trees (see Section 2.3.3), farmers list an average of 1.64 risks at baseline and 1.75 at endline one year later. Farmers who took up the contract show a larger increase in knowledge than those who did not take-up. This is driven both by actual increases in knowledge by some adopters and high rates of forgetting by non-adopters. Appendix Table B.7 shows the correlations between tree survival outcomes and survey measures of knowledge and learning. Conditional on take-up, tree survival is higher among those who knew more about the technology at baseline, measured either by the number of risks to tree survival that they were able to name (column 1) or by their prior adoption of the tree (column 2). We proxy learning about the technology by changes in the number of risks listed at endline versus baseline. Farmers who increased the number of risks listed actually cultivated more trees, which suggests that learning about additional costs does not deter follow-through (column 1). While having a farmer in the group who had adopted the technology before the program, and may be a source of information about tree cultivation, is positively correlated with additional trees, the relationship is imprecisely estimated. These correlations should not be interpreted as causal, but do suggest that a modest amount of the information that arrives after take-up may be in the form of learning about the technology. Including all four knowledge and learning measures in the regression (column 4) increases the R-squared from 0.0951 to 0.1460.

2.7.3 Procrastination

An alternative explanation for the observed pattern of high take-up and low compliance is procrastination or hyperbolic time preferences. Sustained effort choices are frequently associated with time inconsistent behavior, in which the individuals initially takes up, intending to follow-through. But when the time comes to act, costs loom larger (or benefits smaller) than anticipated at the time of initial take-up decision. Mahajan and Tarozzi (2011) document time inconsistent technology adoption for insecticide treated bednets in India, with low rates of net re-treatment.

We examine evidence for procrastination or hyperbolic time preferences by constructing two measures of procrastination from the survey data. The first of these relies on a series of baseline survey questions about the respondent's tendency to spend money quickly or delay purchases or actions. These questions are combined into a binary measure of procrastination, which is likely to miss naive procrastinators.[58] As part of the endline survey, a series of questions about procrastination in other activities, including paying school fees, purchasing agricultural inputs and milling maize were added to the survey. These are combined into a second binary measure, which is more likely to capture naive procrastinators.

We begin by examining whether these measures of procrastination are correlated with contract take-up or tree survival conditional on take-up, controlling for other characteristics. They are not. We next investigate the insight that farmers prone to procrastination may be differentially sensitive to a contract structure that requires them to pay more upfront for inputs if the potential rewards arrive only after a year of effort. We regress take-up on an interaction of each of the two procrastination measures and the subsidy level. For the self-described procrastinators, there is a weakly greater likelihood of take-up at higher subsidy levels. However, the interaction is insignificant with the measure more likely to capture naive procrastinators. These results are summarized in Appendix Table B.3, and suggest a relatively minor role for procrastination in driving take-up or follow-through outcomes.

2.8 Conclusion

This paper shows that uncertainty can play an important role in the adoption of technologies that require costly effort over time. The literature on technology adoption so far has emphasized liquidity constraints and learning as potential reasons why take-up subsidies need not lead to too much take-up by adopters with a low likelihood of following through with the technology. We provide a broad framework for adoption decisions that allows for time-fixed and time-varying heterogeneity as well as multiple dimensions of static heterogeneity across potential

[58]Specifically, the baseline questions prompted (1) If I get money, I tend to spend it too quickly, (2) I often change my mind and do not follow-through with my original intention and (3) I tend to postpone activities until later. All responses were on a four point Likert item from strongly disagree to strongly agree. A binary measure was constructed to equal one if the respondent agreed or strongly agreed with any of the statements.

adopters. Using a conceptual model, we show that uncertainty in the opportunity cost of adoption can increase take-up rates at the cost of reducing average follow-through rates. Our conceptual model allows for unknown but permanent components of opportunity costs of adoption, such as learning, but may also include transient shocks such as illness in the family. We show that the presence of uncertainty at the time of take-up may explain why high prices are unable to screen for high-follow through types at the time of take-up.

Findings from a field experiment show evidence that is consistent with the presence of uncertainty in agricultural technology adoption. In Zambia, farmers decide whether to take-up a nitrogen fixing tree under considerable uncertainty about the benefits and costs of following through to keep the tree alive. The study introduces exogenous variation into the initial subsidy for take-up and random variation in the reward associated with following through for the first year. We find that, although farmers are responsive to both subsidies for take-up and rewards for follow-through, lower subsidies (higher prices) do not help screen individuals who are more likely to follow-through. We also find that 35 percent of farmers who paid a positive price for take-up end up with no trees after a year. These observations are consistent with the arrival of new information between the take-up and follow-through decisions.

The experimental variation is used to identify a structural model of intertemporal decision making under uncertainty, which explains our field results and quantifies the uncertainty that the farmers face at the time of take-up. The structural model also helps distinguish between static and time-varying sources of heterogeneity that may explain the absence of screening effects of prices. We find that, in our setup, both static and time-varying heterogeneity reduce the screening effect of the take-up price. Reducing uncertainty in half would increase the number of farmers that reach the tree survival threshold by 18 percentage points (36 percent).

Like all empirical case studies, our data are specific to our setting. However, the combination of the experimental data with a structural model allows us to simulate adoption outcomes under different levels of uncertainty. Our results are consistent with several different interpretations of what changes for farmers between the time of the take-up and the follow-through decisions. First, farmers may experience shocks, such as an illness in the household or the arrival of agricultural pests, that affect the opportunity cost of following through with the technology.

Second, farmers may acquire additional information about the net costs of tree cultivation after take-up through learning by doing or learning from peers. Finally, farmers may exhibit time inconsistency and choose to procrastinate after take-up. From an ex ante modeling perspective, our general framework is agnostic about which explanation is correct. Though we cannot distinguish among them using our experimental design, supplementary evidence suggests that idiosyncratic shocks are important in our setting and that learning opportunities, while present, are minimal.

Our study is an example of how experimental variation can be used to identify dynamic structural models. The use of experimental variation in treatments at two different points in time offers an alternative to a panel data structure, since statistically independent samples are exposed to different treatment combinations. To our knowledge, this is the first paper to introduce experimental variation in order to satisfy the exclusion restrictions needed for sequential identification.

From a policy standpoint, uncertainty has the effect of lowering adoption outcomes per dollar of subsidy invested, while increasing the expected private profits to the adopter, because the downside risk of take-up is bounded at zero. To the extent that subsidies rely on public funds, an increase in uncertainty represents an ex ante transfer from the public to the private domain, driven entirely by adopter's ability to re-optimize follow-through once new information becomes available. While stronger contracts that would force adopters to follow-through once they take-up a subsidized technology would address the problem of high take-up (and subsidization) coupled with low follow-through, they would do so at a clear cost to the adopter. Future research to explore more innovative solutions to encouraging both take-up and follow-through in the presence of uncertainty offers a promising direction for both environmental and development policies. For example, cheaper monitoring solutions that facilitate subsidies (rewards) for follow-through outcomes are shown to have positive effects on both take-up and follow-through in our setting.

Chapter 3

The Effects of Mineral Profits on Districts in South Africa

3.1 Introduction

The role of mineral extraction in developing countries is highly contested. On the one hand, many local economies in developing countries rely heavily on mining, the process of extracting valuable minerals from rock. However, mineral wealth has also been associated with negative consequences, including exacerbating inequality and reducing economic growth.

South Africa provides an important context in which to study mineral resources, given its diverse and highly valuable stocks. Over the past century, various minerals - including diamonds, gold, coal, and most recently platinum - have made immense contributions to gross domestic product (GDP). Although the role of mining has declined with time, as of 2008 mineral resources still represented nearly ten percent of GDP in direct contributions and over forty percent of exports.[1] While mineral wealth is essential to the South African economy, robust quantitive evidence regarding the effect of minerals on local households and on specific resource-rich regions is limited (Stokke, 2008).[2]

[1]The percentage of GDP and the percentage of exports declined since the 1970s but remains high, at 9.5 percent of GDP and 41 percent of total exports (USGS, 2010; South Africa Chamber of Mines, 2012). However, the largest fraction of mineral wealth is derived from coal, which is excluded from this study.

[2]A notable exception is Stokke (2008), which uses a dynamic general equilibrium model to evaluate growth

Simple trade models predict that an increase in natural resource prices should increase wealth within the region that produces those goods (Ricardo, 1817). However, many historical approaches that evaluate the importance of mineral wealth to economic growth find a surprising contradictory pattern. Specifically, using cross-country econometric approaches that control for a variety of national characteristics, Auty (1993, 1998) and Sachs and Warner (1995) first documented that many countries with natural resource endowments fail to perform better than countries without resource endowments. This phenomenon, dubbed the "natural resource curse", motivated a sizable empirical literature testing this relationship and proposing several channels. The results, discussed in greater detail below, suggest that the resource curse may not be generalizable, instead depending on factors such as external conditions and outcomes evaluated.[3]

The inconclusive evidence, and the potential omitted variable bias associated with cross-country analysis, make recent research approaches focusing on within country outcomes particularly appealing (e.g., Aragon and Rud, 2011; Dube and Vargas, 2013; Asher and Novostad, 2014; Allcott and Keniston, 2014). South Africa, for example, is often considered resource-cursed given that the country as a whole has experienced slower growth than both its middle-income peers and other Sub-Saharan African countries (Elbra, 2013). However, it is difficult to disentangle slow growth as a result of a resource curse from other factors, most notably the fall of the Apartheid regime and HIV/AIDs. For this reason, evaluating how changes in local mineral wealth affects economic outcomes may provide better evidence of the role that mineral resource wealth has on economic outcomes. This type of analysis is also useful given that it examines a price shock to a concentrated sector, but measures effects on average district-level employment. Thus, it provides us some evidence of whether average households benefit from the rents to private firms.

This paper explores how within-country variation in mineral wealth affects district employment and poverty outcomes. Using U.S. Geological Survey data on mineral deposits

after positive shocks to gold prices in the 1970s. This paper is discussed in the literature review.

[3]See Frankel (2010) or van Ploeg (2011) for recent summaries of the empirical economic literature relating to the resource curse.

and historical world prices, I construct a plausibly exogenous variable reflecting district-level aggregate mineral wealth, and I use this variation to evaluate economic outcomes over three Census rounds (1991, 1996, 2001). The analysis demonstrates that positive shocks to aggregate mineral wealth generate higher overall employment, driven by increased mining employment. This increase in mining is accompanied by reductions in agricultural employment and slight reductions in manufacturing employment, consistent with the pattern induced by a resource curse. However, in the short run, districts appear to be better off, with a smaller fraction living below the income poverty line. Individual level data also matches this pattern, with greater average hours worked and higher salaries in districts experiencing a mineral price shock.

At the local level and in the short run, this pattern suggests households are better off as the price of their mineral stocks increases. However, the reduction in agricultural employment could reflect a trade-off that some consider a symptom of de-industrialization, whereby regions experience growth in the mining sector, at the expense of growth in agricultural, manufacturing and industrial sectors (Neary and Corden, 1982). While this paper lacks sufficient data to evaluate longer term outcomes, this shift could help explain why South Africa has remained tied to mineral production and not sufficiently diversified into more skilled sectors. This phenomenon is a “curse” some economists fear binds South Africa to poverty (e.g., Hausmann and Klinger, 2008).

3.2 Background

In Section 3.2.1, I describe the minerals and mining industry within South Africa. In Section 3.2.2, I describe the economic literature on the link between mineral wealth and local economic outcomes, with a focus on developing country findings.

3.2.1 History of Mining in South Africa

Mineral production was essential to the development of the modern South African economy, beginning with the discovery of gold and diamonds in the 1800s. However, over time the contribution of these deposits has fallen and the mineral economy has become more diversified.[4] At its peak in 1970, mining contributed between 13-20 percent to total GDP, depending on the source (South Africa Chamber of Mines, 2014). However, the contribution of mining to the economy has since gradually declined, though as of 2008 it still contributed more than five percent directly to GDP.

The reasons for this decline are varied: deeper deposits have become harder to access, economic demand for various minerals has waxed and waned, and other sectors have changed in their output. Similarly, mining's share of employment has fallen over time, from around 14 percent of the formal sector in 1970 to just under 9 percent in 1998 (Blignaut and Hassan, 2002). In sum, its percentage of GDP each year varies depending on the world economy, political situation, physical production and other factors. Despite this, mineral wealth's important role in the economy remains, both during the period of this analysis (1991-2001) and today.

Historically, the minerals industry in South Africa was predominantly private, until some state-controlled mines developed in the late 2000s.[5] Prior to the 1980s, the vast majority of mineral rents went directly to mining companies. During the 1980s, the ability of the government to accrue rents increased, peaking at around thirty percent in the late 1980s before declining again in the 1990s. The government revenues, which were collected through various royalties and taxes, formally went toward the general national budget rather than toward any specific region or department, and there is little evidence that they were earmarked to any environmental or institutional cause (Blignaut and Hassan, 2002). As discussed below, this has implications for the interpretation of the results. In particular, shocks to mineral wealth may be generate rents for private firms, and it is ex-ante ambiguous how these rents will be distributed among private firms, governments, and local households.

---

[4]Other minerals, particularly in the platinum-metals group (PGMs), have experienced greater growth and made up the largest portion of total export value as of 2008 (USGS, 2010).

[5]The exception to this is that the government owned Eskom during this period, the national electric power utility and Southern Oil Exploration Co. However, I exclude energy-related minerals given concerns that many non-market forces affect their prices, especially in a government monopoly. For example, about seventy percent of domestic coal production as of 1998 went to internal power generation (Coakley, 1998).

3.2.2 Background Literature

Standard trade models suggest that resource abundance should improve economic outcomes within the endowed region. Yet as described in the introduction, a sizable theoretical and empirical literature, starting with Auty (1993) and Sachs and Warner (1995), suggests the possibility of a "natural resource curse" whereby mineral rich countries actually do worse than their non-mineral rich peers. As van der Ploeg (2011) discusses in a survey of the literature, empirical evidence for a resource curse pattern is mixed and depends on factors including sample time period, country, definitions of explanatory variables, outcome variables, and other factors. Below I discuss some of the theories related to the "resource curse" that scholars believe are most likely to apply in South Africa.

First, a popular macroeconomic rationale for a resource curse is the phenomenon known as "Dutch Disease." In this model, because mineral profitability depends on global prices, countries experience booms and busts that negatively affect growth. As local mines experience booms, this generates an influx of foreign currency, which leads to currency appreciation and makes every other industry within the country relatively more expensive. As a result, non-resource businesses paradoxically become less competitive. Economic shifts away from traditional manufacturing and agriculture, and toward mining industries, may generate lower longer-term growth due to the volatility of mining relative to more stable sectors (e.g., Deaton and Miller, 1995; Wick and Bulte, 2009).

A second often cited reason for a resource curse is a model whereby resource wealth leads to overspending on low-return investments. Specifically, Bevan, Collier and Gunning (1991) evaluate commodity booms in 23 developing countries and find that governments that experience mineral windfalls feel pressure to overspend, even if they try to save and invest windfalls. Because on average, these investments are lower value and lower return, and because higher government spending is difficult to reverse, governments become stuck in poor outcomes. This model tends to arise regardless of whether the windfall initially goes to the government or private sector, a fact further discussed in Deaton and Miller (1995). The only known empirical analysis of the resource curse in South Africa, a dynamic general equilibrium model by Stokke (2008), finds precisely this: "I apply the model to South Africa and analyze the macroeconomic impact of increases in gold prices during the 1970s. Political pressure for rapid domestic spending following a surge in resource rents tends to generate myopic government behavior with immediate expansion of government consumption."

Third, a group of newer papers use within-country variation in mineral resource wealth to evaluate various outcomes. Consider first that at the micro-level, shocks to the value of local mineral wealth may act essentially as cash transfers to those in regions experiencing relative wealth booms.[6] Income shocks could generate windfall gains to local district households by generating greater mining employment and/or higher incomes. However, the extent of this effect will depend on the distribution of rents. The more that a small group of firms can capture the rents, and those rents are not re-injected in the local economy through any number of channels, the more likely it is that this effect is not observed. Supporting the idea that distribution of wealth matters, a growing literature suggests that international commodity price shocks may increase local civil conflict (Angrist and Kugler, 2008) and corruption (Asher and Novosad, 2014). Interestingly, this pattern may depend on the type of commodity shock. Dube and Vargas (2012) compare how price shocks to coffee (a labor intensive good) compare to price shocks in oil (much less labor intensive and concentrated ownership). They find that while the price of oil is positively related with conflict, the price of coffee is negatively related with conflict. Given limited information on how mineral wealth is distributed to households in South Africa, it is difficult to predict whether shocks to concentrated sectors affect broader population.

Finally, most research suggests that countries with non-democratic institutions are more likely to experience negative outcomes and corruption associated with natural resource wealth, as compared to countries with democratic institutions (Bhattacharya and Hodler, 2010; Arezki and Bruckner, 2012). As with many of its neighbors, there is evidence that mining companies in South Africa have participated in corrupt practices and financial misdealing with the government (e.g., Dolan and Stoddard, 2013; Standing, 2007). This, in addition to the legally-enforced segregation in South Africa under Apartheid, could prevent districts from reaping

[6] Note that shocks still make others relatively worse off (as any cash transfer does), so in theory they could have a similar "dutch disease" effect. But, in the short term, it unlikely to be making some districts strictly worse off in an absolute sense.

the benefits of mineral resources and instead explain the observed low growth.

Finally, this paper focuses on employment outcomes. Thus, some of the closest related papers relying on within country variation are from the United States. Black, McKinnish, and Sanders (2005) examine coal booms in Appalachia in the 1980s and Allcott and Keniston (2014) look at oil and gas booms from the 1960s-2000s. Like both of these papers, I find an increase in total employment. However, unlike these papers, I do not find employment spillovers to other sectors and instead find trade-offs. This could be because of lower overall employment or because of weaker links between sectors, which I consider in the discussion section.

