

Conservation Agriculture and Climate Resilience

Jeffrey D. Michler^{*a}, Kathy Baylis^b, Mary Arends-Kuenning^b, and Kizito Mazvimavi^c

^a*University of Arizona, Tucson, USA*

^b*University of Illinois, Urbana, USA*

^c*International Crops Research Institute for the Semi-Arid Tropics, Bulawayo, Zimbabwe*

November 2018

Abstract

Agricultural productivity growth is vital for economic and food security outcomes which are threatened by climate change. In response, governments and development agencies are encouraging the adoption of ‘climate-smart’ agricultural technologies, such as conservation agriculture (CA). However, there is little rigorous evidence that demonstrates the effect of CA on production or climate resilience, and what evidence exists is hampered by selection bias. Using panel data from Zimbabwe, we test how CA performs during extreme rainfall events - both shortfalls and surpluses. We control for the endogenous adoption decision and find that use of CA in years of average rainfall results in no yield gains, and in some cases yield losses. However, CA is effective in mitigating the negative impacts of deviations in rainfall. We conclude that the lower yields during normal rainfall seasons may be a proximate factor in low uptake of CA. Policy should focus promotion of CA on these climate resilience benefits.

JEL Classification: O12, O13, Q12, Q54, Q56

Keywords: Conservation Farming; Technology Adoption; Agricultural Production; Weather Risk; Zimbabwe

*Corresponding author email: jeffrey.michler@usak.ca. The authors owe a particular debt to Albert Chirima, Tarisayi Pedzisa, Alex Winter-Nelson, and Yujun Zhou. This work has benefited from helpful comments and criticism by Anna Josephson, David Rohrbach, David Spielman, and Christian Thierfelder as well as seminar participants at the CSAE conference in Oxford, the AERE conference in Pittsburgh, the AAAE conference in Addis Ababa, and the joint SPIA/PIM Conference on technology adoption and impact, in Nairobi. This research was supported by ISPC-SPIA under the grant “Strengthening Impact Assessment in the CGIAR System (SIAC). We gratefully acknowledge financial support for the data collection from the United Kingdom Department for International Development (DFID) and United Nations Food and Agriculture Organization (FAO) through the Protracted Relief Program (PRP) in Zimbabwe, 2007-2011. Neither the funders nor our implementation partners had any role in the analysis, writing of the paper, or decision to submit for publication.

1 Introduction

Smallholder agriculture is not just a source of food but a driver of economic development, particularly for the 75 percent of the world's poor who live in rural areas. However, agricultural production is straining natural resources, suggesting that productivity improvements are required to feed a growing population (FAO et al., 2017). Agriculture and food security are further threatened by climate change, particularly in Sub-Saharan Africa, and particularly for smallholder farmers (Morton, 2007; Schlenker and Lobell, 2010; Wheeler and von Braun, 2013). Climate change has already substantially reduced production in many parts of the world (Lobell et al., 2011). In response, governments and development agencies are encouraging the adoption of 'climate smart' agricultural technologies, including conservation agriculture, with the goal of bolstering productivity, enhancing resilience to weather shocks, and reducing negative externalities (FAO, 2013; Lipper et al., 2014).

Conservation agriculture (CA) is based on three practices promoted as a means for sustainable agricultural intensification: minimum tillage, mulching with crop residue, and crop rotation. The goal of these practices is to increase yields through improvements in soil fertility, and reduce risk to yields from rainfall shocks (Brouder and Gomez-Macpherson, 2014).¹ In addition to the stated private benefits of increased productivity and enhanced yield resilience, proponents claim that CA has positive environmental externalities by increasing soil organic carbon (Smith et al., 1998; Lal and Stewart, 2010).² Yet adoption rates of CA have been low, and disadoption rates are high, limiting the scope of any positive externalities and raising the question of whether farmers receive the promised private benefits (Pedzisa et al., 2015a). The little evidence that exists regarding CA's 'climate smart' properties mainly relies on agronomic analysis of field station trials, missing potential effects of farmer behavior. As Andersson and D'Souza (2014) and Pannell et al. (2014) observe, the few studies that do use observational data often fail to account for selection bias or fully control for all potential sources of endogeneity.

¹This last benefit comes, in the case of sub-optimal rainfall, through the use of planting basins to achieve minimum tillage, in combination with mulching, which increases water infiltration and moisture conservation (Thierfelder and Wall, 2009). In the case of excess rainfall, the retention of crop residue reduces water-driven soil erosion (Schuller et al., 2007).

²Recent agronomic research has shed doubt on the validity of the claim that CA sequesters significant amounts of carbon in soil (Baker et al., 2007; Piccoli et al., 2016; Powlson et al., 2016).