3.3 Data

To explore the effects of mineral wealth gains on district economic outcomes, I combine international data on the locations and prices of minerals with Census survey data. I supplement this with more detailed household-level data to explore changes in income and hours worked.

3.3.1 Minerals Data

Minerals are an important subset of natural resources, generally found in rock and extracted via mining. This analysis utilizes data on the locations of mineral deposits collected by the U.S. Geological Survey (USGS), generally considered the most comprehensive source. The data includes all important mineral sources and it is presented in Table 3.1.[7] After excluding non-economically important minerals, a list of 25 minerals/mineral groups remains. The USGS also collects comprehensive data on international mineral prices, which I also rely upon. I combine the location of mineral deposits with panel data on annual average global prices for these minerals.[8] Though mineral resources are more concentrated in the eastern portion of the country, about 20 percent of all districts have any mineral deposits.

[7]Following similar papers, I exclude energy-producing minerals (coal, natural gas, petroleum, and uranium) and agriculture. This is because non-market forces may play a considerable role in local prices and rents for these minerals.

[8]I use nominal prices, but results are robust to using prices adjusted for US inflation. I do not yet have South African annual inflation factors but I will use these in future versions.

For a subset of the analysis, I divide minerals into "major" minerals and "minor" minerals. Major minerals include all those whose production ranks in the top 10 globally. Major minerals include gold, chromite, diamond, iridium, manganese, platinum group metals, titanium, vanadium, vermiculite, and zirconium. Minor minerals include aluminum, arsenic, asbestos, silver, cobalt, copper, iron, nickel, phosphate, lead, rare earth metals, antimony, silica, tungsten, and zinc.

Table 3.1: *List of Mineral Deposits*

Mineral	Number of Deposits	Number of Districts With Any Deposits	Production Rank
Gold[a]	123	30	6
Aluminum	2	2	12
Arsenic	10	8	Not top 10
Asbestos	5	5	Not top 30
Silver	33	20	19
Cobalt	9	3	21
Chromite	29	9	1
Copper	44	17	24
Diamond	9	7	7 (gemstone)
Iron	41	19	7 (iron ore)
Manganese	10	4	2
Nickel	36	12	11
Phosphate	5	3	14
Lead	9	5	13
Platinum Group Metals[b]	85	30	1
Platinium Only	129	37	1
Rare Earth Metals	1	1	Not top 10
Antimony	7	6	25
Silica	16	8	12
Titanium	13	11	3
Vanadium	13	10	3
Vermiculite	4	3	1
Tungsten	3	2	20
Zinc	4	3	27
Zirconium	10	7	2
Total Deposits	703	72	–

[a] Until 1996, South Africa was first in global gold production but fell during the 1990s.
[b] PGMs include iridium, platinum, palladium, osmium, rhodium, and ruthenium.
Notes:
1. The table reports deposits known as of 2010 and reported to USGS. World rankings are as reported by USGS in each mineral's 2012 yearbook. Some countries do not specify for some minerals, so rankings are based on reporting/USGS analysis. Almost all top producers included.
2. Energy-producing minerals (coal, natural gas, petroleum, and uranium) and agricultural natural resources are not included.
3. In some cases, USGS only has information on the top producers. In cases where South Africa is not the list, I denoted it as not in the top reported minerals. Preliminary checks against country production data confirm rankings.

3.3.2 Census Survey Data

I pool three rounds of the South Africa Census (1991, 1996, 2001) to create employment, poverty, and household measures.[9] Using survey responses by the household head, I construct district-level variables. The variables include the fraction of households in which the household head is employed, the fraction of households in which the household head is employed in mining, agriculture or manufacturing, the fraction of households above/below the poverty line, and the fraction of households in different salary categories.[10] While less information is consistently available in the Census data, the larger sample sizes and longer time span is advantageous.

Districts, the unit of analysis, are the magisterial district boundaries in place prior to the end of Apartheid. These are historically demarcated boundaries which can be linked to newer district boundaries. Districts operate as roughly independent labor markets, which makes it easier to compare outcomes across districts (Dinkelman, 2011). On average, districts contain about 100,000 people; however, district size varies considerably, with the smallest district having around 3,000 people and the largest district having over 800,000 people.

3.3.3 Household Data

I also pool five rounds of October Household Survey data as a source of more detailed hours worked and salary information. This data is similar to the World Bank LSMS surveys, and the survey was conducted in South Africa to inform reconstruction and development after the end of Apartheid. The survey aims to be a nationally-representative sample of household workers.

[9] For the 1996 and 2001 Census rounds, only ten percent samples are publicly available so those are utilized. For the 1991 Census, an almost-full sample is available but I randomly sample ten percent of the data for ease of analysis. In the 1991 Census, a handful of districts (less than 2%) are under-surveyed given "separate state" status under Apartheid. These districts were, in general, worse off than average districts in South Africa and removed from all years of the analysis.

[10] In the Census, both the respondent's occupation and the economic sector of the respondent's employer is reported. In some instances, the way the respondent describes their employment is different from the economic sector/industry of the household head's employer. For this version of the paper, I use the industry of the respondent's employer. In addition, the variables are constructed as "fraction of all household heads" within an industry category, rather than "fraction of employed household heads" in a given category. The variables are correlated and outcomes look similar either way.

3.4 Empirical Strategy

Below I first describe the motivation for my empirical approach. I then describe the steps I take to construct the variables for the analysis. Third, I describe the two specific analysis I perform in detail. The first set of regressions document the differences between districts with and without mineral resources. The second set of regressions measure the effect of changes over time in the value of local mineral resources on outcomes.

3.4.1 Motivation for Empirical Approach

The goal of this analysis is to identify how changes in mineral wealth affects districts, on average. Early research interested in this question compared countries with greater numbers of operational mines and high production to those with low numbers of known open mines and/or mine production. This approach, while often a useful step, suffers from endogeneity and reverse-causality. Minerals found in locations that are closer to economic activity or trade routes are more likely to be developed into mines. In addition, opening and developing productive mines requires institutions and infrastructure (e.g., as described in van der Ploeg, 2011).

To avoid this source of bias, I instead use the presence of mineral resource deposits, rather than active or productive mines. While it is still possible that locations with mineral deposits are systematically different from locations without minerals, mineral deposits themselves are less likely to be a *product* of infrastructure and institutions.[11] In addition, my approach, described below in Section 3.4.3, interacts mineral deposit locations with changes in international commodity prices to predict local economic outcomes (similar in spirit to Deaton and Miller, 1995; Angrist and Kugler, 2008; Dube and Vargas, 2008; Asher and Novosad, 2014). Using this approach, the identification comes from variations in the value of local minerals. As a robustness check, I also present the results using only districts with any mineral deposits.

[11] It is possible that mineral locations are correlated with other natural conditions, such as water or land slope. In this case, the omitted geographic variable could generate the economic difference, rather than the mineral wealth. If the paper were only to compare mineral locations to non-mineral locations, these omitted variables would violate the exclusion restriction. However, in this paper, this effect is not as concerning because the identification comes from time-varying international price changes.

The key assumption of this empirical approach is that country-level producers do not have enough market power to drive changes in international commodity prices. This assumption is generally valid, given that even in cases where South Africa is the top producer of a given commodity, they still control only a fraction of the market and their ability to drive prices is limited.[12] However, if this is violated, and changes in individual country production levels are generating changes in global prices, then omitted variable bias could be concerning. Specifically, one would worry that whichever factor that is changing production (e.g., institutions, political shifts) is also affecting district-level outcomes through channels other than through changes in mineral wealth. As a check on this, I also present regression results that measure the effects of minor minerals only and the effects of major minerals only.

3.4.2 Mineral Deposits and District Characteristics

Given the lack of data on the role that mineral resource wealth has had in South Africa, I first document any differences between districts with minerals and those without any mineral deposits. Similar to a cross-country analyses, a within-country cross-sectional analysis may demonstrate that districts with mineral resources are different from districts without mineral resources. I test the correlation between mineral resource deposits and district characteristics using a standard OLS specification:

$$Y_d = \beta_o + \beta_1 * Deposit_d + \beta_2 X_d' + \varepsilon_d \tag{3.1}$$

where $Deposit_d$ is an indicator for any mineral resource deposit and X_d is a set of standard district controls in district d. In this cross-sectional specification, the outcome variable is a variety of district-level characteristics constructed from the 1996 Census. In the analysis, robust standard errors clustered are provided.[13]

[12] The formula for whether or not a mineral is going to affect international prices (through market power) is based on the elasticity of demand for the product. Given mine operators are independent, it is rare for a mineral operator outside of an organization (e.g., OPEC) to have a significant influence on the year-level average international price.

[13] I cannot cluster by mineral because districts have a composite of multiple minerals or by province because South Africa has only nine provinces.

3.4.3 Measuring the Effect of Variation in Mineral Prices

The main objective of the paper is to estimate the extent to which changes in mineral wealth affect local economic outcomes. For each district, I first identify the types of minerals produced at each deposit location with a district. I then construct two district-level measures of the relative "price shock" for each given mineral.[14] The first measures short-term (one year) mineral price shocks, and the second measures longer-term (five year) mineral price shocks. For each mineral, I define the one-year price shock for that given mineral as:

$$PriceShock1yr_{m,t} = \frac{Price_{m,t}}{Average\ Price_m}$$

where $Price_{m,t}$ is the price for a given mineral m in year t and $Average\ Price_m$ is the average price from 1980-2010.[15] I construct the five-year price shock similarly, taking the average price from the current year plus the previous four years in the numerator:

$$PriceShock5yr_{m,t} = \frac{\frac{1}{5} \overset{[}{i} = t-4]i = t \sum Price_{m,i}}{Average\ Price_m}$$

Each price shock variable provides a measure of the current price relative to the average price during the period. Values of the price shock greater than one indicate a relatively higher price over the past one/five years for a given mineral. Likewise, values lower than one indicate relatively lower prices for a given mineral.

These price shock values are then incorporated into a general measure of changes to mineral wealth for the district d in year t. Again, the one-year mineral index indicates the short-term changes to mineral wealth to the district and the five-year mineral index indicates relatively longer-term shifts in wealth (though both measures are still "short term" in relation to broader economic analysis). For each district, I identify the mineral deposits within that district. I then calculate a mineral index, which is the average price shock for all present minerals, weighted

[14] This construction, using a combination of price shocks to all minerals, follows the approach of Asher and Novosad (2014).

[15] I use nominal US dollars. Future versions will control for South African inflation/purchasing power, but that data is currently not available (effects would be small given year fixed effects). The results look the same if I use real (US inflation-adjusted) US prices.

by the number of deposits of that mineral within the district. The resulting mineral price index for a given period is as follows, for one-year and five-year indices, respectively:

$$MineralPriceIndex1yr_{dt} = \frac{\sum_{m=1}^{m=M} (PriceShock1yr_{m,t} * Number\ Deposits_{m,d})}{Total\ Number\ Deposits_d}$$

$$MineralPriceIndex5yr_{dt} = \frac{\sum_{m=1}^{m=M} (PriceShock5yr_{m,t} * Number\ Deposits_{m,d})}{Total\ Number\ Deposits_d}$$

where m reflects each given mineral, M is all minerals, and the variable is constructed for a district d and year t.

I then estimate the effect of mineral price index on outcomes in the district:

$$Y_d = \beta_o + \beta_1 (MineralPriceIndex_{dt}) + \lambda_d + year_t + \varepsilon_d \tag{3.2}$$

where I insert the one-year and five-year indices into the $MineralPriceIndex_{dt}$ term.[16]

In Figure 3.1, I present the means and distributions of the constructed indices for those districts with any mineral deposits. The figure demonstrates considerable variation in the value of mineral wealth, particularly when only considering minor minerals. In addition, though not shown, significant variation exists both across districts and within districts over time, driven by the varying compositions of mineral deposits in different districts.

[16] I have also run the above equation with additional controls for the overall district employment, race breakdown, income, and education levels (using 1991 Census data). The results remain similar.

Figure 3.1: *Average Value of the Instrument Across Districts*

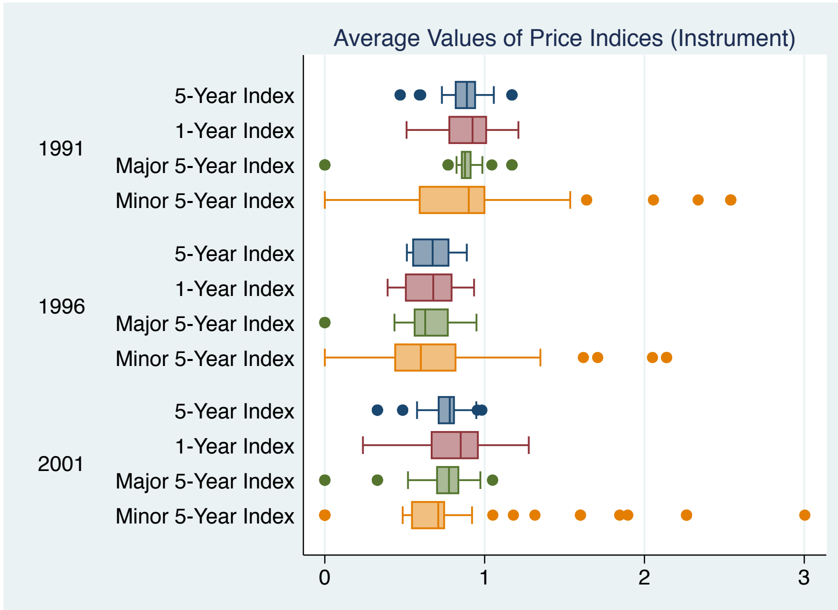

The above figure demonstrates that the instrument has considerable variance across districts, but somewhat smaller variance across years (though it still moves 15-25 percent between the years with Census data, on average). The figure also demonstrates that smaller minerals experienced greater variance.

In Figure 3.2, I present average value of the five-year mineral price shock for each mineral, for the three years with Census data. The prices shocks vary considerably across minerals. As the graph shows, during the years of matching Census data (1991, 1996, and 2001), the minerals did worse than their long-term average prices. In addition, although not shown here, price shocks also vary significantly by year, with no single mineral dominating the results that follow.

Figure 3.2: *Variation in Minerals Price Index within South Africa*

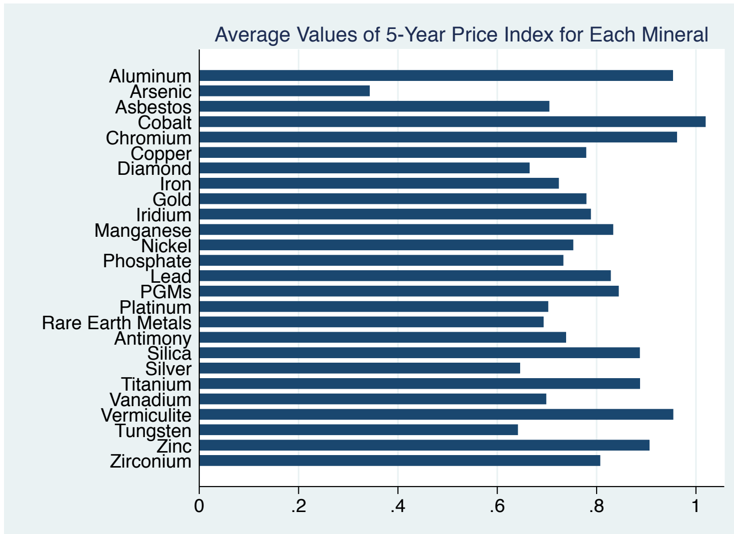

The above figure reflects the average value of the 5-Year Mineral Price Index for each economically important mineral. As is clear from the figure, most prices were on average lower during this period than during the entire period.

3.5 Results

3.5.1 Districts with Mineral Deposits

In Table 3.2, I present the results of equation (1), using a cross-section of data from the 1996 Census. This table provides several basic facts about districts with mineral endowments. First, although districts with mineral deposits have greater populations, they also have larger district areas. Because the land effect dominates, overall population density is negatively correlated with the presence of minerals. Second, districts with minerals have employment rates about ten percentile points higher than districts without mineral deposits. In particular, though overall mining employment is small, households are more than twice as likely to be employed in mining. Districts with any mineral deposits also have slightly higher incomes, greater piped water and electricity access, and higher education levels. The results look similar if I restrict the "Any Deposit" variable to only major or minor minerals (given many of them are co-located).

Finally, note that this preliminary version of the paper does not consider migration, although

Table 3.2: *Cross-Section Relationship between Deposit and District Characteristics*

Outcome Variable:	Coefficient on Mineral Indicator:	Coefficient on Mineral Indicator, with controls	Outcome Mean (SD):
Log(District Population)	0.720***	0.343***	10.88
	(0.136)	(0.0966)	(1.229)
Land Area (sq km)	3,530***	3,011***	2,730
	(907.3)	(953.8)	(3,170)
Density (N/sq km)	-147.5**	-230.4**	178
	(58.78)	(93.50)	(640)
Employment	0.102***	0.0185**	0.500
	(0.0243)	(0.00830)	(0.199)
Mine Employment	0.0189***	0.0162***	0.00777
	(0.00331)	(0.00319)	(.0155)
Agriculture Employment	-0.0226**	0.00317	0.0824
	(0.0105)	(0.00763)	(.0888)
Female-Headed House	-0.0469***	-0.00686	0.3721
	(0.0159)	(0.00560)	(0.1337)
African Race	0.0647**	0.0589***	0.6969
	(0.0304)	(0.0194)	(.3162)
Poverty Indicator	-0.0657***	-0.00248	0.4388
	(0.0194)	(0.00865)	(0.1703)
High School Education	0.0549***	0.0166*	0.1917
	(0.0145)	(0.00900)	(.1002)
Piped Water	0.0678**	0.00892	0.5853
	(0.0339)	(0.0175)	(.2958)
Bucket Toilet	-0.0493***	0.000165	0.1209
	(0.0188)	(0.0176)	(.1728)
Electricity for Cooking	0.0963***	-0.0241*	0.3600
	(0.0293)	(0.0126)	(.2379)
Observations	354	354	

Notes:
1. Each observation is a district. The outcome variables are listed in the left-most column. Each row is a regression of the outcome variable on an indicator variable which equals one if a district has any mineral deposits. District population is the total district population, and the subsequent rows are the fraction of the district households with a given characteristic. In column (1), the outcome variable is regressed on the "Any Deposit" variable. In column (2), the regression includes controls for the other variables.
2. The "Any Deposit" variable is an indicator equal to one if any mineral deposit exists within the district (major or minor).
3. Robust standard errors are used. Astericks denote significance: *p < 0.10,** p < 0.05,*** p < 0.01.

many men migrate across districts to work in the mines. This factor will generate spillover effects, improving outcomes in districts without minerals. Given this, the current analysis provides a conservative estimates of how these districts may differ. In subsequent analysis described below, it also generates more conservative estimates.

3.5.2 Causal Effects of Resource Wealth Changes

In Table 3.3, I present the results of equation (2), measuring the effect of the five-year mineral price index on local economic outcomes. The takeaway from Table 3.3 is that sustained high mineral prices shift the average employment composition of the district: an increasing fraction participate in mining and a decreasing fraction participate in agriculture.

Table 3.3: *The Effect of Five-Year Price Shock on District Employment and Poverty*

Outcome Variable:	Any Employment	Employment in Mining	Employment in Agriculture	Employment in Manufacture	Poverty Indicator
	(1)	(2)	(3)	(4)	(5)
Panel A. All Districts					
Price Index (5-year)	0.0967*	0.167***	-0.145***	-0.0156	-0.0783***
	(0.0429)	(0.0593)	(0.0535)	(0.0145)	(0.0189)
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Dependent Mean	0.3276	0.07638	0.1429	.0774	.5072
Observations	1,021	1,014	1,014	1,014	1,021
Panel B. Only Districts with Any Minerals					
Price Index (5-year)	0.0266	0.125**	-0.0458	-0.00731	-0.0213
	(0.0610)	(0.0599)	(0.0567)	(0.00926)	(0.158)
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Dependent Mean	0.3813	0.1083	0.1086	.0870	.4765
Observations	211	209	209	209	211

Notes:
1. Each observation is a district-year from 1991, 1996, and 2001. In column (1), the outcome variable is the fraction of district households which reports any employment. In Column (2), (3), and (4) the outcome is fraction of district households which report employment within given sectors. Column (5) is a variable reflecting the fraction of households with income below the annual poverty indicator. Columns (2), (3) and (4) are missing seven districts in 1991, which did not report breakdowns by employment type. Panel A includes all districts, while Panel B restricts the analysis to only those districts with any mineral deposits.
2. District and year fixed effects are included in all regressions.
3. The dataset was created from the 1991, 1996, and 2001 Census and the USGS Mineral Yearbooks.
4. Robust standard errors, clustered at the district level, are used. Astericks denote significance: *** p<0.01, ** p<0.05, * p<0.1.

To interpret the results further, consider that the coefficient on the index in each regression reflects the effect that a one unit increase in the price shock index variable would have on outcomes. If a district has a single mineral deposit, a one unit increase indicates an average price over the five-year period that is 100 percent higher than the overall average price. For a district with many mineral deposits, the average price of all minerals, weighted by number of deposits, is the "average" district price and a one unit increase indicates a 100 percent increase in this average price.

In Panel A, which includes all districts, a district with a 100 percent increase in the price index would have had employment that was 16.7 percentile points higher than places with average prices. However, the average price shock during this period was less than one (0.157), suggesting that on average, prices were 84 percent lower than longer-term averages and mineral employment was 14 percentile points lower.

In Table 3.4, I perform the same estimation, but I use two index variables that are constructed separately, one using only major minerals and the other using only minor minerals. For this analysis, major minerals are those ranked within the top ten, as shown in Table 3.1 and minor minerals are those not ranking in the top ten. This breakdown is useful for a few reasons. First, major minerals may generate greater effects given larger baseline production volume and *total* value produced. However, the results for major minerals should be evaluated with more caution because shocks to major mineral prices are more likely to violate the exclusion restriction. In other words, a top-ten producer is more likely to affect global prices through changes in its production level. If an omitted variable (e.g., a strike) affects production and other local outcomes, and that production affects global prices, the estimated price index coefficient is confounded. Luckily, most financial analysis suggests that annual prices are unlikely to move substantially based on a single country's activities, and particularly unlikely to move given a single mine's activities.

In odd-numbered columns of Table 3.4, I regress the outcome variable on a price shock variable constructed using major minerals. In even-numbered columns of Table 3.4, I regress the outcome variable on a price shock index constructed using only minor minerals. This table demonstrates that the pattern holds within each subgroup of minerals, and the results are not statistically different for each group. This is not completely surprising, given many major minerals were declining during this period and price variation may have been correlated across different minerals (e.g., between major and minor minerals, which are also sometimes co-located). The reduction in agricultural and manufacturing employment levels are somewhat more clear when only considering minor minerals. This may be because the greater price variation among minor minerals provides more variation for the identification.

In Table 3.5, I look at a limited sample of household-level data from 1995 to 1999. This data includes more detailed labor information, in particular on hourly work and salary level. This data shows that across all districts, when districts experience mineral price shocks, households work more hours. Column (3) is largely driven by the extensive margin, with more households working. However, Column (4) shows that among those that already work, hours also increase (intensive margin response).

Table 3.4: *The Effect of Five-Year Price Shock on District Employment and Poverty, by Major/Minor Minerals*

Outcome Variable:	Any Employment		Employment in Mining		Employment in Agriculture		Employment in Manufacturing		Poverty Indicator	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Major Min. Price Index(5-year)	0.0314		0.136*		-0.126		-0.0124		-.102***	
	(0.0617)		(0.0776)		(0.0768)		(0.00957)		(.012)	
Minor Min. Price Index(5-year)		0.0490		0.133*		-0.155**		-0.0167*		-.042***
		(0.0615)		(0.0780)		(0.0648)		(.00978)		(.0111)
Mean Dependent Variable	0.3276		0.07638		0.1429		.0774		0.5072	
Observations	1,021		1,014		1,014		1,014		1,021	

Notes:
1. Each observation is a district-year from 1991, 1996, and 2001. In columns (1)-(2), the outcome variable is the fraction of district households which reports any employment. In columns (3) -(8), the outcome is fraction of district households which report employment within given sectors. The outcome variable in columns (9)-(10) reflect the fraction of households with income below the annual poverty indicator.
2. District and year fixed effects are included in all regressions.
3. The dataset was created from the 1991, 1996, and 2001 Census and the USGS Mineral Yearbooks.
4. Robust standard errors, clustered at the district level, are used. Astericks denote significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 3.5: *The Effect of One-Year Price Shock on Individual Employment and Poverty (1995-1999)*

Outcome Variable:	Currently Working	Work in the Past Year	Hours of Work (all)	Hours of Work (conditional on job)	Salary
	(1)	(2)	(3)	(4)	(5)
Price Index (1-year)	0.0974***	0.0961***	5.796***	1.857***	0.546***
	(0.0174)	(0.0168)	(0.924)	(0.511)	(0.0914)
Province-Year Controls	Yes	Yes	Yes	Yes	Yes
Mean Dependent Variable	0.53	0.56	24.92	46.56	2.76
Observations	120,174	120,174	120,174	63,903	120,174

*** p<0.01, ** p<0.05, * p<0.1
Notes:
1. Each observation is a household head. Robust standard errors clustered at the province-year level.
2. Data source is the October Household Survey from 1995, 1996, 1997, 1998, and 1999.

Interestingly, when the results are separated out by distinct minerals, it is not one single type of mineral deposit that drives the results. Instead, in different years and in different districts, shocks to different minerals appear be generating to the results. In future, I hope to incorporate greater institutional and political context regarding the different minerals, in order to understand which ones may be most likely to generate the variation in outcomes.

In Table 3.6, I present the regression results for how mineral shocks affect the 0.25 quantile, the 0.5 quantile and the 0.75 quantile of average district employment. As the results show, employment improvements are driven by districts that already have higher levels of employment.

Interestingly, with the current construction of the index, I do not find that one-year or five-year shocks to mineral prices increase access to piped water or reduce the distance to medical facilities. It is likely, however, that greater access to public services depends on both increased income (demand) and increased construction of these public services (supply). It is likely that there is a lag in the supply of public services, which could be why I do not observe a significant change in these household services outcomes. However, in future versions I will

Table 3.6: *Quartile Regression of District Employment on Income Shocks*

Outcome Variable:	Fraction of District Employed		
Quartile:	0.25	0.5	0.75
Price Index (5-year)	0.0248	0.0632	0.214**
	(0.260)	(0.680)	(0.107)
Fixed Effects	Yes	Yes	Yes
Observations	1,021	1,021	1,021

*** p<0.01, ** p<0.05, * p<0.1
Notes:
1. Each observation is the fraction of the district that is employed at all.
2. Data source is the Census from 1991, 1996, and 2001.

examine other household-level services in more detail, testing factors such as different lags, different constructions of the price index, and running the analysis with a more robust set of controls.

3.6 Discussion and Conclusion

This analysis provides evidence that shocks to mineral wealth within a district increase employment, particularly in mining, and reduce the number of households living in poverty. Households also report working more hours and earning greater salaries. The gains are large, providing an important transfer of wealth to mineral-endowed districts. However, future work is needed to understand the mechanism through which the observed outcome arises. For example, governments may spend more on districts that are experiencing a resource boom, and these results would be inclusive of this effect. However, given that changes to aggregate mineral wealth are driven in various years by different minerals, and given that government royalties are distributed broadly, this may be a somewhat less likely channel. Rather, it seems more likely that an increase in natural resource prices causes mining companies to grow, either on the intensive margin by hiring new workers at existing mines or on the extensive margin by opening up new mines or supporting businesses.

These results are essentially "short run." The mineral wealth shocks are measured over the last five years, and the outcome data spans ten years. In this short run analysis, districts with mineral deposits benefit from their value. However, longer-term macroeconomic arguments that South Africa has experienced a resource curse, as measured through metrics such as GDP, may still be valid. Specifically, the analysis demonstrates trade-offs, rather than complementarities, between mining and other employment sectors. An increase in mineral wealth reduced the number of households employed in agriculture and to a lesser extent, manufacturing. This trade-off could be evidence of a longer-run natural resource curse, whereby South Africa over-invests in natural resources at the expense of agriculture and manufacturing, as Hausmann and Klinger (2008) argue. Given the gains measured in this paper, reconciling short term employment and income increases with possibly longer-term lower growth is key challenge for future research.

Bibliography

Akinnifesi, F. K., Ajayi, O., Sileshi, G., Chirwa, P. W., and Chianu, J., **J**. (2010). Fertiliser trees for sustainable food security in the maize-based production systems of East and Southern Africa: A review. *Agronomy for Sustainable Development, 30*(3):615–629.

Allcott, H., and Keniston, D. (2014). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. NBER Working Paper Series No. 20508.

Almond, D., and Currie, J. (2010). Human capital development before age five. *Handbook of Labor Economics, 4*: 1315–1486.

Angrist, J. D., and Kugler, A. D.(2008). Rural windfall or a new resource curse? Coca, income and civil conflict in Colombia. *The Review of Economics and Statistics, 90* (2): 191–215.

Aragon, F., and Rud, J. P. (2011). Natural resources and local communities: Evidence from a Peruvian gold mine. Working Paper, Simon Fraser University (April).

Arceo-Gomez, E. O., Hanna, R. N., and Oliva, P. (2012). Does the effect of pollution on infant mortality differ between developing and developed countries? Evidence from Mexico City. NBER Working Paper 18349.

Arezki, R., and Bruckner, M. (2012). Commodity windfalls, democracy and external debt. *The Economic Journal 122*: 848–866.

Arrow, K. J., and Fisher, A. C. (1974). Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics, 88*(2):312–319.

Asher, S., and Novosad, P. (2014). Dirty politics: Natural resource wealth and politics in India. Working Paper, Oxford University (April).

Ashraf, N., Berry, J., and Shapiro, J. M. (2010). Can higher prices stimulate product use? Evidence from a field experiment in Zambia. *American Economic Review, 100*(5):2383–2413.

Ashraf, N., Jack, B. K., and Kamenica, E. (2013). Information and subsidies: Complements or substitutes? *Journal of Economic Behavior & Organization, 88*:133–139.

Ashraf, N. Bandiera, O. and Jack, B. K. (2014). No margin, no mission? a field experiment on incentives for public service delivery. *Journal of Public Economics, 120*:1–17.

Auty, R. M. (1993). Sustaining development in mineral economies: the resource curse thesis. Routledge, New York.

Auty, R. M., and Mikesell, R.F. (1998). Sustainable development in mineral economies. Oxford University Press, New York.

Avery, A. A. (1999). Infantile methemoglobinemia: reexamining the role of drinking water nitrates. *Environmental Health Perspectives 107*: 583–586.

Banjeree, A., Galiani, S., Levinsohn, J., and Woolard, I. (2007). Why has unemployment risen in the new South Africa? NBER Working Paper No. 13167.

Barnes, R., and Fagg, C.(2003) Faidherbia albida: Monograph and annotated bibliography.

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

Berry, J. Fischer, G., and Guiteras, R. P. (2012). Eliciting and utilizing willingness to pay: Evidence from field trials in northern Ghana. Working Paper.

Bevan, D., Collier, P., and Gunning, J. W. (1991). Temporary trade shocks in developing countries: Consequences and policy responses. Oxford, Centre for the Study of African Economies.

Black, D., McKinnish, T., and Sanders, S. (2005). The economic impact of the coal boom and bust. *The Economic Journal 115:* 449-476.

Blaine, S. (2013). Agriculture is putting SA in deep water. *Business Day Live.* Retrieved from http://www.bdlive.co.za/business/agriculture/2013/04/17/ agriculture-is-putting-sa-in-deep-water.

Blignaut, J., Ueckermann, L., and Aronson, J. (2009). Agriculture production's sensitivity to changes in climate in South Africa. *South African Journal of Science 105*: 61-68.

Blignaut, J. N., and Hassan, R. M. (2002). Assessment of the performance and sustainability of mining sub-soil assets for economic development in South Africa. *Ecological Economics, 40* (1): 89-101.

Borja-Aburto, V. H., Castillejos, M., Gold, D. R., Bierzwinski, S., and Loomis, D. (1998). Mortality and ambient fine particles in Southwest Mexico City, 1993–1995. *Environmental Health Perspectives, 106* (12), 849–855.

Bradshaw, D., Groenewald, P., Laubscher, R., Nannan, N., Nojilana, B., and Norman, R.

(2003). Initial burden of disease estimates for South Africa: 2000. *South African Medical Journal, 93*: 6828.

Brainerd, E., and Menon, N. (2014). Seasonal effects of water quality: The hidden costs of the Green Revolution to infant and child health in India. *Journal of Development Economics, 104*, 49-64.

Brollo, F., Nannicini, T., Roberto, P., and Tabellini, G. (2013). The political resource curse. *American Economic Review, 103* (5): 1759-96.

Brown, S. (1997). Estimating biomass and biomass change of tropical forests: A primer.

Bruckner, M., and Ciccone, A. (2010). International commodity prices, growth and the outbreak of civil war in Sub-Saharan Africa. *The Economic Journal, 120*: 519–534.

Bruhn, M., and McKenzie, D.(2009). In pursuit of balance: Randomization in practice in development field experiments. *American Economic Journal: Applied Economics, 1*(4):200–232.

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

Cardenas, J. C., and Carpenter, J. (2008). Behavioural development economics: Lessons from field labs in the developing world. *Journal of Development Studies, 44*(3):311–338.

Carter, R., Laajaj, R., and Yang, D. (2014). Subsidies and the persistence of technology adoption: Field experimental evidence from Mozambique. Working Paper.

Chassang, S., Miquel, G. P., and Snowberg, E. (2012). Selective trials: A principal-agent approach to randomized controlled experiments. *American Economic Review, 102*(4):1279–1309.

Chay, K. Y., and Greenstone, M. (2003). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession. *The Quarterly Journal of Economics, 118*(3), 1121-1167.

Clay, K., Troesken, W., and Haines, M. R. (2010). Lead and mortality. NBER Working Paper Series No. 16480.

Cohen, J., and Dupas, P. (2010). Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics, 125* (1):1–45.

Conning, J., and Udry, C. (2007). Rural financial markets in developing countries. In *Handbook of Agricultural Economics, 3*, : 2857–2908.

Cosgrove, W. J., and Rijsberman, F. R. (2000). The use of water today. In *World Water Vision: Making Water Everybody's Business* (chapter 2). World Water Council. Retrieved from: http://www.worldwatercouncil.org/fileadmin/wwc/Library/WWVision/Chapter2.pdf.

Council for Scientific and Industrial Research (CSIR) (2010). *A CSIR Perspective on Water in South Africa - 2010.* CSIR Report No. CSIR/NRE/PW/IR/2011/0012/A. Compiled by: S. Oelofse and W. Strydom. Retrieved from: www.csir.co.za

Croen L. A., Todoroff, K., and Shaw G. M.(2001). Maternal exposure to nitrate from drinking water and diet and risk for neural tube defects. *American Journal of Epidemiology, 153*:325–331.

Currie, J., and Neidell, M. (2005). Air pollution and infant health: what can we learn from California's recent experience? *The Quarterly Journal of Economics, 120*(3): 1003-1030.

Currie, J., and Vogl, T. (2012). Early life health and adult circumstances in developing countries. NBER Working Paper 15827.

Cutler, D., and Miller, G. (2005). The role of public health improvements in health advances: the twentieth-century United States. *Demography, 42* (1): 1-22.

Dam Safety Office (2013). List of Registered Dams (March 2013). Retrieved from: https://www.dwaf.gov.za/DSO/Default.aspx.