In this paper, we compare CA’s effect on overall yield to its potential in reducing yield losses due to deviations in average rainfall, controlling for potential selection bias. While CA is hypothesized to be ‘climate smart,’ increased resilience during periods of rainfall stress may come at the cost of yields during regular growing conditions. We follow Di Falco and Chavas (2008) by defining resilience in terms of practices that limit productivity loss when challenged by climatic events. We use four years of plot-level panel data covering 4, 171 plots from 729 households across Zimbabwe to estimate how yields respond during periods of high and low rainfall shocks, compared to the norms households experience over time. The data include plot-level inputs and output for five different crops: maize, sorghum, millet, groundnut, and cowpea. Unlike earlier work that estimates CA’s effect on a single crop, this rich set of production data allows us to examine how the impacts of CA vary in the multi-cropping system common across Sub-Saharan Africa.

Causal identification of CA’s impact on yields is complicated by non-random adoption of the technology.³ We identify two potential sources of endogeneity that might bias our results. First is the presence of unobserved household heterogeneity that influences both adoption and yields, which we control for with household fixed effects. Second is the possible presence of unobserved time-varying shocks that might affect a household’s access to and use of CA while being correlated with yields. To instrument for household adoption, we use data on CA promotion through the distribution of subsidized inputs as part of the Zimbabwe Protracted Relief Programme (PRP). The PRP was a four year, £28 million project aimed at providing short-term nutritional, economic, and agricultural interventions to one-third of all smallholder households in the country, about 1.7 million people (Jennings et al., 2013). In addition to these two primary concerns, we conduct a number of robustness checks, including controlling for plot-level unobservables, changes to the functional form, testing for the impact of rainfall event outliers, and changing our definition of rainfall shock.

We find that adoption of CA in years of average rainfall results in no yield gains, and in some cases yield losses, compared to conventional practices. Where CA is effective is in mitigating the

³Our use of the term “adoption” refers to the *use* (as opposed to the *first time* use) of CA on a plot. Similarly, “disadoption” refers to the *use* of conventional cultivation techniques on a plot that was previously cultivated using CA.

negative impacts of deviations in rainfall. The magnitude of these resilience effects vary by crop, and whether rainfall is in shortfall or surplus. We then use the estimated coefficients on the CA terms to predict the returns to CA at various values of rainfall. A one standard deviation decrease or increase in rainfall is required before the returns to CA become positive. Our results, using observational data, support a recent meta-analysis of agronomic field experiments which found that, on average, CA reduces yields but can enhance yields in dry climates (Pittelkow et al., 2015).

Pannell et al. (2014) discuss several limitations with the state of economic research on CA. A key limitation is a lack of clarity regarding what qualifies as CA adoption. Our survey data allow us to partially address this concern. While the data cover four cropping seasons, only in the final two seasons were households asked about which specific practices they implemented on plots that they identified as cultivating with CA. Among households that used CA in the last two years of the survey, 97 percent of them used planting basins to achieve minimum soil disturbance, 27 percent mulched with crop residue, and 36 percent practiced crop rotation. In total, 22 percent of self-identified CA adopters engaged in at least two practices while only five percent of CA adopters actually engaged in all three practices. Therefore, in our context, we use a practical definition of CA as, at the very least, minimum tillage, the original and central principle of CA. We then define traditional cultivation practices as everything other than self-identified CA adoption, which for the vast majority of plots amounts to conventional tillage practices. While our pragmatic approach is not ideal, it is in line with previous literature at the farm level and with meta-analyses of agronomic field experiment data.

This paper makes three contributions to the literature. First, by interacting the instrumented measure of CA adoption with deviations in rainfall we directly test whether CA can mitigate yield loss due to adverse weather events. The only evidence of the resilience of yields under CA comes from station and on-farm field trial data or from observational studies that fails to adequately control for both time-variant and time-invariant sources of endogeneity. In this literature, researchers acknowledge that selection bias is a problem because farmers choose the technology based on its expected benefits, which are heterogeneous. Much of the literature relies on cross-sectional data and techniques that involve estimating selection correction terms (Kassie et al., 2009; Teklewold et al.,

2013; Teklewold et al., 2013; Kassie et al., 2015; Kassie et al., 2015; Abdulai, 2016; Manda et al., 2016). Like any other IV, selection models rely on exclusion restrictions for proper identification of time-variant endogenous choice variables. However, reliance on cross-sectional data means that one cannot easily control for endogeneity coming from time-invariant sources. A second, related, set of literature relies on panel data to control for time-invariant heterogeneity (Arslan et al., 2014; Ndlovu et al., 2014; Arslan et al., 2015; Pedzisa et al., 2015a; Pedzisa et al., 2015b). This literature, though, fails to control for time-variant factors causing endogeneity in the adoption decision. We use realizations of on-farm yields over four years and control for both sources of endogeneity. Ours is one of the first papers to test the ‘climate smart’ properties of CA as they are experienced by farm households.

Second, we expand the analysis of CA by encompassing a variety of crops. Most previous work focuses on the impact of CA on yields for a single crop, often maize (Teklewold et al., 2013; Brouder and Gomez-Macpherson, 2014; Kassie et al., 2015; Manda et al., 2016). But, farmers in Sub-Saharan Africa frequently grow multiple crops in a single season, meaning that decisions regarding cultivation method and input use are made at the farm level, not the crop level. To examine the farm-wide impacts, we adopt a flexible yield function that allows coefficients on inputs to vary across crops. This approach accommodates heterogeneity in the input-response curves by allowing slopes and intercepts to vary across crops without forcing us to split the random sample based on non-random criteria. Our results provide new insight on the use and impact of CA in a multi-cropping environment.