Davies, B.R., O'Keeffe, J. H., and Snaddon, C.D. (1995). River and stream ecosystems in Southern Africa: predictably unpredictable. In: Cushing CE, Cummins KW and Minshall GW (eds.) *River and Stream Ecosystems.* Amsterdam, The Netherlands: 537-599.

Deaton, A., and Miller, R. (1995). International commodity prices, macroeconomic performance, and politics in Sub-Saharan Africa. Priceton Studies in International Finance, No. 79.

Department of Water Affairs. (1996). South Africa Water Quality Guidelines: Domestic Use. CSIR Environmental Services, Pretoria, South Africa. Retrieved from: http://www.dwaf.gov.za/iwqs/wq_guide/index.html

Department of Water Affairs. (2012). South Africa Rivers. Retrieved from: http://www.dwaf.gov.za/iwqs/gis_data/river/rivs500k.html.

Dercon, S. (2002). Income risk, coping strategies, and safety nets. *The World Bank Research Observer, 17* (2):141–166.

Dinkelman, T. (2011). The effects of rural electrification on employment: New evidence from South Africa. *American Economic Review, 101*(7): 3078-3108.

Dixit, A., and Pindyck, R. (1994). Investment under uncertainty. Princeton University Press.

Dube, O., and Vargas, J. F. (2013). Commodity price shocks and civil conflict: Evidence from Colombia," *Review of Economic Studies, 80*: 1384-1421.

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

Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica, 82*(1):197–228.

Ebenstein, A., Zhang, J., McMillan, M., and Chen, K. (2011). Chemical fertilizer and migration in China. NBER Working Papers 17245.

Ebenstein, A. (2012). The consequences of industrialization: Evidence from water pollution and digestive cancers in China. *Review of Economics and Statistics, 94* (1): 186-201.

Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., and Cullen, M. R. (2013). Selection on moral hazard in health insurance. *American Economic Review, 103* (1):178–219.

Elbra, A. (2013). The forgotten resource curse: South Africa's poor experience with mineral extraction. *Resources Policy, 38*: 549–557.

Esrey S., Potash, J., Roberts, L., and Schiff, C. (1991). Effects of improved water supply and sanitation on ascariasis, diarrhoea, drancunculiasis, hookworm infection, schistosomiasis and trachoma. *Bulletin of the World Health Organization, (60)*: 609-21.

Fewtrell, L. (2004). Drinking-water nitrate, methemoglobinemia, and global burden of disease: A discussion. *Environmental Health Perspectives, 112*(14): 1371–1374.

Field, E., Glennerster, R., and Hussam, R. (2011). Throwing the baby out with the drinking water: Unintended consequences of arsenic mitigation efforts in Bangladesh. Working Paper.

Fink, G., and Masiye, F. (2012). Assessing the impact of scaling-up bednet coverage through agricultural loan programmes: evidence from a cluster randomised controlled trial in Katete, Zambia. *Transactions of the Royal Society of Tropical Medicine and Hygiene, 106*(11):660–667.

Fink, G., Jack, B. K., and Masiye, F. (2014). Seasonal liquidity constraints and agricultural productivity: Evidence from Zambia. Working Paper.

Fischer, G., Karlan, D., McConnell, M., and Raffler, P. (2014). To charge or not to charge: Evidence from a health products experiment in Uganda. Working Paper.

Food and Agriculture Organization of the United Nations (FAO). (2005). Fertilizer use by crop in South Africa. Land and Water Development Division.

Food and Agriculture Organization of the United Nations (FAO). (2007). Dams and agriculture in Africa. Prepared for the Aquastat program.

Foster, A. D., and Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annual Review of Economics, 2*(1):395–424.

Funke, N., Nortje, K., Findlater, K., Burns, M., Turton, A., Weaver, A., and Hattingh, H., H.

(2007). Redressing inequality: South Africa's new water policy. *Environment Magazine (April).* Retrieved from:
http://www.environmentmagazine.org/Archives/ Back%20Issues/April%202007/Funke-full.html.

Garry, V. F., Schreinemachers, D., Harkins, M. E., and Griffith, J. (2002). Pesticide appliers, biocides, and birth defects in rural Minnesota. Environmental Health Perspectives, 104(4): 394–399.

Galiani, S., Gertler, P. J., and Schargrodsky, E. (2005). "Water for Life: The Impact of the Privatization of Water Services on Child Mortality." *Journal of Political Economy, 113*: 83-120.

Gertler, P., and Gruber, J. (2002). Insuring consumption against illness. *American Economic Review, 92*(1): 51–70.

Greenstone, M., and Hanna, R. N. (2014). Environmental regulations, air and water pollution, and infant mortality in India. *American Economic Review, 104(10)*: 3038–3072.

Greenstone, M., and Jack, B. K. (2014). Envirodevonomics: A research agenda for an emerging field. Forthcoming in *Journal of Economic Literature.*

Hausmann, R., and Klinger, B. (2008). South Africa's export predicament. *Economics of Transition, 16* (4): 609–637.

Heckman, J. J., and Navarro, S. (2005). Empirical estimates of option values of education and information sets in a dynamic sequential choice model. Working paper.

Heckman, J. J., and Navarro, S. (2007). Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics, 136*(2):341–396.

Hughes, D. (2004). Incorporating groundwater recharge and discharge functions into an existing monthly rainfall–runoff model. *Hydrological Sciences Journal, 42* (2): 297-311.

Huizenga, J. (2011). Characterisation of the inorganic chemistry of surface waters in South Africa. *Water South Africa, 37* (3):401-410.

Hunter, P. R., MacDonald, A. M., and Carter, R. C. (2010). Water Supply and Health. *PLOS Medicine, 7* (11).

Hunter, M., and Posel, D. (2012). Here to work: the socioeconomic characteristics of informal dwellers in post-apartheid South Africa. *Environment and Urbanization, 24*: 285-304.

Jack, B. K. (2013). Private information and the allocation of land use subsidies in malawi. *American Economic Journal: Applied Economics, 5*:113–35.

Jalan, J., and Ravallion, M. (2003). Does piped water reduce diarrhea for children in rural India? *Journal of Econometrics, 112*:153-173.

Jayachandran, S. (2009). Air quality and early-life mortality: Evidence from Indonesia's wildfires. *Journal of Human Resources, 44* (4).

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

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

Knobeloch, L., Salna, B., Hogan, A., Postle, J., and Anderson, H. (2000). Blue babies and nitrate-contaminated well water. *Environmental Health Perspectives, 108*(7): 675-678.

Kremer, M., Leino, J., Miguel, E., and Zwane, A. P. (2011). Spring cleaning: Rural water impacts, valuation, and property rights institutions. *Quarterly Journal of Economics, 126*(1): 145-205.

Levine, D. I., Beltramo, T., Blalock, G., and Cotterman, C. (2012). What impedes efficient adoption of products? evidence from randomized variation in sales offers for improved cookstoves in Uganda. Unpublished.

Lipscomb, M., and Mobarak, A. M. (2011). Decentralization and the political economy of water pollution: Evidence from the re-drawing of county borders in Brazil. Working Paper.

Mahajan, A., and Tarozzi, A. (2011). Time inconsistency, expectations and technology adoption: The case of insecticide treated nets. Working Paper.

Manassaram, D. M., Backer, L. C., and Moll, D. M. (2006). A review of nitrates in drinking water: maternal exposure and adverse reproductive and developmental outcomes. *Environmental Health Perspectives, 114* (3): 320–327.

Mantel, S., Hughes, D., and Muller, N.W. (2010). Ecological impacts of small dams on South African rivers: Part 1: Drivers of change – water quantity and quality. *Water South Africa 36*:3.

Munger, R., Isacson, P., Hu, S., Burns, T., Hanson, J., and Lynch, C.F. (1992). Birth defects and pesticide-contaminated water supplies in Iowa. *American Journal of Epidemiology*: 136:959.

Neary, J. P., and Corden, W. M. (1982). Booming sector and de-industrialization in a small open economy. *The Economic Journal, 92*: 825-848.

Nilsson, C., and B. M. Renofalt. (2008). Linking flow regime and water quality in rivers: a challenge to adaptive catchment management." *Ecology and Society, 13*(2): 18.

Oberholster, P. J., and Ashton, P. J. (2008). State of the Nation report: An overview of the current status of water quality and eutrophication in South African rivers and reservoirs. Parliametary Grant Deliverable.

Oberholster, P., Aneck-Hahn, N., Ashton, P. J., Botha, A., Brown, J., Dabrowski, J., de Klerk, A., de Klerk, L., Genthea, B., Geyer, H., Hall, G., Hill, L., Hoffman, A., Kleynhansf, C., Lai, J., le Roux, W., Luus-Powell, W.Masekoameng, E. McMillan, P. Myburgh, J., Schachtschneider, K., Somerset, V., Steyl, J., Surridge, A. K. J., Swanevelder, Z. H., van Zijl, M. C., Williams, C., and Woodborne, S. (2011). Risk assessment of pollution in surface waters of the upper Olifants River system: Implications for aquatic ecosystem health and the health of human users of water" Council of Scientific and Industrial Research report number CSIR/NRE/WR/IR/2011/0041/B.

O'Connor, T. (2001). Effect of small catchment dams on down- stream vegetation of a seasonal river in semi-arid African savanna. *Journal of Applied Ecology, 38*:1314-1325.

O'Keeffe, J., and De Moor, F. C. (1988). Changes in the physico-chemistry and benthic invertebrates of the Great Fish River, South Africa, following an interbasin transfer of water. *Regulated Rivers: Residual Management, 2*: 39-55.

Olmstead, S. (2009). The Economics of Water Quality. *Review of Environmental Economics and Policy, 4*(1): pp. 44–62.

Oster, E., and Thornton, R. (2012). Determinants of technology adoption: Peer effects in menstrual cup take-up. *Journal of the European Economic Association, 10*(6):1263–1293.

Phillips, H. (1999). Demographic survey data in South Africa: An evaluation of methodology and quality. *SA Journal of Demography, 7*(1):1-10.

Pindyck, R. (1993). Investments of uncertain cost. *Journal of Financial Economics, 34*(1):53–76.

Platzky, L., and Walker, C. (1985). *Surplus People* (1st ed.). Johannesburg: Ravan Press.

Pruss, A., Kay, D., Fewtrell, L., and Bartram, J. (2002). Estimating the burden of disease from water, sanitation, and hygiene at a global level. *Environmental Health Perspectives, 110*(5): 537-542.

Ricardo, D. (1817). *On the Principles of Political Economy and Taxation.* London: John Murray.

Sachs, J., and Warner, A. (1995). Natural resource abundance and economic growth, Development Discussion Paper No. 517a. Harvard Institute for International Development, Harvard University.

Saka, A., Bunderson, W., Itimu, O., Phombeya, O., and Mbekeani, Y. (1994). The effects of acacia albida on soils and maize grain yields under smallholder farm conditions in Malawi. *Forest Ecology and Management, 64*:217–230.

Schoonbee, H., Adendorff, F. C., de Wet, L. M., de Wet, L. P. D., Fleischer, C. L., van der Merwe, C. G., van Eeden, P. H. and Venter, A. J. A. (2006). The occurance and accumulation of selected heavy metals in fresh water ecosystems affected by mine and industrial polluted effluent. Water Resources Commission Project No. 312.

Seckler, D., Barker, R., and Amarasinghe, U. (1999). Water scarcity in the twenty-first century. *International Journal of Water Resources Development 15*: 29–42.

Shiklomanov, I. A. (1999). World water resources and water use: Present assessment and outlook for 2025. State Hydrological Institute, St. Petersburg, Russia.

Silverbauer, M. (2009). Methods for visualizing complex water quality. PhD Thesis. Department of Zoology, University of Cape Town.

South Africa Department of Enviromental Affairs and Tourism (DEAT). (2006). *A Report on the State of the Environment.* Pretoria, South Africa.

South Africa National Committee on Irrigation and Drainage (SANCID) (2014). Country Profile - South Africa. Retrieved from: http://www.icid.org/cp_safrica.html.

Standing, A. (2007). Corruption and the extractive industries in Africa, Institute for Security Studies.

Stange, K. M.(2012). An empirical investigation of the option value of college enrollment. *American Economic Journal: Applied Economics, 4*(1):49–84.

Stock, J. H., and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In Andrews DWK Identification and Inference for Econometric Models, New York: Cambridge University Press: 80-108.

Stokke, H. E. (2008). Resource boom, productivity growth and real exchange rate dynamics: A dynamic general equilibrium analysis of South Africa. *Economic Modelling, 25*(1): 148-160.

Strobl, E., and Blanc, E. (2013). Is small better? A comparison of the effect of large and small dams on cropland productivity in South Africa" World Bank Policy Research Working Paper 6567.

Strobl, E., and Strobl, R. (2010). "The Distributional impact of large dams: Evidence from cropland productivity in Africa. *Journal of Development Economics* 96 (2): 1–19.

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

Tarozzi, A., Mahajan, B., Blackburn, B., Kopf, D., Krishnan, L., and Yoong, J. (2013). Microloans, bednets and malaria: Evidence from a randomized controlled trial. *American Economic Review*.

Tilman, D., Cassman, P. A., Matson, P. A., Naylor, R., and Polasky, S. (2002). Agricultural sustainability and intensive production practices. *Nature, 418*:671-677.

Train, K. E. (2009). Discrete choice methods with simulation. Cambriddge University Press.

Tredoux, G. (2004). Nitrate and associated hazard quantification and strategies for protecting rural water supplies: Report to the Water Resources Commission. WRC Report No. 1058/1/04.

United Nations Development Program (UNDP) (2006). *Human Development Report: Beyond Scarcity: Power, Poverty and the Global Water Crisis.* Rome, Italy: Kevin Watkins (lead author). Retrieved from: http://hdr.undp.org/en/content/human-development-report-2006.

United Nations Water. (2011). Water Quality: Policy Brief.
Retrieved from: http://www.unwater.org/downloads/waterquality_policybrief.pdf

U.S. Geological Survey. (2013). Metal prices in the United States through 2010: U.S. Geological Survey Scientific Investigations Report 2012–5188, p. 204.

Van der Ploeg, F. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature 49*(2): 366-420.

Van Niekerk, H., Silberbauer, M. J., and Hohls, B.C. (2009). Monitoring programme revision highlights long-term salinity changes in selected South African rivers and the value of comprehensive long-term data sets. *Environmental Monitoring and Assessment, 154* (1-4):401-411.

Vitousek, P. M., Naylor, R., Crews, T., David, M. B., Drinkwater, L. E., Holland, E., Johnes, P. J., Katzenberger, J., Martinelli, L. A., Matson, P. A., Nziguheba, G., Ojima, D., Palm, C. A., Robertson, G. P., Sanchez, P. A., Townsend, A. R., and Zhang, F. S. (2009). Nutrient inbalances in agricultural development. *Science, 324:* 1519-1520.

Water Resources Group 2030. (2009). Charting our water future: Economic frameworks to inform decision-making. Retrieved from:
http://www.2030waterresourcesgroup.com/water.

Walmsley, R. D. (2003). Development of a strategy to control eutrophication in South Africa: a review and discussion document. Water Quality Management Series (Project 1: Phase 1). Department of Water Affair and Forestry.

Wick, K., and Bulte, E. (2009). The curse of natural resources. *Annual Review of Resource Economics (1)*: 139-156.

Winchester, D., Huskins, J., and Ying, J. (2009). Agrichemicals in surface water and birth defects in the United States. *Acta Paediatrica 98*: 664–669.

The World Health Organization. (2012). Estimates for the use of Improved Drinking Water Sources: South Africa. The World Health Organization.

Appendix A

Appendix to Chapter 1

A.1 Dam Data

As Figure A.1 shows, South Africa has more dams than any other country in Africa. For this analysis, dam information was obtained from the Dam Safety Office within the Department of Water Affairs.[1] About 15 percent of dams were dropped because they were missing critical information on the dam completion form, which could not be verified (690/4,830). For almost all dams (658/690), the missing critical information was date of completion. Based on conversations with the DWAF and google searches to validate information, these dams are generally old dams constructed before the period of this analysis. Some may be no longer functional, although I cannot confirm this. Of the rest, 19 dams were missing information on dam purpose and the remaining 13 were duplicates or had missing or conflicting location information. Of those dropped, 499 reported a purpose of "irrigation." The resulting dataset included 4,140 dams, which were restricted to the subset of irrigation dams, 3,254. These dams were matched to magisterial districts using their GPS coordinates. Figure A.2 presents dam construction by year. Dam construction surged in the 1980s, but fell in recent years.

[1] Data is available online: http://www.dwaf.gov.za/DSO/Default.aspx. The data was accessed on August 23 2013, including all dams listed as of March 2013. Correspondence with the Department of Water Affairs confirmed data validity.

Figure A.1: *Dams in Africa as of 2007*

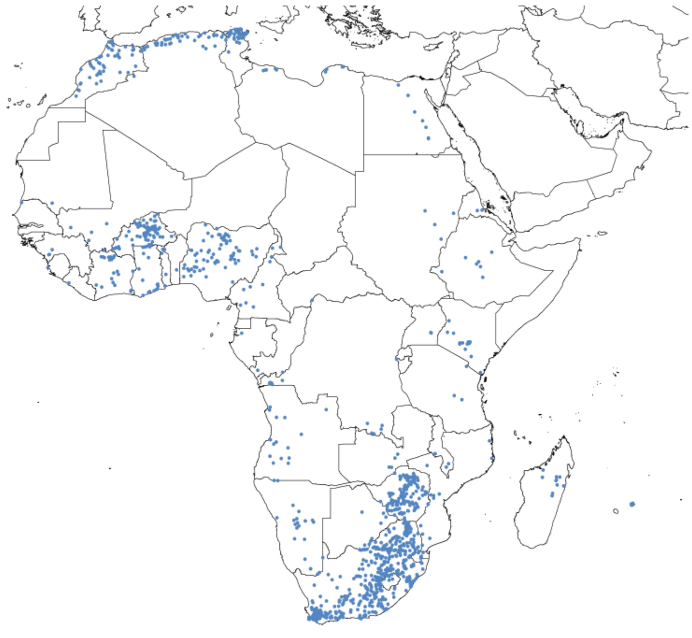

Source: FAO, 2007

Figure A.2: *Dam Construction by Year and Purpose*

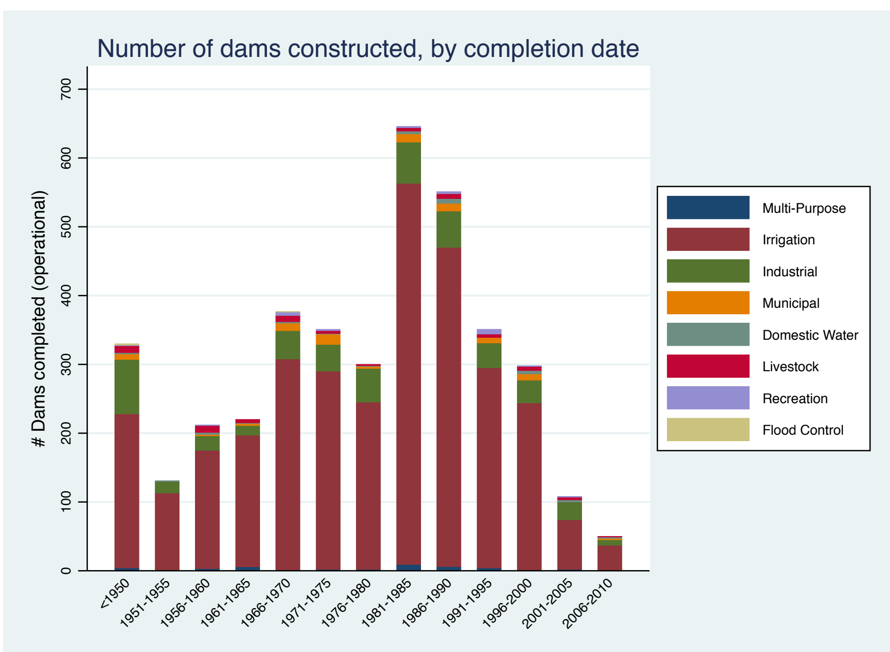

Table A.1: *Irrigation Dam Summary Statistics*

Characteristics	Dams (as of 1980) Mean [SD]	Dams (built 1980-2010) Mean [SD]	Dams (as of 2010) Mean [SD]
Wall Height (meters)	10.989	10.792	10.896
	[8.331]	[6.434]	[7.393]
Capacity ('000 cubic meters)	11208	2181	6428
	[172966]	[22415]	[121.086]
Catchment Area* (sq km)	2849	1148	2026
	[21833]	[16762]	[19671]
Surface Area (sq km)	3.678	0.288	1.913
	[51.786]	[2.418]	[36.061]
Large Size (>30 m)	0.039	0.016	0.027
	[0.195]	[0.127]	[0.162]
Multiple Reported Purposes (1=yes)	0.043	0.039	0.041
	[0.204]	[0.194]	[0.199]
Embankment Construction (1=yes)	0.936	0.96	0.948
	[0.245]	[0.197]	[0.222]
Used for Domestic Water Supply (1=yes)	0.023	0.012	0.018
	[0.150]	[0.111]	[0.131]
Number of Dams	1,571	1,762	3,254

Notes:
1. The above summary statistics refer to any dam that has a listed purpose as "irrigation."
2. Most of the size data was available for all dams. The exception is catchment size, which is only reported for about 20% of dams (601/3,254).

A.2 Water Quality Indicators

The water quality data is collected and managed by the South African Department of Water Affairs. Each observation includes station number, GPS coordinates, date and time, and readings for approximately 20 different water quality measures. At each station, the following indicators are collected: sulfates, nitrates, total dissolved solids (TDS), pH, electrical conductivity, fluoride, chloride, orthophosphate, sodium, potassium, calcium, magnesium, ammonium, and silica. I report those indicators that were most likely to be directly affected by agricultural activity. TDS is also a composite of several salt ions, so it moves roughly the same as calcium, magnesium, sodium, chlorine. Because water quality indicators are often directly linked to each other, almost all indicators move when dams are constructed and joint tests indicate overall water quality changes.

I construct my district-year measurements as follows. I take the median reading per day (average if there are two readings). I then average readings across the month. Blanks are dropped. I average readings across months to get year averages. In the end, water quality data exists for 267 of the 354 districts. Missing districts largely have low populations and/or no rivers. Although data collection was consistent during this period, occasional duplicate or missing values exist and can be observed through sample size variation. This is usually the result of equipment and/or human errors. Table A.2 provides summary statistics for water quality indicators.

Table A.2: *Descriptive Statistics of Chemical Indicators*

Indicator	Mean (sd)	Min	10%	90%	Max	Drinking Standard	N
Nitrate and Nitrite as N	0.80 (3.04)	0.01	0.04	1.92	219.52	10	24,335
Sulfate as SO4	96.23 (275.24)	0.25	3.34	224.50	14,187.45	200	23,714
Dissolved Major Salts	574.10 (1898.3)	8.67	49.14	1,003.99	44,435.80	1,000	20,371
Potassium	(5.39) 19.91	0.05	0.58	9.99	498.18	100	21,997
Chloride as Cl	172.62 (971.05)	0.53	4.67	192.20	36,020.00	1,200	23,333
Phosphate	0.26 (5.33)	0.00	0.01	0.41	825.00	NA	24,854

Notes:
1. This data describes the distribution of station-year averages. The mean value averages all station-year mean values, the maximum is the highest of all station-year mean values, and the minimum is the lowest of all station-year mean values.
2. The concentrations are listed in mg/l units.
3. The Drinking Standards are those listed in the South Africa Department of Water Affairs Guidelines (DWAF, 2006).

A.3 Geographic Variables

The data includes several geographic control variables. The river and land gradient measures are constructed using elevation data from ArcGIS. Geographic information systems, like ArcGIS, provide elevation data derived from the Shuttle Radar Topography Mission (SRTM). The elevation data is provided in a raster (pixel) format. I use the ArcGIS tool, Slope (3D analyst) to construct the approximate slope at each pixel using the raster elevation data. I then use this to construct the average district slope and river gradient slope (the latter is restricted to pixels along the river network provided by the Department of Water Affairs). The other geographic controls, including river length and district area, are also constructed based on magisterial districts using ArcGIS.

A.4 Homeland District Mapping

Former homeland regions make up some fraction of a magisterial district. The map of homeland regions was obtained from the Department of Water Affairs and is shown in Figure 1.1. The analysis defines any district that is at least 30 percent former homeland area as a "former homeland district." Figure A.3 shows that most districts either had either no homeland land or a significant portion.

Figure A.3: *Percent Homeland Area in Each Magisterial District*

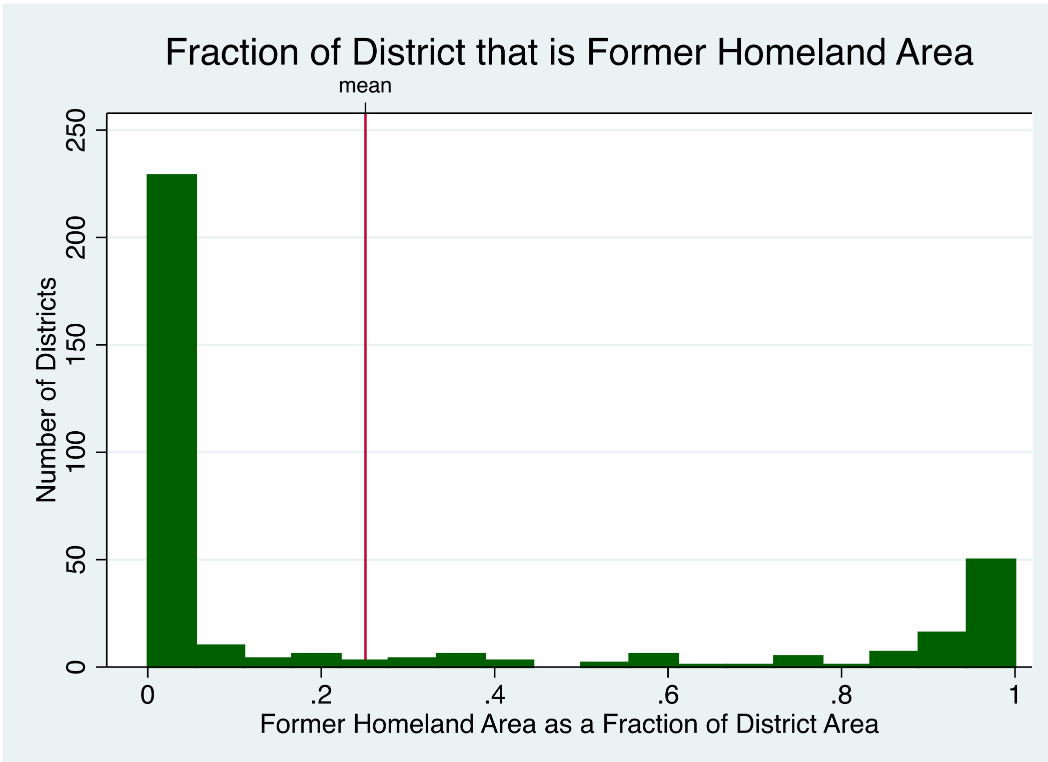

A.5 Demographic and Health Data

The analysis pools three rounds of Demographic and Health Survey (DHS) data: 1987, 1998, and 2003. The sample includes observations from 319 out of 354 districts. For each round, I drop all infants born during the year of surveying, given that the data is censored for these children. For the 1987 survey, I also drop 56 infants because the district could not be identified. For the 1987 data, I only use the previous two children, given only these observations had reliable information. For the 1998 and 2003 survey, each mother has an average of 2.7 children and I use all infants reported. Despite restricting the 1987 data to the previous two children, more mothers were sampled in this round so it makes up roughly one-third of the data.

There are some concerns that the 1987 survey data round under-reported responses from mothers in the homelands. This is a limitation of all South African survey data collected prior to 1994. Rigorous quality testing of the data suggests it is of "reasonable quality" compared to other international DHS surveys (Phillips, 1999). For the main analysis, the IV identification relies on changes in 1994 and 1998. As a result, the 1987 survey round does not identify the effect, but rather provides statistical precision. The analysis can be run without this survey round and the main results look similar: coefficients on dams are the same magnitude and direction, but slightly smaller and significant at 10 percent.

I also ran the regression of infant mortality on dams for each survey round separately. I obtain positive coefficients on dams within the district and smaller coefficients on upstream dams when I use any single survey, though results are less precise given the smaller sample.

Below are definitions of the DHS variables used in the main text:

1. **Employed:** An indicator equal to one if the mother answered yes to having any occupation or to her partner having any occupation (including unskilled labor). If there is no response for the partner, I examine only the mother.

2. **Female HH:** An indicator equal to one if the "Sex of head of household" is reported to be "Female".

3. **Urban:** An indicator equal to one if the enumeration area is urban or semi-urban.

4. **Electricity:** An indicator equal to one if the mother answers yes to "Does your household have electricity?".

5. **Piped:** An indicator equal to one if the mother answered "piped to residence" or "piped to yard" to "What is the main source of drinking water for members of your household?".

6. **Education:** A variable from 1-10 for the education standard obtained by the mother. Any education level above 10 was recoded as a 10, given that this was the highest obtainable response in the 1987 survey round.

A.6 Census Data

The analysis relies on two Census rounds (1996, 2001) to explore migration and district compositional changes. Each Census round includes a ten percent sample of the total population. For the migration analysis, I use the 2001 Census, which contains variables on previous residence. Below are definitions of the Census variables (at the individual adult level) used to describe non-migrants and migrants in Table 1.12:

1. **Population:** District-level population constructed using estimates provided by the census.

2. **Employed:** Reports any employment. The way this question was asked varied slightly between 1996 and 2001. The 1996 Census asked simply: Does the person work?. Thus, responses could have been formal or informal labor, including on one's own land. The 2001 Census more explicitly asked: Did the person do any work for pay, profit or family gain for one hour or more?. Because of the change in the way this question was asked, the first-difference results in Appendix Table A.4 should be interpreted cautiously.

3. **High school (HS) educated:** Education of at least "matric" level. This is short for "matriculation" and it is roughly equivalent to grade 12, standard 10, or form 5.

4. **Piped:** Primary domestic water source is piped into dwelling or onto site.

5. **Below poverty line:** Monthly derived household income is 6000 rand or less.

In Appendix Table A.4, I use Census household data to construct district-level characteristics. These district-level characteristics are defined as above, but use household responses. For variables that vary by household member (e.g., employment), the status of the household head is used.

A.7 October Household Survey Data

I pool five rounds of October Household Survey Data to obtain information on water-related choices and detailed household labor activities. The water-related variables were constructed using household-level responses. The employment-related variables were constructed using individual-level responses for household adults (anyone over 15 years old). The regressions are weighted using household weights, which are based on the inverse of the probability a household is chosen. The variables used in Table 1.9 and Table 1.10 are defined below:

1. **Main Source of Domestic Water**: A response variable to: "What is this household's main source of water?" This variable could take one of eleven choices, including piped, public tap, dam, or river.

2. **Dist. >= 200m to Water**: A binary response variable equal to one for any household that responded 200 meters or more to the question: "If the water source is outside the dwelling, how far is the water source from the dwelling?"

3. **Pay for Any Water:** A binary response variable equal to one if the household answered yes to the question: "Does the household have to pay for any water?".

4. **Enough Water:** This question was only asked in 1995-1996. An indicator equal to one was created if a household answered "Always" or "Mostly yes" to the question: "Is the water obtained enough for household purposes?"

5. **Currently Working:** A binary variable equal to one if an adult is currently working, or currently has a job even if sick in the past week. Work could include: "Any work for pay, profit or family gain? For example, formal work for a salary or wage; informal work such as making things for sale, selling things or providing a service; work on a farm or land, whether for a wage or as part of household's farming activities".

6. **Working in the Last Year:** A binary variable equal to one if the person has worked in the last year. The same definition of work as above.

7. **Total Hours in Last 7 Days:** A numerical write-in response to: "What is the total number of hours that (the person) actually worked per week?".

8. **Salary Below Poverty:** Variable based on the question, "What was total salary at main job (including overtime and bonus, before deductions)?". Given skip patterns, no response was given if an individual was unemployed. For those adults, the salary was recoded as zero if the adult reported unemployment and therefore did not answer the salary question. The binary "salary below poverty" indicator was generated from this question, and it is equal to one if the reported salary is less than 6,000 rand per month.

A.8 Robustness Results for Subsets of the Population

Table A.3 presents disaggregated results for population subsets. Dams generate positive and significant infant mortality increases for those in rural areas and for those using public taps. Similarly, the infant mortality result disappears for those relying on piped water. These results should not be interpreted as causal, however, because each subset is a selected sample. In other words, households using piped water are also more likely to be wealthier, better educated, and different along other unobservable dimensions. In addition, the smaller sample size, particularly for those that use river sources, makes the results less precise.

Although the survey does not provide a variable for migration status, it asks mothers if they have been in their residence for more than two years. For this subgroup, the effect remains positive but imprecisely estimated. While the sample is small and I cannot distinguish the effect from zero, it is still positive. In Section 1.6, I conduct more in depth analysis using Census data.

Table A.3: *Infant Mortality Results by Type*

Dependent Variable: Infant Mortality					
Subset:	Residence for >2yrs	Rural	Piped	Public tap	River
	(1)	(2)	(3)	(4)	(5)
Panel A. Instrumental Variables					
Dams in District	0.00171	0.00229*	-1.88e-05	0.00215*	0.00682*
	(0.00149)	(0.00121)	(0.00141)	(0.00119)	(0.00395)
Dams Upstream	-0.000493	-0.000471	6.88e-05	0.000570	-0.00149
	(0.000575)	(0.000537)	(0.000390)	(0.00111)	(0.00164)
Panel B. OLS					
Dams in District	0.00183	0.00315**	0.00103	0.00265	0.00606***
	(0.00226)	(0.00133)	(0.00129)	(0.00252)	(0.00231)
Dams Upstream	-0.000448	1.29e-05	0.000240	-0.000296	0.000209
	(0.00132)	(0.000537)	(0.000357)	(0.000920)	(0.000850)
Controls	Yes	Yes	Yes	Yes	Yes
Observations	11,367	17,084	14,152	9,299	7,031
Mean Dependent Variable	0.049	0.056	0.030	0.058	0.069

Notes:
1. The outcome variable is based on self-reported birth date and death dates from 1980-2002.
2. Column (1) restricts the data to those who report living in their residence for greater than two years. This is the closest proxy to non-migrant status available in the data. It is only available in the data from 1987-2003. Column (2) restricts the data to those living in rural areas. Column (3), (4), and (5) restrict the sample based on the household's primary domestic water source (piped to house/yard, public tap, or river).
3. The IV regressions include the following control variables, each interacted with the policy variable: total length of seasonal and perennial rivers, total district area, district slope category (0-1.5%, 1.5-3%, 3-6%, >6
4. The OLS regressions include district and year fixed effects.
5. Robust standard errors clustered at the district level are reported in parentheses. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.
6. Infant mortality data is from the South Africa Demographic and Health Surveys (1987, 1998, 2003).

A.9 Census First Difference Results

Table A.4 presents first difference regression results using the 1996 and 2001 Census. The IV specification is as shown in estimating equations (3) and (4), except that the outcome variable is the change in district characteristic between 1996 and 2001. Similarly, the OLS specification is below:

$$\Delta y_d = \gamma_o + \gamma_1 \Delta D_d + \gamma_2 \triangle D_d^U + \gamma_3 \left(\Delta D_d * H_d\right) + \gamma_4 \left(\Delta D_d^U * H_d\right) + H_d + X_d + \lambda_p + \varepsilon_d$$

where Δy_d is the change in fraction of the district which a given characteristic, ΔD_d is the change in number of dams, and ΔD_d^U is the change in number of dams upstream. As in the main analysis, H_d indicates whether a district is a former homeland and λ_p is province fixed effects. The limitation with this analysis is that I cannot control for district-level trends. To alleviate this concern, I add X_d to control for district characteristics in both the IV and OLS. The controls are based on 1996 district-level characteristics, including population, female-headed households, fraction of high school educated household heads, fraction of households with piped water, and fraction below the poverty line indicator.