Third, we provide suggestive evidence on the reason for the low adoption rate of CA among farmers in Sub-Saharan Africa. Given early evidence on the positive correlation between CA and yields (Mazvimavi and Twomlow, 2009), and its promotion by research centers, donor agencies, and governments (Andersson and D’Souza, 2014), the slow uptake of CA has presented the adoption literature with an empirical puzzle. Once we control for endogeneity in the adoption decision, CA frequently has a negative impact on yields during periods of average rainfall. We conclude that the lack of yield gains during average rainfall seasons may be a limiting factor in the uptake of CA, and by extension limits any positive environmental externalities CA may provide. Focusing on CA’s

potential as a ‘climate smart’ agricultural technology by advertising the resilience benefits of CA marks a path forward for the promotion of CA in regions facing climate extremes.

2 Theoretical Framework

We begin by defining two stochastic Cobb-Douglas yield functions along the lines of Just and Pope (1978), Barrett et al. (2004), and Suri (2011):

$$Y_{kit}^{\text{CA}} = e^{\beta_{kt}^{\text{CA}}} \left(\prod_{j=1}^J X_{jkit}^{\gamma_{jk}^{\text{CA}}} \right) e^{u_{kit}^{\text{CA}}}, \quad (1)$$

$$Y_{kit}^{\text{TC}} = e^{\beta_{kt}^{\text{TC}}} \left(\prod_{j=1}^J X_{jkit}^{\gamma_{jk}^{\text{TC}}} \right) e^{u_{kit}^{\text{TC}}}, \quad (2)$$

where Y_{kit} is yield for crop k cultivated by household i at time t and X_{jkit} is a set of j measured inputs, including a crop-specific intercept, used on the k^{th} crop by household i at time t . We allow the yield functions for conservation agriculture (CA) and traditional cultivation (TC) to have different parameters for inputs (γ_j^{CA} and γ_j^{TC}), although the same set of potential inputs are used in both cultivation methods. The disturbance terms (u_{kit}^{CA} and u_{kit}^{TC}) are sector-specific errors composed of time-invariant farm and household characteristics and time-variant production shocks.

Taking logs of equations (1) and (2) gives us:

$$y_{kit}^{\text{CA}} = \beta_{kt}^{\text{CA}} + x'_{jkit} \gamma_{jk}^{\text{CA}} + u_{kit}^{\text{CA}}, \quad (3)$$

$$y_{kit}^{\text{TC}} = \beta_{kt}^{\text{TC}} + x'_{jkit} \gamma_{jk}^{\text{TC}} + u_{kit}^{\text{TC}}, \quad (4)$$

where we can decompose the disturbance terms into:

$$u_{kit}^{CA} = \theta_{ki}^{CA} + \xi_{kit}^{CA}, \quad (5)$$

$$u_{kit}^{TC} = \theta_{ki}^{TC} + \xi_{kit}^{TC}. \quad (6)$$

The θ_{ki} 's are crop-level productivity terms based on factors not chosen but known by the household. The transitory errors (ξ_{kit}^{CA} and ξ_{kit}^{TC}) are assumed to be uncorrelated with each other, uncorrelated with the X_{jkit} 's, unknown by households, and have a zero mean and finite variance.

We can define the relative productivity of a plot cultivated using CA compared to traditional methods as $(\theta_{ki}^{CA} - \theta_{ki}^{TC})$, which relies on the decomposition of θ_{ki}^{CA} and θ_{ki}^{TC} into:

$$\theta_{ki}^{CA} = b_k^{CA} \left[\sum_{t=1}^T (R_{kit}^{CA} - R_{kit}^{TC}) \right] + \tau_i, \quad (7)$$

$$\theta_{ki}^{TC} = b_k^{TC} \left[\sum_{t=1}^T (R_{kit}^{CA} - R_{kit}^{TC}) \right] + \tau_i. \quad (8)$$

The b_k 's are a measure of each crop's relative advantage under a given cultivation method and for a certain level of rainfall, R . The τ_i term is the household's absolute advantage in cultivating crops, which does not vary by crop type or cultivation method and is constant over time.