The results show that districts receiving more dams experience relative population declines and employment increases. However, the fraction directly employed in agriculture does not increase substantially. The fraction piped goes up slightly (but imprecisely) and the fraction educated goes down slightly. These small changes do not suggest major shifts in the district composition associated with migration.

Table A.4: *District-Level Changes (First Difference)*

Dependent Variable:	Δ log(Population)	Δ Fraction Employed	Δ Fraction Ag. Employ	Δ Fraction Piped	Δ Fraction Educated
	(1)	(2)	(3)	(4)	(5)
Panel A. Instrumental Variables					
Dams in District	-0.0672**	0.0160*	-0.00205	0.0128	-0.00253*
	(0.0315)	(0.00845)	(0.00464)	(0.00789)	(0.00150)
Dams Upstream	0.0354*	-0.0101*	-0.000564	-0.00943*	0.00160
	(0.0211)	(0.00551)	(0.00275)	(0.00573)	(0.00120)
Panel B. OLS					
Dams in District	-0.0217	0.00581	0.00275	0.00622	-0.000222
	(0.0189)	(0.00678)	(0.00317)	(0.00411)	(0.000923)
Dams Upstream	0.0111	-0.00677	-0.00244	-0.00733*	-0.000753
	(0.0160)	(0.00601)	(0.00241)	(0.00382)	(0.000794)
Dams in District	0.0186	-0.00488	0.00120	-0.0175	0.000539
*Homeland	(0.0314)	(0.0122)	(0.00381)	(0.0169)	(0.00224)
Dams Upstream	-0.00500	0.00917	-0.00108	0.0216	0.00116
*Homeland	(0.0233)	(0.00933)	(0.00306)	(0.0139)	(0.00169)
Controls	Yes	Yes	Yes	Yes	Yes
Observations	354	354	354	354	354
Mean Dep. Variable:	0.1780	-0.3382	-0.01266	0.0282	-0.1253

Notes:
1. Each observation is a magisterial district. The outcome variable is the change in the fraction of the surveyed district individuals with a given characteristic. The first column is the change in the logged district population. The variables in the subsequent columns are: any employment, agricultural employment, piped water, and high school education.
2. For the IV and OLS, I add district-level controls from the 1996 Census: district population, fraction of female-headed households, fraction of high school educated household heads, fraction of households with piped water, and fraction below the poverty line indicator. I include province fixed effects in the OLS and IV.
3. The IV regressions include geographic controls: total river length, elevation category, slope category, total district area, upstream total river length, upstream slope category, upstream elevation category, and upstream total area. Former homeland controls include a categorical variable for the fraction homeland and a categorical variable for fraction upstream homeland.
4. Robust standard errors clustered at the district level are in parentheses. Asterisks denote significance: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

A.10 Employment Outcomes by Gender and Age

The below tables evaluate the effect of dams on employment, broken out by gender in Table A.5 and age of woman in Table A.6. The results demonstrate that while employment by women increases in response to dams, the increases are lower for women under thirty and near zero for women over fifty, the two segments of women that are likely to be primary caregivers.

Table A.5: *Employment Outcomes by Gender*

	MEN ONLY				WOMEN ONLY			
Outcome Variable:	Currently Employed	Employed in the Last Year	Total Hours in Last 7 Days	Salary Below Poverty	Currently Employed	Employed in the Last Year	Total Hours in Last 7 Days	Salary Below Poverty
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Instrumental Variables								
Dams in District	0.0204	0.00318	1.165	-0.0451*	0.0613***	0.0524*	3.474**	-0.0675***
	(0.0218)	(0.0331)	(1.755)	(0.0233)	(0.0236)	(0.0316)	(1.452)	(0.0231)
Dams Upstream	0.0357	0.0665	1.486	0.0233	0.00208	0.0367	-1.200	0.000975
	(0.0277)	(0.045)	(1.371)	(0.0369)	(0.0305)	(0.0375)	(1.013)	(0.0314)
Panel B. OLS								
Dams in District	0.00927**	0.0102**	0.671***	-0.00940***	0.00680*	0.0105**	0.383***	-0.00847***
	(0.00368)	(0.00406)	(0.189)	(0.00298)	(0.00352)	(0.00455)	(0.126)	(0.00287)
Dams Upstream	0.00504***	0.00577***	0.237***	-0.00555***	0.00485***	0.00620***	0.183*	-0.00389**
	(0.00163)	(0.00164)	(0.0841)	(0.00202)	(0.00177)	(0.00184)	(0.0998)	(0.00168)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	157,327	157,327	156,475	157,327	191,259	191,259	190,462	191,259
Mean Dep. Var.	0.435	0.455	20.289	.742	0.255	0.274	10.695	.873

Notes:

1. Each observation is an individual 15 years or older, within households surveyed by the OHS.
2. The IV regression includes all geographic controls for the district and upstream, steep gradient*year fixed effects, and district fixed effects.
3. The OLS controls include district and year fixed effects.
4. Robust standard errors clustered at the district level. Asterisks denote significance: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.6: *Women Employment Outcomes by Age*

	Women Only								
	Currently Employed			Total Hours Worked in Last 7 Days			Salary Below Poverty		
Age of woman:	<30	30-50	>50	<30	30-50	>50	<30	30-50	>50
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A. Instrumental Variables									
Dams in District	0.0627***	0.124***	0.00361	3.274**	6.421**	0.163	-0.0482**	-0.136***	-0.0219
	(0.0235)	(0.0394)	(0.0299)	(1.452)	(2.657)	(1.470)	(0.0233)	(0.039)	(0.018)
Dams Upstream	0.0150	-0.0297	-0.0177	0.0629	-3.495*	-1.582**	-0.0212	0.0319	0.0131
	(0.0273)	(0.0554)	(0.0387)	(0.927)	(2.087)	(0.754)	(0.0335)	(0.0515)	(0.023)
Panel B. OLS									
Dams in District	0.0114***	0.00548	0.00334	0.553***	0.269	0.0726	-0.0121***	-0.00745*	-0.0219
	(0.00325)	(0.00468)	(0.00251)	(0.154)	(0.198)	(0.114)	(0.00295)	(0.00382)	(0.00222)
Dams Upstream	0.00442	0.00521*	0.00572***	0.109	0.158	0.317***	-0.00161	-0.00558***	-0.00497***
	(0.0027)	(0.00286)	(0.00159)	(0.141)	(0.127)	(0.119)	(0.00183)	(0.00209)	(0.00173)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	79,727	67,990	41,723	79,539	67,503	41,614	79,727	67,990	41,723
Mean Dep Var	.161	.431	.149	6.81	18.10	6.11	.918	.784	.936

Notes:
1. Each observation is a female 15 years or older, surveyed by the October Household Survey.
2. The IV regression includes all geographic controls for the district and upstream, steep gradient*year fixed effects, and district fixed effects.
3. The OLS controls include district and year fixed effects.
4. Robust standard errors clustered at the district level. Asterisks denote significance: *** p<0.01, ** p<0.05, * p<0.1.

Appendix B

Appendix to Chapter 2

B.1 Conceptual Model

This appendix includes the formal proof of Propositions 1 through 4 in the main text. We start by characterizing agents' decisions and types in a more formal way.

B.1.1 Expected Value of Take-Up

The expected net benefit of take-up (which appears in the take-up decision inequality, equation (1)) can be rewritten as

$$\mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) = \Pr(R - F_0 - F_1 > 0|F_0) \times \left[R - F_0 - \mathbb{E}_{F_1|F_0}(F_1|R - F_0 - F_1 > 0)\right],$$

where $\Pr(R - F_0 - F_1 > 0|F_0)$ indicates the type-specific (i.e., conditional on F_0) probability of follow-through and $R - F_0 - \mathbb{E}_{F_1|F_0}(F_1|R - F_0 - F_1 > 0)$ is the net benefit, conditional on follow-through.

B.1.2 Adoption Types

Under the distributional assumptions stated in the main text:

- $F_0 \perp F_1$
- F_1 takes one of two values: $F_1 = \{f_L, f_H\}$, with $f_L < f_H$, and $\mathbb{E}_{F_1|F_0}(F_1) = g_1(f_L)f_L +$

$g_1(f_H)f_H$, where $g_1(.)$ is the probability mass function of F_1

- F_0 is continuously distributed across agents with cumulative distribution function $G_0(.)$,

we can classify agents in three follow-through types: non-adopters, contingent adopters and always adopters.

Non-Adopters Non-adopters are characterized by the condition on F_0,

$$R - F_0 < f_L \tag{B.1}$$

such that even when the realization of F_1 is low (f_L), their net benefit of follow-through is negative. The share of non-adopters is given by $1 - G_1(R - f_L)$. Their probability of follow-through is always 0 and so is their expected private benefit. Non-adopters take-up only if $c - A > 0$, or if the subsidy exceeds the cost of take-up. Note that even when they take-up (purchase the technology), they never follow-through.

Contingent Adopters Contingent adopters are characterized by the condition

$$f_L < R - F_0 < f_H. \tag{B.2}$$

Contingent adopters follow-through when the realization of F_1 is F_L, but not when the realization is F_H. The share of contingent adopters is given by $G_0(R - f_L) - G_0(R - f_H)$, with expected private benefit given by

$$\mathbb{E}_{F_1|F_0}\left[\max(R - F_0 - F_1, 0)|R - F_0 - F_1 > 0\right] = g_1(f_L)\left(R - F_0 - f_L\right)$$

where $g_1(F_L)$ is their probability of follow-through. The take-up decision of these agents is characterized by condition $F_0 \leq R - f_L - \frac{c-A}{\delta g_1(f_L)}$.

Always Adopters Always adopters are characterized by the condition

$$f_H < R - F_0. \tag{B.3}$$

Hence, they follow-through whether the draw of F_1 is f_L or f_H: $\Pr(R - F_0 - F_1 > 0|F_0) = 1$. The share of always adopters is given by $G_0(R - f_H)$, and their private benefit given by

$$\mathbb{E}_{F_1|F_0}[\max(R - F_0 - F_1, 0)|R - F_0 - F_1 > 0] = R - F_0 - \mathbb{E}(F_1).$$

They take-up only if $F_0 < R - \mathbb{E}(F_1) - \frac{c-A}{\delta}$.

Selection and Follow-Through

Conditions (B.1), (B.2), and (B.3) determine thresholds over the support of F_0 that delimit the shares of always adopters, contingent adopters and never adopters for a given distribution of F_0. Figure 1 illustrates these thresholds on the probability density function of F_0, $g_0(F_0)$. Note that the bell shaped distribution for F_0 shown in Figure 1 is not a necessary assumption of the model, and is used only to visualize the shares of each agent type as the area under the curve delimited horizontally by the thresholds in grey: $R - f_H$ and $R - f_L$.

The thresholds in black correspond to the take-up decision for each agent type. The take-up threshold for contingent adopters, $R - f_L - \frac{c-A}{\delta g_1(f_L)}$, is always to the right of the threshold for contingent adopters provided that the subsidy, A, is less than or equal to the total cost of the technology, c. Hence, the bigger the subsidy, A, the bigger the share who take-up, but follow-through only if $F_1 = f_L$. The take-up threshold for always adopters, $R - \mathbb{E}(F_1) - \frac{c-A}{\delta}$, may be to the left or to the right of the threshold, $R - f_H$, which defines the group of always adopters. If $\frac{c-A}{\delta} \leq f_H - \mathbb{E}(F_1)$, all always adopters will take-up. However, if $\frac{c-A}{\delta} > f_H - \mathbb{E}(F_1)$, a bigger subsidy may increase take-up among always adopters.

In sum, the subsidy A affects follow-through rates conditional on take-up by determining the shares of always adopters and contingent adopters that take up. When the subsidy is small, such that $A < c - \delta(f_H - \mathbb{E}(F_1))$, not all always adopters take-up. When the subsidy is between $c - \delta(f_H - \mathbb{E}(F_1))$ and c all always adopters take-up, but just a fraction of contingent adopters take-up. For subsidies larger than c, all always adopters, all contingent adopters and some non-adopters take-up.

Proposition 1 *Follow-through conditional on take-up increases as a function of the total (potentially subsidized) take-up cost.*

Conditional adopters are the population of interest for understanding the relationship between uncertainty and technology adoption: they constitute the only group whose follow-through decision is affected by the shock realization. The share of conditional adopters who take-up is given by

$$\frac{\mathrm{g}_1(f_L)\left[\mathrm{G}_0\left(R - f_L - \frac{c-A}{\delta g_1(f_L)}\right) - \mathrm{G}_0(R - f_H)\right] + \mathrm{G}_0\left(R - f_H\right)}{\mathrm{G}_0\left(R - f_L - \frac{c-A}{\delta g_1(f_L)}\right)} \tag{B.4}$$

if $\frac{c-A}{\delta} < f_H - \mathbb{E}(F_1)$ and is 100 percent if $\frac{c-A}{\delta} \geq f_H - \mathbb{E}(F_1)$. These two expressions show how follow-through depends on A through the take-up decision of the different types of agents: the larger the subsidy, A, the larger the share of contingent adopters that take-up, reducing the overall rate of follow-through among those who take-up.

Proposition 2 *An increase in uncertainty reduces follow-through conditional on take-up.*

Note that in expression (B.4), an increase in the spread of F_1 (distance between f_H and f_L) results in a bigger increase in the denominator than in the numerator, since part of the numerator is multiplied by $\mathrm{g}_1(f_L)$, which is a number between 0 and 1. Hence, uncertainty worsens follow-through conditional on take-up.

Proposition 3 *An increase in uncertainty weakens the relationship between take-up cost and conditional follow-through shown in Proposition 1.*

Uncertainty increases the share of contingent adopters. This is easy to see since the share of contingent adopters is determined by the probability mass over the support of F_0 between $R - f_H$ and $R - f_L$. The greater the spread of F_1, the bigger the share of contingent adopters, and the less the take-up decision predicts follow-through. In the extreme case of no uncertainty, there are no contingent adopters ($f_H = f_L$) and all we have is either always adopters or never adopters. In this case, A increases take-up among always adopters, but does not lower follow-through conditional on take-up unless adopters are paid to take-up the technology ($A > c$).

Option Value of the Contract

The option value associated with the take-up decision when agents are free to follow-through or not at time 1, i.e. under limited liability, is given by

$$OV(F_0) = \mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) - \max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right) \tag{B.5}$$

with $\mathbb{E}_{F_1|F_0}(R - F_0 - F_1) = R - F_0 - \mathbb{E}(F_1)$, where $\max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right)$ represents the expected profit associated with making the follow-through decision at time 0, or the value of the static contract. Note that for non-adopters, the decision to not follow-through does not change with new information. Hence, $\mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) = \max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right) = 0$. Similarly, always adopters' decision does not change with new information. Hence, $\mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) = \max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right) = R - F_0 - \mathbb{E}(F_1)$. Therefore, the only group with a positive option value is contingent adopters. For them, $\mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) > \max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right)\mathbb{E}_{F_1|F_0} \max(R - F_0 - F_1, 0) > \max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right)$, since

$$\mathbb{E}_{F_1|F_0}(R - F_0 - F_1) = R - F_0 - \mathbb{E}(F_1).$$

The share, $G_0(R - \mathbb{E}(F_1)) - G_0(R - F_H)$ would take-up under a contract with commitment (i.e., a static contract where take-up and follow-through decisions are made simultaneously at time 0) since their expected benefit under commitment, $R - F_0 - \mathbb{E}(F_1)$, is greater than zero. The share $G_0(R - F_L) - G_0(R - \mathbb{E}(F_1))$ would only take-up in the contract without commitment, since their expected benefit under commitment, $R - F_0 - \mathbb{E}(F_1$, is less than zero. Hence, for contingent adopters,

$$\max\left(\mathbb{E}_{F_1|F_0}(R - F_0 - F_1), 0\right) = \begin{cases} R - F_0 - F_1 & \text{if} \quad F_0 < R - \mathbb{E}(F_1) \\ 0 & \text{if} \quad F_0 > R - \mathbb{E}(F_1) \end{cases}$$

From the definition in (B.5), it follows that the option value for contingent adopters with

$F_0 < R - \mathbb{E}(F_1)$ is given by

$$\begin{aligned} g_1(f_L)\,(R - F_0 - f_L) - & \quad (R - F_0 - \mathbb{E}(F_1)) \\ &= \quad [g_1(f_L) - 1]\,(R - F_0) + g_1(f_H) f_H \\ &= \quad g_1(f_H)\,(f_H + F_0 - R)\,; \end{aligned}$$

while the option value for contingent adopters with $F_0 > R - \mathbb{E}(F_1)$ is equal to their expected private benefit without commitment: $g_1(f_L)\,(R - F_0 - f_L)$.

In summary, the option value as a function of F_0 is given by

$$OV(F_0) = \begin{cases} 0 & \text{if} \quad F_0 > R - f_L \\ g_1(f_L)\,(R - F_0 - f_L) & \text{if} \quad R - \mathbb{E}(F_1) < F_0 \leq R - f_L \\ g_1(f_H)\,(f_H + F_0 - R) & \text{if} \quad R - f_H < F_0 \leq R - \mathbb{E}(F_1) \\ 0 & \text{if} \quad F_0 \leq R - f_H \end{cases}$$

Proposition 4

The option value associated with take-up is increasing in uncertainty, which results in higher take-up at all take-up cost levels.