We can re-define this comparative advantage gain as:

$$\theta_{ki} \equiv b_k^{TC} \left[\sum_{t=1}^T (R_{kit}^{CA} - R_{kit}^{TC}) \right] \quad (9)$$

and define a new parameter, $\phi_k \equiv b_k^{CA}/b_k^{TC} - 1$. This allows us to re-write equations (7) and (8) as:

$$\theta_{ki}^{CA} = (\phi + 1)\theta_{ki} + \tau_i, \quad (10)$$

$$\theta_{ki}^{TC} = \theta_{ki} + \tau_i. \quad (11)$$

Substituting the decomposed error terms back into equations (3) and (4) results in:

$$y_{kit}^{\text{CA}} = \beta_{kt}^{\text{CA}} + x'_{jkit} \gamma_{jk}^{\text{CA}} + (\phi + 1)\theta_{ki} + \tau_i + \xi_{kit}^{\text{CA}}, \quad (12)$$

$$y_{kit}^{\text{TC}} = \beta_{kt}^{\text{TC}} + x'_{jkit} \gamma_{jk}^{\text{TC}} + \theta_{ki} + \tau_i + \xi_{kit}^{\text{TC}}. \quad (13)$$

Using a generalized yield function of the form

$$y_{kit} = h_{kit} y_{kit}^{\text{CA}} + (1 - h_{kit}) y_{kit}^{\text{TC}}, \quad (14)$$

we can substitute equations (12) and (13) into equation (14) to get:

$$y_{kit} = \beta_{kt}^{\text{TC}} + \theta_{ki} + (\beta_{kt}^{\text{CA}} - \beta_{kt}^{\text{TC}}) h_{kit} + x'_{jkit} \gamma_{jk}^{\text{TC}} + \phi_k \theta_{ki} h_{kit} + x'_{jkit} (\gamma_{jk}^{\text{CA}} - \gamma_{jk}^{\text{TC}}) h_{kit} + \tau_i + \epsilon_{kit}. \quad (15)$$

Here h_{kit} is an indicator of whether or not crop k was cultivated by household i at time t using CA. The θ_{ki} term is the impact of rainfall, or rainfall shocks, on crop k grown by household i . We are primarily interested in the coefficient ϕ_k on the composite term $\theta_{ki} h_{kit}$, which measures the impact of rainfall on yield given a household was using CA.

The expected gain from using CA is therefore:

$$B_{kit} = y_{kit}^{\text{CA}} - y_{kit}^{\text{TC}} = \phi_k \theta_{ki} + (\beta_{kt}^{\text{CA}} - \beta_{kt}^{\text{TC}}) + x'_{jkit} (\gamma_{jk}^{\text{CA}} - \gamma_{jk}^{\text{TC}}). \quad (16)$$

Neither the household fixed effect term nor the transitory error term appear in equation (16). This is because τ_i is restricted to impact yields identically regardless of cultivation method and because of the assumptions placed on ξ_{kit}^{CA} and ξ_{kit}^{TC} . These logarithmic representations of the yield function map back to a generalized Cobb-Douglas yield function in which returns are heterogeneous based on crop type and cultivation method. While it is recognized that the Cobb-Douglas is not a flexible functional form, in studies of developing country agriculture it is often still preferred to the translog due to data limitations, the simplicity of the production technology, and the frequent

issue of multicollinearity in estimation (Michler and Shively, 2015).

Our primary empirical specification is contained in equation (15). To implement the estimation procedure we make two simplifying assumptions. First, we allow $\beta_{kt}^{CA} - \beta_{kt}^{TC} = \beta_{kt}$ and assume $\beta_{kt} = \beta_k \forall t$. This amounts to assuming that, holding rainfall constant, productivity gains from adoption of CA vary by crop but not over time. While there is evidence from experimental plots that CA improves productivity over time, this typically requires ten to fifteen years (Giller et al., 2009). Given that our panel only covers four years, we do not expect productivity gains to vary over this short of a time frame. Second, we assume $\gamma_{jk}^{CA} = \gamma_{jk}^{TC} = \gamma_{jk}$. This allows the yield response curves to vary by input and by crop but not by cultivation method. We address this second assumption in greater detail in Section 4.1.

Finally, a potential issue with identifying the impact of CA adoption on yields is that the decision to adopt, h_{kit} , may be correlated with transitory shocks contained in ϵ_{kit} . In our empirical implementation we instrument for the likely endogeneity of the adoption term using an approach outlined by Wooldridge (2003). We discuss this approach more fully in Section 4.2.

3 Data

This study uses four years of panel data on smallholder farming practices in Zimbabwe collected by the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT). The data cover 783 households in 45 wards from 2007-2011.⁴ The wards come from 20 different districts which were purposefully selected to provide coverage of high rainfall, medium rainfall, semi-arid, and arid regions. Thus the survey can be considered nationally representative of smallholder agriculture in Zimbabwe. For our analysis we use an unbalanced panel consisting of a subset of 728 randomly selected households. The 55 excluded households come from the 2007 round, which we drop completely, because in that year the survey only targeted households who had received NGO support as part of the Zimbabwe PRP.

⁴Wards are the smallest administrative district in rural Zimbabwe.

3.1 Household Data

Our data provide us with detailed cultivation data for five crops on 4,171 unique plots (see Table 1). Maize, the staple grain of Zimbabwe, is the most commonly cultivated crop. Just over half of all observations are maize, and 98 percent of households grow maize on at least one plot in every year. The next most common crop, in terms of observations, is groundnut, which is often grown in rotation with maize. The third most common crop is sorghum, which is frequently grown as an alternative to maize in the semi-arid regions of Zimbabwe. As an alternative to groundnut and sorghum, households will often grow cowpea or pearl millet, respectively. Both of these crops are much less common and combine to account for only 15 percent of observations.

Examining the production data by year reveals a high degree of annual variability (see Table 2).⁵ Yields in 2009 and 2010 were about 30 percent higher than yields in 2008 or in 2011, despite much higher levels of fertilizer and seed use in the low yield years. CA practices vary both by crop and over time (see Figure 1). In 2008, adoption of CA was at over 40 percent for those cultivating maize, sorghum, groundnut, and cowpea and it was about 30 percent for those cultivating millet. Since then, adoption of CA has declined for all crops so that the average adoption rate is now only 17 percent, although use of CA for maize cultivation remains relatively high at 25 percent. Previous literature has hypothesized that this abandonment of CA by households in Zimbabwe is due to the withdrawal of NGO input support as the PRP was scaled down (Ndlovu et al., 2014; Pedzisa et al., 2015b). We make use of this insight to help identify CA adoption.

3.2 Rainfall Data

To calculate rainfall shocks we use satellite imagery from the Climate Hazards Group InfraRed Precipitation with Station (CHIRPS) data. CHIRPS is a thirty year, quasi-global rainfall dataset that spans 50°S-50°N, with all longitudes and incorporates 0.05° resolution satellite imagery with in-situ station data to create a gridded rainfall time series (Funk et al., 2015). The dataset provides daily rainfall measurements from 1981 up to the current year. We overlay ward boundaries on the 0.05° grid cells and take the average rainfall for the day within the ward. We then aggregate the

⁵See Appendix A for a more detailed discussion of year-over-year changes in the household data.

ward level daily rainfall data to the seasonal level.⁶

Figure 2 shows historic seasonal rainfall distribution by ward over the 15 year period 1997-2011. We observe large seasonal fluctuations as well as longer term cycles in rainfall patterns. The four-year period 1997-2000 saw relatively high levels of rainfall. This was followed by a six-year period in which rainfall was relatively scarce. Since 2007, and throughout the survey period, rainfall was again relatively abundant. It is important to note that only one year in the survey period, 2011, experienced below average rainfall. To some extent, our results will be driven by the yield outcomes from the single drought year in our study.⁷

We follow Ward and Shively (2015) in measuring rainfall shocks as normalized deviations in a single season’s rainfall from expected seasonal rainfall over the 15 year period 1997-2011:

$$R_{wt} = \left| \frac{r_{wt} - \bar{r}_w}{\sigma_{r_w}} \right|. \quad (17)$$

Here shocks are calculated for each ward w in year t where r_{wt} is the observed amount of rainfall for the season, \bar{r}_w is the average seasonal rainfall for the ward over the period, and σ_{r_w} is the standard deviation of rainfall during the same period. Our rainfall shock variable treats too little rain as having the same effect as too much rain. While this may not appear to be meaningful in an agronomic sense, proponents of CA claim that CA will have a positive impact on yields in both abnormally wet and abnormally dry conditions (Thierfelder and Wall, 2009).

We also calculate a measure of rainfall shortage as:

$$\underline{R}_{wt} = \begin{cases} \left| \frac{r_{wt} - \bar{r}_w}{\sigma_{r_w}} \right| & \text{if } r_{wt} < \bar{r}_w \\ 0 & \text{otherwise} \end{cases}, \quad (18)$$

as well as a measure of rainfall surplus:

⁶The rainy season in Zimbabwe runs from November to March, with households planting immediately ahead of the first rains and harvesting in May or June.

⁷For more detail regarding seasonal variation in rainfall see Appendix A. For results from robustness checks using various subsets of years see Appendix D.

$$\bar{R}_{wt} = \begin{cases} \left| \frac{r_{wt} - \bar{r}_w}{\sigma_{rw}} \right| & \text{if } r_{wt} > \bar{r}_w \\ 0 & \text{otherwise} \end{cases} . \quad (19)$$

These measures help clarify if CA’s impact on yield resilience is primarily due to mitigating loss from drought or from excess rainfall.

One potential concern with our rainfall terms is that they are measures of annual deviations in meteorological data which are designed to proxy for agricultural drought or overabundance of rainfall. In Zimbabwe, agricultural drought can take the form of late onset of rains or mid-season dry spells. A further complication is that rainfall in Zimbabwe is often of high intensity but low duration and frequency, creating high runoff and little soil permeation. To test for this, we examine the impact of both seasonal and monthly deviations in rainfall on crop output. We find little empirical evidence that late onset of rainfall or mid-season dry spells reduce yields. Results from these regressions, as well as a more detailed discussion of historical rainfall patterns, are presented in Appendix A. Based on these results we conclude that our use of seasonal deviations is a strong proxy for rainfall-induced stresses to agricultural production.