For a given agent with $F_0 = f_0$, option value increases with uncertainty. As uncertainty increases (the distance between f_H and f_L), so does the likelihood that $R - f_H < f_0 \leq R - f_L$, which in turn increases the likelihood that the agent becomes a contingent complier. Hence, as uncertainty increases, the share of agents with a positive option value from take-up also increases. As expected, the option value has an asymmetric relationship with the upper and lower bounds of the shock distribution. One can increase the option value indefinitely by lowering f_L (which is equivalent to increasing the realization of the positive shock, since f_L enters as a cost in the profit function). However, lowering f_H leads to an increase in the option value up to the point where $R - \mathbb{E}(F_1) < f_0$; beyond this, the option value remains constant and equal to $g_1(f_L)\,(R - F_0 - f_L)$, which is equal to the expected private benefit of the contract to contingent adopters.

As a function of R, the option value for a given individual with $F_0 = f_0$ is zero up to the point where $R - f_L$ is larger than f_0. Beyond this, the agent becomes a contingent adopter and the option value is increasing with R up to $R = f_0 + \mathbb{E}(F_1)$, where it peaks and then falls up to $R = f_0 + f_H$. After this, the option value becomes 0 again since the value of R is large enough to guarantee follow-through.

B.2 Estimation

The estimation of the model outlined in the main text is done via simulated maximum likelihood.[1] This appendix details the estimation procedure used to recover the structural parameters.

B.2.1 Additional Parameters

Our field experiment design included two additional treatment arms in addition to the ones described: a "surprise reward treatment" group and a monitoring group. In the structural estimation, we modify the profit function to account for the variation in choices that these treatment arms introduce.

Surprise Reward Treatment Half of the farmers who attended training (52.5 percent) were assigned to a "surprise reward treatment" and did not learn about the threshold reward for follow-through ($\geq$ 35 trees) until after their decision to take-up was made. As explained, this treatment arm allows us to explore whether liquidity constraints explain the absence of selection effects in the data. In order to keep track of the information differences at the time of take-up in the estimation, we allow for these individuals to have a separate component in the profit from planting any positive amount of trees (a constant "surprise treatment" effect, α_S). If these individuals had identical beliefs about the costs and benefits derived from the trees (which in practice means that random parameters F_0, F_1 and T were drawn from the same distribution as those in the standard treatment, who learned about the reward before choosing to take-up), the surprise treatment effect would be zero. However, we observe reduced form evidence that there was an expectation of a higher profit among those who did not know about

[1] See Train (2009).

the reward before taking up: their take-up rate is higher than the rate among farmers who received a reward of zero. The average take-up among those in the surprise reward treatment was approximately equal to the the take-up rate of farmers in the standard treatment who drew a reward of ZMK 40,000 before they made their take-up decision.[2] Hence, in our estimation, the surprise treatment is left unrestricted and is estimated to be 91.79 (s.e. 8.11) in the main model and 54.42 (s.e. 10.235) in the model with a mean shift in F. Note that this latter coefficient is close to the reduced form effect.

Monitoring Group A small share of the program participants, 15.8 percent, were randomly selected to receive regular visits to monitor tree-related activities, which allows us to more closely observe time use. This group experienced higher follow-through rates than farmers who were not assigned to the monitoring group,. Though the treatment was not designed to have an impact and the monitors were explicitly told not to communicate information about tree cultivation to the farmers, monitoring may have influenced farmers in a number of ways. For example, monitoring could have increased the subjective value of the trees by making them seem "more important" or decreased the cost of caring for them by periodically reminding individuals of their location and commitment. Farmers were not aware that they would be monitored when they made their take-up decision. In order to account for the observed effect of monitoring in the estimation, we allow the profit of those in the monitoring group to have a separate component that takes the value of zero if no trees are cultivated and of α_M when any positive number of trees is cultivated. This parameter is estimated to be 238.40 (s.e. 36.844) in the main model and 229.53 (s.e. 37.22) in the model with a mean shift in F.

B.2.2 Objective Function Details under Simulated Maximum Likelihood

We use simulation methods to evaluate the objective function, equation (9) (from the main text) , for any given value of the parameters. We use simulated maximum likelihood because there are several quantities in our objective function that do not have a closed form expression.

[2]This calculation is performed from the results of a linear regression of take-up on the reward among those who had knowledge of the reward before deciding to take up.

As is usually done in random parameter models, we integrate away the unobserved random parameters when writing the analytic probabilities for each outcome. These integrals, once more, do not have a closed form solution. Hence, we use numerical integration to write the probability of choosing N trees conditional on parameters $\mu_F, \sigma_{F_0}, \sigma_{F_1}, \mu_T, \sigma_T, \alpha_S,$ and α_M. Before writing the expression for the simulated probabilities of choosing N trees, we note one more aspect of our estimation strategy.

When using simulation methods to estimate discrete choice models with random parameters, numerical integrals are used to approximate theoretical probabilities. This often results in a stepwise as opposed to smooth objective function, since small probabilities are hard to approximate numerically and can be very noisy. In order to smooth the kinks in our objective function, we add an extreme value distributed error term at the end of the profit function. This allows us to compute probabilities between 0 and 1 for each draw of the random parameters, which results in a smoother objective function. Monte Carlo simulations suggest this method will not introduce bias our results provided that we choose a relatively small scale parameter, λ, which we refer to as smoothing factor. In the estimation we use a smoothing factor of 0.5.

Thus, using Train (2009) notation, the simulated probabilities of choosing N trees at $t = 1$ are

$$\check{P}_i(N^* = n|\theta) = \frac{1}{K}\sum_{k=1}^{K} \frac{\exp\left(\frac{1}{\lambda}\Pi(n|F_{0k}, F_{1k}, T_k, R_i)\right)}{\sum_{j=0}^{50}\exp\left(\frac{1}{\lambda}\Pi(j|F_{0k}, F_{1k}, T_k, R_i)\right)}$$

where k indexes each draw of the full random parameter vector, (F_{0k}, F_{1k}, T_k), given the vector of parameters $\theta = (\mu_F, \sigma_{F_0}, \sigma_{F_1}, \mu_T, \sigma_T, \alpha_S, \alpha_M)$, and farmer-specific treatments A_i and R_i.

Similarly, the simulated probability of take-up at $t = 0$ is given by

$$\begin{aligned} \check{P}_i(TakeUp|\theta) &= \\ \tfrac{1}{K}\textstyle\sum_{k=1}^{K} &\mathbf{1}\left(A_i - c + \delta\check{\mathrm{E}}\left[\max_N \Pi(N|T_k, F_{0k}, F_1, T_i, R_i)|\, F_{0k}, T_k\right] > 0\right) \end{aligned} \tag{B.6}$$

where k indexes each draw of the partial random parameter vector, (F_{0k}, T_k), given the vector of parameters $\theta = (\mu_F, \sigma_{F_0}, \sigma_{F_1}, \mu_T, \sigma_T, \alpha_S, \alpha_M)$, and farmer-specific treatments A_i and R_i. Note that the expected profit conditional of random variables F_0 and T, and observed treatments A and R also involves an integral without a closed form solution. We therefore use the simulated version of it in expression (B.6). More specifically,

$$\begin{aligned} \check{\mathbb{E}}\left[\max_N \Pi(N|T_k, F_{0k}, F_1, T_i, R_i)|\, F_{0k}, T_k\right] &= \\ \tfrac{1}{M}\textstyle\sum_{m=1}^{M} \max_N & \quad \Pi(N|F_{0k}, F_{1m}, T_k, R_i, A_i) \end{aligned} \tag{B.7}$$

where m indexes each of M draws from a normal distribution with mean F_{0k} and variance given by $\sigma^2_{F_1}$.

For estimation purposes, we use $K = 1500$ and $M = 100$. Each of the k draws are independent across observations. However, the M draws used in (B.7) are kept constant across observations. This reduces our computing power substantially without affecting the independence assumptions across observations (note that (B.7) is conditioned on F_{0k} and T_k, which are drawn independently for each farmer).

B.2.3 Maximization Algorithm

In order to guarantee that the point estimates correspond to the global maximum of the likelihood function, we first conducted a grid search that would inform our starting values for the numerical maximization. The grid search was conducted over 80 thousand different combinations of the parameters and, to minimize computing time, was conducted with a lower value of K and M (400 and 50 respectively).

In addition, we conducted a three stage recursive maximization (minimization of the negative likelihood) where in each stage we maximized the simulated likelihood along a subset of the parameter vector holding the rest constant. This method worked better than the single step maximization in Monte Carlo simulations. The subsets of parameters in each of the three stages were (μ_T, σ_T, ρ), $(\sigma_{F0}, \sigma_{F1})$, and $(\mu_F, \alpha_M, \alpha_S, \mu_{shift})$ respectively.[3] The three stages were repeated sequentially until a convergence criterion involving changes in the parameter values was reached.[4] In a final stage, we used the resulting parameter estimates as starting values in a single step numerical maximization. This last step yielded small changes in

[3] μ_{shift} corresponds to the common uniform shock in the mean shifter model discussed.

[4] The convergence criterion we used was that the square sum of differences between the new parameters and the starting values (the estimated parameters from the last optimization round) was less than 0.0001. The number of iterations was very robust to the critical value chosen and never reached more than four iterations.

the parameter values (the largest change was less than 6 percent and corresponded to the monitoring parameter, α_M; the second largest change was of 4 percent and corresponded to the standard deviation of F_0, σ_{F0}).

B.2.4 Standard Errors

Standard errors were computed using the variance of the numerically approximated scores, which should converge to the negative of the Hessian in the limit provided that the point estimates are the argmax of the log-likelihood function (Train, 2009). We chose this method instead of the numerical Hessian because it allowed us to choose the size of the step (h) when calculating the numerical score. Simulated methods often result in "roughness" of the likelihood function, which, in our case, led to a non-positive definite numerical Hessian.[5] In order to verify that we were at a (local) minimum, we plotted the likelihood to verify its curvature along each parameter, one at a time.

B.3 Deterministic Tree Survival Assumption

One of the assumptions in the specification of the farmer's optimization problem is that survival of trees is deterministic, conditional on effort. We allow for the cost of tree cultivation to be quadratic in the number of trees, which would capture increasing marginal costs of tree cultivation arising from increasing marginal opportunity cost of time. Our assumption on deterministic survival can be thought of as a two stage optimization process, where the farmer decides on the optimal number of trees to keep alive first, and then allocates the amount of costly effort that guarantees survival to each of those trees. This assumption is less restrictive than one would think.

The two-stage optimization process is roughly consistent with standard optimization under probabilistic survival with a few restrictions: that the probability of tree survival for a single tree as a function of effort, $p(e)$, (a) is independent across trees; (b) attains 1 at some level of

[5] The default numerical gradient calculation also led to gradient components that far from zero. In contrast, all elements of the numerical gradient we "manually" calculated were very close to zero.

effort, $\tilde{e}$, and (c) is a convex function of effort up to $\tilde{e}$; that is $\lim_{e \to \tilde{e}} p'(e) > 0$. In addition we maintain the standard interior solution assumptions of the profit function: (d) increasing and convex cost of effort, $c(e)$ (i.e. $c'(e) > 0$, $c''(e) > 0$), and (e) diminishing marginal returns to the additional tree. We can denote this last assumption as $g_i > g_{i+1}$, where g_i denotes the marginal benefit of the ith tree that survives. Assumption (c) guarantees that the optimal allocation of effort across two or more trees, given an optimal level of total effort $\bar{e}$, is such that the farmer will allocate $\tilde{e}$ to as many trees as possible up to $k\tilde{e} \leq \bar{e}$. If $k\tilde{e} < \bar{e}$, then only the last tree $(k+1)$ will be allocated the remaining effort, $\bar{e} - k\tilde{e}$, making its survival probability less than one. This optimal allocation of effort is thus consistent with deterministic survival of all trees the farmer cultivates, except for possibly the very last tree.[6]

It could be, however, that no amount of effort guarantees the survival of a given tree: i.e., the probability function reaches a maximum of $p\left(\tilde{\tilde{e}}\right) < 1$ at $\tilde{\tilde{e}}$. In order to explore whether such a model fits our data better, we simulate farmer's behavior assuming this is the case. We keep the parameters that govern farmers' heterogeneity and shocks from our estimated mean-shift model, and we add probabilistic survival to the argument of the indicator function for reaching the 35-tree threshold.[7] Table B.1 replicates the reduced form comparison exercise

[6] The proof behind this optimal distribution of effort across trees consists of showing that there are no interior solutions to the optimization problem where more than one tree is allocated an amount of effort between 0 and $\tilde{e}$. We can prove this by contradiction for the case of two trees. The proof can be easily extended to an unlimited number of trees.

The farmer's maximization problem in the case of two trees is given by

$$\max_{e_1, e_2} \pi(e_1, e_2) = g_1 p(e_1) + g_2 p(e_2) - c(e_1 + e_2)$$

where $g_1 > g_2$ (because of assumption (e)), $p(.)$ meets assumptions (a), (b) and (c), and $c(.)$ meets assumption (d).

For a solution to this problem where both trees receive an amount of effort between 0 and $\tilde{e}$ to exist (i.e. $0 < e_1^* < \tilde{e}$ and $0 < e_2^* < \tilde{e}$), the following condition needs to be satisfied

$$g_1 p'(e_1^*) - c'(e_1^* + e_2^*) = g_2 p'(e_2^*) - c'(e_1^* + e_2^*)$$

which can be simplified to

$$g_1 p'(e_1^*) = g_2 p'(e_2^*) l \tag{B.8}$$

Because $g_1 > g_2$, and $p''(e) > 0$ for $0 < e < \tilde{e}$, condition (B.8) requires that $e_1^* < e_2^*$. However, it is easy to see that given a constant total amount of effort, e^*, no optimal distribution of this effort, (e_1^*, e_2^*) will be such that $e_1^* + e_2^* = e^*$ and $e_1^* < e_2^*$ as $g_1 p'(e_1) > g_2 p'(e_2)$ for all $e_1 \leq e_2$. I.e., given a constant amount of total effort, the farmer can always do better reallocating some effort to the tree that has the higher return. Thus, no interior solution exists where more than one tree is receiving an amount of effort less than the minimum amount that guarantees survival, $\tilde{e}$.

[7] Recall that the continuous component of the profit function confounds marginal costs and benefits. Thus we cannot introduce probabilistic survival to the benefit portion, without affecting the cost per-tree, which should

in the main text under this alternative assumption. For ease of comparison, Panel A shows the reduced form results using the observed data (i.e. is identical to Panel A in the main text). Panel B shows the reduced form results with simulated data under our baseline deterministic tree survival assumption and the estimated parameters of our mean shifter model (i.e. is identical to Panel C). Panels C - E implement the same regressions, with simulated data from a model that keeps our estimated parameters constant (Panel B), but models tree survival outcomes as stochastic and governed by either a binomial distribution (Panels C and D) or a beta binomial distribution (Panel E).[8]

remain deterministic.

[8]We keep the estimated parameters under the deterministic survival assumption instead of reestimating them under the stochastic survival assumption due mainly to computing time constraints. Thus, the fit of the model may further improve if we let other parameters adjust instead of keeping them constant. However, the little sensitivity of the reduced form responses we see in Table B.1 leads us to believe that we would not gain much in terms of fit by reestimating the model under the stochastic survival assumption.

Table B.1: *Stochastic Tree Survival*

	(1) Take-up	(2) 35-tree threshold	(3) # trees | # trees>0	(4) 1.(zero trees)		(5) Take-up	(6) 35-tree threshold	(7) # trees | # trees>0	(8) 1.(zero trees)
Panel A. Observed Data (Repeats Panel A in Reduced Form Table)									
Take-up subsidy	0.022***	-0.004	-0.229	-0.003	Reward	0.001*	0.001***	0.044***	-0.001***
	(0.005)	(0.004)	(0.200)	(0.005)		(0.000)	(0.000)	(0.013)	(0.000)
Observations	1,314	1,092	701	1,092		624	1,092	701	1,092
R-squared	0.071	0.002	0.005	0.001		0.006	0.018	0.022	0.019
Panel B. Mean Shift and No Stochastic Survival (Repeats Panel C in Reduced Form Table)									
Take-up subsidy	0.020***	-0.003	-0.002	0.008**	Reward	0.001*	0.003***	0.094***	-0.001***
	(0.002)	(0.003)	(0.124)	(0.003)		(0.000)	(0.000)	(0.012)	(0.000)
Observations	1,314	1,120	605	1,120		624	1,120	605	1,120
R-squared	0.062	0.001	0.000	0.006		0.006	0.107	0.089	0.013
Panel C. Survival probability = 0.98									
Take-up subsidy	0.020***	-0.002	0.062	0.009**	Reward	0.001*	0.003***	0.105***	-0.001***
	(0.002)	(0.003)	(0.133)	(0.003)		(0.000)	(0.000)	(0.012)	(0.000)
Observations	1,314	1,120	603	1,120		624	1,120	603	1,120
R-squared	0.062	0.000	0.000	0.006		0.006	0.108	0.109	0.012
Panel D. Survival probability = 0.95									
Take-up subsidy	0.020***	-0.002	0.083	0.008**	Reward	0.001*	0.003***	0.101***	-0.001***
	(0.002)	(0.003)	(0.137)	(0.003)		(0.000)	(0.000)	(0.013)	(0.000)
Observations	1,314	1,120	596	1,120		624	1,120	596	1,120
R-squared	0.062	0.000	0.001	0.006		0.006	0.095	0.098	0.012
Panel E. Survival probability distributed beta binomal with mean 0.57 and sd 0.37									
Take-up subsidy	0.022***	-0.001	-0.038	0.006*	Reward	-0.000	0.001***	0.057***	-0.000
	(0.002)	(0.002)	(0.151)	(0.003)		(0.000)	(0.000)	(0.013)	(0.000)
Observations	1,314	1,099	518	1,099		624	1,099	518	1,099
R-squared	0.071	0.000	0.000	0.003		0.000	0.019	0.034	0.002

Notes: This table shows coefficients from regressions of each of four indicator variables (take-up, binary 35-tree threshold, tree survival larger than zero, and no tree survival) on each of our randomized treatments (take-up subsidy and threshold reward) for both non-stochastic and stochastic models. Panel A shows these regression outcomes for the true data. Panel B shows the fit of the strcutural model by simulating all four outocomes using the model estimates and examining the how much the linear relationships between outcomes and treatments resemble those in Panel A. These panels recreate Panels A and C of the reduced form table in the main body of the paper. Panel C here estimates binomial survival assuming that the probabilty any one tree survives is 0.98. Panel D estimates binomial survival assuming that the probabilty any one tree survives is 0.95. Panel E assumes that the probability any one tree survives is distributed beta binomal with mean 0.57 and standard deviation 0.37, corresponding to an alpha parameter of 0.428 and a beta parameter of 0.318.