4 Identification Strategy

Our analysis is based on observational data from a setting in which households were randomly selected into the survey but adoption of CA was not random. Thus, care must be taken in understanding what constitutes the relevant comparable group for hypothesis testing. Additionally, identification of the impact of CA on yields is confounded by two potential sources of endogeneity. First, time-invariant unobserved household-level characteristics might influence both CA adoption and yields. Second, time-variant unobserved shocks captured in ϵ_{kit} may simultaneously affect the adoption decision and crop yields.

4.1 Establishing the Relevant Comparison Group

While we want to compare adopters to non-adopters, if CA requires different levels of inputs, or if yields respond in different ways to inputs under CA compared to traditional cultivation (TC), we cannot simply compare across the two groups, irrespective of any selection issues.

In Table 3 we explore the differences in outputs and input use by crop across cultivation practices.⁸ Mean yields for all crops are significantly higher under CA than under TC methods. One obvious potential reason why CA is often associated with higher yields is that households increase the intensity of agricultural input application under CA. Compared to TC methods, CA is associated with significantly higher levels of both basal and top applied fertilizer. Conversely, more seed and land is used in TC compared to CA for most crops. The different rates of input use between CA and TC means that the two cultivation techniques are not directly comparable without taking into consideration measured inputs. Because we observe the quantity of fertilizer, seed applied, and the size of each plot, we can directly control for these differences in our econometric analysis.

Directly controlling for input use in our regression framework only allows us to compare CA adopters to non-CA adopters (ignoring issues of selection), if our assumption that yield response curves do not vary by cultivation method holds. We can verify this by graphing the correlation between yields and inputs. In the top row of Figure 3 we draw scatter plots of yields and seeding rates in panel (1), yields and basal fertilizer application rates in panel (2), and yields and top fertilizer application rates in panel (3). Blue dots represent observations from CA plots and yellow dots represent observations from TC plots. We then fit a quadratic function to the data along with a 90 percent confidence interval band. Panels (1) - (3) imply that CA does not affect the responsiveness of yields to seed but may affect the responsiveness of yields to both basal and top applied fertilizer. Thus, we cannot compare yields of CA plots to TC plots, even when explicitly controlling for the amount of measured inputs applied to the plot.

Our data were collected during the implementation of the Zimbabwe Protracted Relief Program

⁸For each crop-cultivation pair we first test for normality of the data using the Shapiro-Wilk test. In every case we reject the null that the data is normally distributed. Because of this, we rely on the Mann-Whitney (MW) test instead of the standard t -test to determine if differences exist within crops across cultivation practices. Unlike the t -test, the MW test does not require the assumption of a normal distribution. In the context of summary statistics we also prefer the MW test to the Kolmogorov-Smirnov (KS) test, since the MW test is a test of location while the KS test is a test for shape. Results using the KS test are equivalent to those obtained from the MW test.

(PRP). The PRP simultaneously promoted CA and provided access to subsidized fertilizer and seed. While the subsidy program provided nearly universal coverage at the district level, within districts program support varied from ward-to-ward. Thus, a household’s access to subsidized inputs varies depending on which ward a household lives in. To examine if the difference in yield response to inputs across cultivation method is a result of access to inputs, and not a result of CA itself, we estimate four different regressions: log of yield on ward indicators, log of seed on ward indicators, log of basal fertilizer on ward indicators, and log of top fertilizer on ward indicators. We then plot the residuals from these different regressions and again fit a quadratic function to the data (see panels (4) - (6) in Figure 3). The result allows us to compare yield response by cultivation method within a ward, and thus control for different rates of access to inputs across wards. Comparing panels in the top of Figure 3 to panels in the bottom, much of the apparent differences in yield response to inputs is actually a result of ward effects. While differences in yield response do not completely disappear, the differences that remain are small compared to differences coming from variation in access to inputs. We conclude that after controlling for ward, our assumption that yield response curves do not vary by cultivation method appears to hold.⁹

In our regression analysis we control for differences in access to inputs resulting from ward location using household fixed effects. Additionally, household fixed effects remove any other unobserved time-invariant household effects, such as skill at farming, wealth, or risk aversion, that may be correlated with both the τ_i term, the decision to adopt, and yields in equation (15). Thus, we are estimating the yield response to CA within a household during periods of high and low rainfall shocks, compared to the norms a household experiences over time.

Within a household, farmers may choose to apply CA techniques to some plots and not others based on plot-level unobservables. To test whether this is a valid concern, we compare output and input use across two different “treatment” and “control” groups. In the first, we compare plots on which CA was used in all four rounds of the data (always adopters) with plots that

⁹In addition to the possibility that quantity of input use varies by location, there is also the possibility that quality of input use varies. While the data lacks direct measures of quality, it does contain information about the source of inputs. There does not appear to be large differences in the source of fertilizer based on whether or not the plot was cultivated using CA. We do find that newly purchased seed tends to be used on TC plots, but any bias introduced by differences in seed quality would make our estimates of yield resilience due to CA an underestimate.

were cultivated using CA in 2008 but which reverted to TC practices in subsequent years (future disadopters). In the second, we compare plots that never experienced CA (never adopters) with plots that were cultivated using TC in 2008 but which adopted CA in subsequent years (future adopters). The intuition behind comparing these plot-level adoption types is that if differences in output and input use exist only after adoption histories diverge, then plot-level unobservables are not a concern. Examining Table 4, we see that in 2008 output and input use levels tend to be statistically indistinguishable from each other.¹⁰ The two exceptions are the seeding rate between always adopters and future disadopters, and the rate of application of top fertilizer between never adopters and future adopters. By 2011, when adoption histories are different, both yields and input use has diverged. The one exception is basal fertilizer, which is the same between never adopters and future adopters. Thus, there is *prima facie* evidence that differences in yields on CA and TC plots are due to differences in input use and not due to differences in unobserved plot-level characteristics.

In summary, Table 3 demonstrates that a naive comparison that does not consider differences in input use would lead to the erroneous conclusion that CA, by itself, increases yields for all crops. To address this issue, we include measured inputs by crop type in our regression analysis so that we are comparing CA plots to TC plots with similar input use. A second concern is that yields under CA might respond differently from yields under TC. Figure 3 shows that controlling for differences in access to inputs at the ward level explains most of the differences in yield response. Our use of household fixed effects controls for this and means that we are comparing yields on CA plots within a household during periods of high and low rainfall shocks to the norms a household experiences over time. A final concern is that households may choose to apply CA to certain plots based on unobservable plot-level characteristics, thus making plot-level comparisons within a household invalid. Table 4 shows that differences in yields are due to differences in input use and not differences in unobserved plot-level characteristics. As a robustness check, in subsection 5.3 we estimate yield functions using plot-level fixed effects, which allows us to compare yields under CA

¹⁰Similar to our approach in Figure 3, we first regress each variable on ward indicators in order to control for differences in access to inputs across wards. The values displayed in each cell of the table are residuals from these regressions.

and TC methods for plots whose adoption status changed over the period of study.

4.2 Selection Bias

Having established the relevant comparison group within our data, we next deal with the potential of selection bias. Because adoption of CA is not random, it is likely to be correlated with unobserved time-varying factors. In the case of Zimbabwe, CA was promoted as one element in the PRP, which also included subsidies for inputs (Mazvimavi et al., 2008). Thus, adoption rates of CA within a ward are strongly correlated with the level of input subsidies distributed in the ward (Mazvimavi and Twomlow, 2009; Ndlovu et al., 2014; Pedzisa et al., 2015b). Because the intensity of promotion and the level of subsidies changed from year-to-year, adoption of CA is neither random nor static and therefore is likely correlated with unobserved time-varying factors.

To address the endogeneity of CA we instrument for plot-level adoption using the number of households in the ward that receive NGO support as part of the PRP. An appropriate instrument for this model need not be random but must be correlated with the plot-level decision to adopt CA and uncorrelated with plot yields, except through the treatment. As Figure 4 shows, there is strong correlation between the probability that a household adopts CA and their receiving assistance from an NGO, as part of the PRP, in the form of input subsidies. According to the final report on the PRP, rural households were categorized into four standardized groups based on asset ownership, labor availability, and severity of cash constraints (Jennings et al., 2013). While the PRP provided support to 1.7 million people, those who were eligible for input subsidies, and thus most likely to adopt CA, were required to be households with land and labor but no cash. Eligibility for NGO support was well defined but due to budgetary limitations not all eligible households received support, meaning that NGOs may have targeted individual households based on unobservable characteristics.

That not all eligible households received subsidized inputs raises the concern that correlation exists among NGO targeting of households, CA adoption, and yields. To demonstrate that our instrument satisfies the exclusion restriction, we must demonstrate that any correlation between our instrument and yields occurs only through the adoption of CA or can be explicitly controlled for

in our regression framework. One potential pathway of contamination is through NGO targeting of households based on time-invariant characteristics, such as those defined in the eligibility criteria of the PRP. Our use of fixed effects controls for this possibility, making the use of NGO support a plausibly valid instrument. A second potential pathway of contamination is that input subsidies likely have a direct effect on yields, making it an invalid instrument. Here we can directly control for this pathway by including the amount of seed and fertilizer applied on each plot, though we cannot control for quality of inputs. A third potential pathway of contamination is that certain wards or districts received abnormal levels of support because of local levels of land or asset ownership. While average land and asset ownership at the district-level does vary, the percentage of households in a district receiving NGO support is remarkably consistent across districts, never falling below 59 percent and never reaching above 80 percent. A final potential pathway of contamination is that NGO targeting was based on some unobservable time-varying characteristics that are directly correlated with yields. While far from definitive proof, Figure 4 shows that there is little correlation from year-to-year between plot-level yields, CA adoption by household, and the number of households within a ward that received input subsidies as part of the PRP.¹¹ The level of support that neighboring households receive is correlated with a given household receiving NGO support but does not directly impact that household’s plot-level yields, especially once we control for input use, household effects, and local conditions.¹² Thus, the number of households in the ward that receive NGO support satisfies the exclusion restriction.