Panel D assumes that the maximum probability of survival, $p(\tilde{\tilde{e}})$, is 0.98, while Panel D assumes that this maximum survival probability is 0.95. We chose relatively high probabilities for the simulation as lower probabilities result eliminate bunching at 35, and thus are inconsistent with what we observe in our data (see Figure B.1). The beta binomial distribution in Panel E allows for the maximum probability to vary across farmers according to a beta distribution with parameters 0.57 and 0.37. The purpose of this exercise is to examine whether by relaxing the deterministic survival assumption we can do a better job matching the reduced form results in Panel A than do our main estimates, Panel B.

Figure B.1: *Observed Tree Survival Outcomes*

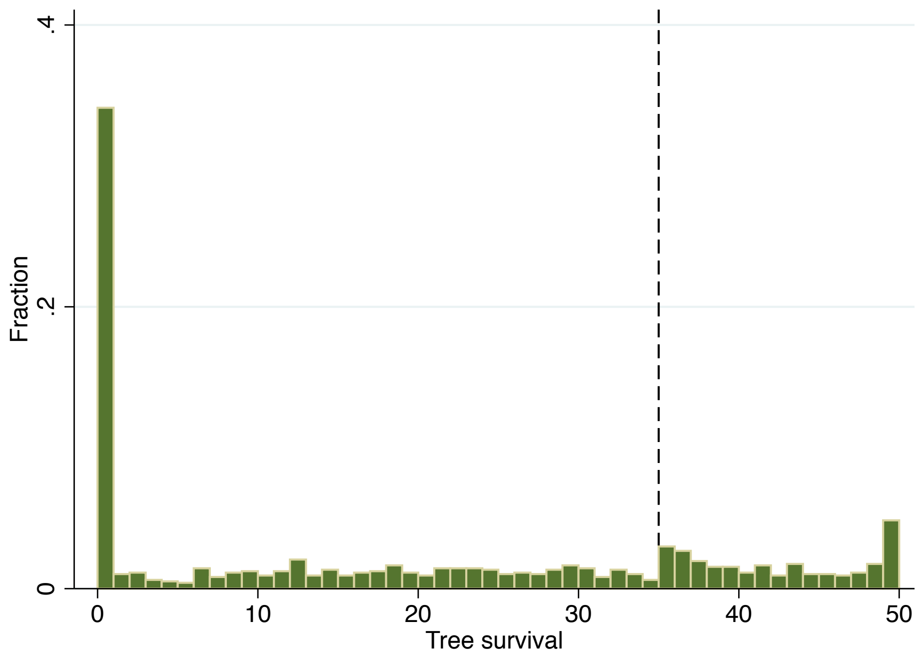

Notes: Histogram of tree survival outcomes for all farmers assigned a positive reward for reaching the 35-tree survival threshold. The threshold is shown by the dashed vertical line.

Overall, we see little improvement when stochasticity is introduced into the tree survival outcomes. The main model performs least well on the relationship between the take-up subsidy and the positive number of trees and zero trees (Panel B, columns 3 and 4). Both models overestimate the effect of the reward on the likelihood of reaching the 35-tree threshold and the number of trees for farmers with any surviving trees (Panel B, columns 6 and 7). The model variants in Panel C and D show no improvement on any of these dimensions, and in some cases worsen the fit. Only Panel E improves on the fit compared to our main model (Panel B), and not by much: the coefficients on the reward for the 35-tree threshold attainment and for positive tree survival are closer to the observed data but qualitative differences remain. Importantly, these improvement come at the expense of a poorer match in other responses that are well-fit by our main model, such as the relationship between the reward and take-up and the relationship between the reward and zero-trees cultivated (columns 5 and 8).

Table B.2: *Knowledge and Experience with the Technology*

	(1)	(2)	(3)	(4)
Knowledge of risks to tree survival (1-5)	4.5035***			4.4380***
	[0.7828]			[0.7762]
Change in risk knowledge	2.4856***			2.4897***
	[0.4819]			[0.4779]
Prior planting of Faidherbia		5.1819**	4.4554**	4.0145**
		[2.1442]	[1.9009]	[1.9328]
Any prior planting in group			1.5780	1.8883
			[1.7852]	[1.7111]

Notes: OLS regressions of tree survival on proxy measures of knowledge and learning. The sample is restricted to farmers in both the baseline and endline surveys. Knowledge of risks to tree survival counts the number of risks that the farmer was able to recall. Change in risk knowlege measures how that number changed between the endline and the baseline. Prior planting of Faidherbia is an indicator for whether the farmer had adopted the technology prior to the program. Any prior planting in group is an indicator for whether anyone in the farmer group had adopted. All columns control for the remaining variables shown in the balance tables and also for treatment variables. * p<0.10 ** p<0.05 *** p<0.01.

Table B.3: *Procrastination*

Dependent variable:	Take-up	Survival	Survival	Take-up	Take-up
	(1)	(2)	(3)	(4)	(5)
	Panel A: Self-described procrastinator				
Binary Procrastination Measure	-0.0080	-0.8574	-1.3361	-0.0273	-0.0016
	[0.0208]	[1.1563]	[1.1410]	[0.0440]	[0.0651]
Take-up subsidy	0.0221***	-0.0513	-0.0184	0.0206***	0.0233***
	[0.0044]	[0.2004]	[0.1952]	[0.0047]	[0.0058]
Reward in '000 ZMK			0.0669***		0.0007**
			[0.0114]		[0.0003]
Procrastination x subsidy				0.0031	-0.0028
				[0.0047]	[0.0069]
Constant	0.4507***	9.0534***	4.3205	0.4622***	0.3980***
	[0.0783]	[3.1850]	[3.1716]	[0.0804]	[0.1088]
N	1275	1071	1071	1275	603
	Panel B: Reports procrastination on other activities				
Binary Procrastination Measure	-0.0302	0.0622	0.1404	-0.0938	-0.0521
	[0.0278]	[1.2728]	[1.2767]	[0.0629]	[0.0767]
Take-up subsidy	0.0231***	-0.1045	-0.0745	0.0197***	0.0205***
	[0.0044]	[0.2021]	[0.1958]	[0.0043]	[0.0050]
Reward in '000 ZMK			0.0686***		0.0006*
			[0.0119]		[0.0003]
Procrastination x subsidy				0.0102	0.0076
				[0.0066]	[0.0080]
Constant	0.4589***	8.8880***	3.8534	0.4766***	0.4333***
	[0.0801]	[3.1334]	[3.1569]	[0.0806]	[0.1014]
N	1223	1030	1030	1223	576

Notes: OLS regressions of take up and survival on indicators of procrastination. Standard errors clustered at the group level are in brackets. Columns 2 and 3 condition on take-up. Column 5 conditions on knowing the reward before take-up (excludes the surprise reward treatment). See text for a description of the procrastination measures used in the regressions. * p<0.10 ** p<0.05 *** p<0.01.

Table B.4: *Incentive Spillovers within Group*

Dependent variable is tree survival	(1)	(2)
Average reward in group (excl. own)	0.262	0.578*
	(0.240)	(0.317)
Own reward	0.319**	0.642**
	(0.0569)	(0.253)
Group reward x own reward		-0.0230
		(0.0182)
N	1088	1088

Notes: OLS regressions of tree survival on average draw in farmer group, conditional on take-up, and own draw. Standard errors are clustered at the group level. * p<0.10 ** p<0.05 *** p<0.01.

Table B.5: *Balance*

	A=0 Mean [SD] (1)	A > 0 (2)	R=0 Mean [SD] (3)	Reward > 0 (4)	Surprise=0 Mean [SD] (5)	Surprise reward (6)	N (7)
Respondent is head of household	0.735	0.001	0.694	0.0000	0.702	0.038	1292
	[0.442]	[0.003]	[0.463]	[0.0003]	[0.458]	[0.025]	
Age, respondent	37.872	-0.074	37.439	0.0071	37.25	1.284*	1266
	[13.716]	[0.100]	[12.808]	[0.0073]	[13.684]	[0.709]	
Female headed household	0.135	0.001	0.149	0.0001	0.119	0.009	1292
	[0.343]	[0.003]	[0.358]	[0.0002]	[0.324]	[0.017]	
Years of education, respondent	5.342	-0.039	5.284	0.0033	5.363	-0.009	1292
	[3.276]	[0.028]	[3.133]	[0.0022]	[3.331]	[0.177]	
Household size	5.465	-0.036**	5.328	0.0003	5.246	0.208*	1292
	[2.142]	[0.015]	[2.390]	[0.0013]	[2.217]	[0.114]	
Ordinal discount rate (1 - 5)	2.454	-0.001	2.538	0.0004	2.43	0.106	1262
	[1.627]	[0.013]	[1.714]	[0.0010]	[1.612]	[0.095]	
Non-agricultural assets	9.806	-0.101**	9.343	-0.0013	9.08	0.314	1292
	[5.789]	[0.041]	[5.625]	[0.0034]	[5.506]	[0.287]	
Years working with Dunavant	4.228	-0.033	3.776	-0.0015	3.865	0.069	1292
	[3.748]	[0.035]	[3.404]	[0.0022]	[3.496]	[0.214]	
Total landholdings (hectares)	3.02	-0.022	2.881	0.0000	2.873	0.056	1290
	[2.248]	[0.023]	[2.188]	[0.0014]	[2.253]	[0.115]	
Number of fields	2.874	0.001	2.866	0.000	2.886	-0.054	1292
	[1.065]	[0.009]	[1.194]	[0.0008]	[1.122]	[0.070]	
Average distance from home to plots	20.532	-0.084	19.416	-0.0202	18.397	0.91	1292
	[24.195]	[0.236]	[22.436]	[0.0122]	[20.625]	[1.068]	
Poor soil fertility	0.108	-0.001	0.104	0.0000	0.106	-0.029*	1292
	[0.310]	[0.002]	[0.307]	[0.0002]	[0.308]	[0.017]	
Regular interaction with lead farmer	0.415	0.004	0.448	-0.0006**	0.412	0.007	1290
	[0.493]	[0.004]	[0.499]	[0.0003]	[0.493]	[0.029]	
Affiliated with CFU or COMACO	0.037	0.002	0.037	0.000	0.042	0.003	1292
	[0.189]	[0.002]	[0.190]	[0.0001]	[0.202]	[0.012]	
Prior knowledge of Faidherbia	0.68	-0.002	0.664	0.000	0.64	0.027	1292
	[0.467]	[0.004]	[0.474]	[0.0003]	[0.480]	[0.025]	
Prior planting of Faidherbia	0.111	-0.001	0.09	0.0001	0.088	0.014	1292
	[0.314]	[0.002]	[0.287]	[0.0002]	[0.283]	[0.014]	
Knowledge of risks to tree survival	1.72	-0.005	1.701	0.0000	1.648	-0.013	1292
	[0.905]	[0.006]	[0.785]	[0.0005]	[0.816]	[0.046]	
N	325	967	134	1041	614	678	

Notes: Means are reported for the base group in columns 1, 3 and 5. Coefficients and standard deviations from a regression of the household variable on treatment are reported in other columns. * p < 0.10 ** p < 0.05 *** p < 0.01.

Table B.6: *Attrition Across Data Collection Phases*

	Takeup Mean [SD]	Baseline	Endline	Tree monitoring
	(1)	(2)	(3)	(4)
Take-up subsidy	6.1564	0.0029*	0.0000	0.0000
	[4.5399]	[0.0016]	[0.0020]	[0.0007]
Reward ('000 ZMK)	69.3347	0.0001	0.0000	0.0000
	[48.4713]	[0.0001]	[0.0001]	[0.0000]
Surprise reward treatment	0.5239	0.0097	0.0124	0.0025
	[0.4996]	[0.0088]	[0.0124]	[0.0061]
N, outcome = 1	1317	1292	1232	1083

Notes: Attrition across data collection rounds by treatment. Column 1 shows means and standard deviations for each treatment. Each cell in columns 2 - 4 shows the coefficient from a regression of an indicator being present at the data collection stage regressed on each treatment with standard errors clustered at the farmer group level. Column 4 is conditional ontake-up (N=1092). For observations missing the reward variable (surprise reward treatment, no take up), a missing variable dummy for the reward is added to the regression. Reported coefficients are among non-missing reward values. * p < 0.10 ** p < 0.05 *** p < 0.01.

Table B.7: *Correlation Between Farmer Observables and Program Outcomes*

	(1)	(2)	(3)	(4)	(5)	(6)
	Take-up		35-tree threshold		Tree survival	
Household head at training	0.0705**	0.0651**	0.0041	0.0010	0.4308	0.2770
	[0.0266]	[0.0229]	[0.0326]	[0.0328]	[1.2126]	[1.1939]
Female household head	0.0298	0.0275	-0.0257	-0.0248	0.2531	0.2885
	[0.0325]	[0.0292]	[0.0379]	[0.0352]	[1.4998]	[1.3798]
Respondent education	0.0009	0.0002	0.0084	0.0079	0.3926*	0.3668*
	[0.0035]	[0.0033]	[0.0044]	[0.0042]	[0.1831]	[0.1744]
Household size	0.0089	0.0104*	0.0082	0.0063	0.2118	0.1240
	[0.0051]	[0.0049]	[0.0058]	[0.0057]	[0.2066]	[0.1994]
Non-agricultural assets	0.0001	0.0017	0.0001	-0.0003	-0.0468	-0.0606
	[0.0020]	[0.0019]	[0.0030]	[0.0029]	[0.1080]	[0.1018]
Years working with Dunavant	0.0050	0.0062	0.0071	0.0077	0.1467	0.1845
	[0.0041]	[0.0034]	[0.0040]	[0.0041]	[0.1623]	[0.1598]
Land size (hectares)	0.0052	0.0052	-0.0037	-0.0043	-0.0041	-0.0261
	[0.0043]	[0.0040]	[0.0064]	[0.0062]	[0.2521]	[0.2424]
Number of fields	0.0108	0.0051	-0.0017	-0.0042	0.7444	0.6132
	[0.0100]	[0.0089]	[0.0127]	[0.0126]	[0.5520]	[0.5467]
Distance from home to plots	-0.0006	-0.0002	0.0003	0.0003	-0.0217	-0.0196
	[0.0007]	[0.0006]	[0.0007]	[0.0007]	[0.0266]	[0.0252]
Poor soil fertility	-0.0310	-0.0200	-0.0171	-0.0237	-1.7101	-1.9494
	[0.0354]	[0.0360]	[0.0517]	[0.0491]	[1.9597]	[1.9100]
Sees YGL often	0.0251	0.0250	-0.0243	-0.0142	0.5124	0.9962
	[0.0194]	[0.0181]	[0.0280]	[0.0270]	[1.1066]	[1.0389]
Affiliated with CFU or COMACO	0.0422	0.0052	0.0793	0.0635	4.9661*	4.0746
	[0.0482]	[0.0430]	[0.0721]	[0.0709]	[2.3327]	[2.3833]
Prior knowledge of Faidherbia	0.0377	0.0352	0.0423	0.0442	0.7249	0.8228
	[0.0288]	[0.0252]	[0.0344]	[0.0321]	[1.3283]	[1.1978]
Prior planting of Faidherbia	-0.0758	-0.0644	0.0643	0.0779	4.3897*	5.0115*
	[0.0461]	[0.0364]	[0.0561]	[0.0573]	[2.1425]	[2.1682]
Knowledge of risks to tree survival	0.0159	0.0157	0.0390**	0.0402**	1.6840**	1.7570**
	[0.0137]	[0.0114]	[0.0127]	[0.0128]	[0.5754]	[0.5794]
Constant	0.6230***	0.3444***	0.0542	-0.0318	8.1421**	2.6194
	[0.0648]	[0.0842]	[0.0620]	[0.0762]	[2.6058]	[3.0433]
R squared	0.0296	0.1898	0.0247	0.0750	0.0314	0.1083
Treatment controls	no	yes	no	yes	no	yes
Obs	1288		1080		1080	
Dep. Var. Mean	0.8385		0.2528		17.6000	

Notes: OLS regressions of outcomes on observables collected as part of the baseline survey, during training. The outcome in columns 3 and 4 is an indicator for reaching the reward threshold (≥ 35 trees). Even columns include controls for the experimental treatments: subsidy level, reward level, reward timing and monitoring. * p<0.10 ** p<0.05 *** p<0.01.

Table B.8: *Effect of Reward Timing on Tree Survival Outcomes*

	(1) Surprise = 0 Mean/[SD]	(2) Surprise = 1 Mean/[SD]	(3) Reward x Surprise Coef/(SE)
R = 0	11.02	11.32	
	[14.33]	[16.00]	
R = (0,70000)]	14.71	15.87	0.85
	[16.86]	[16.87]	(2.72)
R = (70000,150000]	20.32	21.05	0.43
	[17.99]	[17.48]	(2.66)

Notes: Outcome is tree survival (continuous), conditional on take up. Columns 1 and 2 show means and standard deviations in each reward category, by the reward timing condition. Surprise = 1 indicates that farmers learned about the reward only after the take-up decision. Column 3 reports estimated coefficients and standard errors clustered at the farmer group level for a linear regression of tree survival on reward category interatcted with the surprise reward treatment. We report the coefficient on the interaction term only.