4.3 Estimation Procedure

Our primary specification follows directly from equation (15):

$$y_{kit} = \beta_k h_{kit} + \theta_k + \phi_k \theta_k h_{kit} + x'_{jkit} \gamma_{jk} + \tau_i + \delta_t + \epsilon_{kit}. \quad (20)$$

¹¹The correlation coefficient between CA and plot yields is $\rho = 0.1585$ while the correlation coefficient between our IV and plot yields is $\rho = 0.0018$.

¹²One may be concerned with potential spillover effects, in that a household that adopts CA due to NGO support might result in increased yields for a neighbor. For the case of CA in Zimbabwe this is unlikely for two reasons. First, the practices that constitute CA are location specific, in that use of planting basins or residue or rotation on one plot will have no impact on the productivity of another plot. Second, even if this was not the case, households in the survey are relatively dispersed throughout wards, meaning that the plots of one household are not contiguous with neighboring households.

Recall that h_{kit} is an indicator for whether crop k on plot i at time t is under CA cultivation and β_k is the associated impact of CA on yields. Note also that θ_k is the impact on crop k of rainfall shock R_{wt} in ward w at time t , as defined in equation (9). ϕ_k is the measure of the yield resilience of crop k when it is cultivated using CA and experiences rainfall shock R_{wt} , x'_{jkit} is a vector of crop, plot, and time specific inputs, including crop-specific intercept terms, and τ_i is a time-invariant household effect. To our original yield function we add year dummies as additional controls.

Given that adoption of CA, h_{it} , is interacted with either crop type (to give h_{kit}) or with crop type and the exogenous rainfall shock (to give $\theta_k h_{kit}$), we follow Wooldridge (2003) in instrumenting for only the potentially endogenous term. Prior to estimating equation (20) with two-stage least squares (2SLS), we estimate a zero-stage probit that predicts \hat{h}_{it} from the following latent variable model of adoption:

$$h_{it} = 1 [\zeta Z_{wt} + \theta_k + x'_{jkit} \alpha_{jk} + \tau_i + \delta_t + v_{it} \geq 0], \quad (21)$$

where Z_{wt} is number of households in ward w that receive NGO support at time t , ζ is the associated coefficient, and $v_{it} \sim \mathcal{N}(0, \sigma_{v_{it}}^2)$ and is independent of u_{it} and τ_i . Since our zero-stage equation is nonlinear, we use a Mundlak-Chamberlain device to control for household unobservables instead of the fixed effects used in the 2SLS. This amounts to replacing τ_i with its linear projection onto the time averages of observable choice variables $\tau_i = \bar{\mathbf{x}}_i \lambda + c_i$, where $c_i \sim \mathcal{N}(0, \sigma_{c_i}^2)$ (Mundlak, 1978; Chamberlain, 1984). Using the Mundlak-Chamberlain device allows us to avoid the incidental variable problem created by using fixed effects in a nonlinear model while still controlling for unobservables (Wooldridge, 2010). While our zero-stage and 2SLS equations take different approaches to controlling for τ_i , identification of h_{it} still fully relies on Z_{wt} . Any difference that exists between the fixed effects and the Mundlak-Chamberlain device is captured in c_i , which in expectation has zero mean. Our use of the Mundlak-Chamberlain device will be less efficient than fixed effects to the extent that $\sigma_{c_i}^2 > 0$, but this introduces no bias into our estimate of h_{it} . Most importantly, the use of the Mundlak-Chamberlain device does not change the causal identification as it strips out time-invariant effects in the same way as using fixed effects and therefore introduces no bias into our subsequent 2SLS estimates. This has become the preferred method for controlling for

time-invariant unobserved heterogeneity in non-linear equations (Ricker-Gilbert et al., 2011; Suri, 2011; Bezu et al., 2014; Verkaart et al., 2017).

We then calculate the Inverse Mills Ratio (IMR), or sample selection correction term, using \hat{h}_{it} and construct instruments for h_{kit} by interacting the IMR with crop type and construct instruments for $\theta_{ki}h_{kit}$ by interacting the IMR with crop type and the exogenous rainfall shock. Wooldridge (2003) shows that this approach produces consistent estimates and improves on efficiency when compared to instrumenting the entire interaction term. We then use these instruments to estimate equation (20) using 2SLS.

This allows us to conduct several diagnostic tests for our instruments. First, we calculate the Kleibergen-Paap LM statistic, which is a test for underidentification (i.e., is the correlation between the endogenous variables and the instruments statistically different from zero). Second, we conduct two different weak instrument tests, which are a higher hurdle than the test for underidentification. The problem with weak instrument tests is that there is currently no universally accepted way in which to construct p -values for these statistics when 1) errors are not i.i.d. and 2) the equation is exactly identified (Bazzi and Clemens, 2013). This is because the Stock and Yogo (2005) critical values frequently used in weak instrument tests are calculated under the assumption that the errors in the relevant regression are i.i.d. and that there are at least two more overidentifying restrictions than the number of endogenous variables. We address this problem by conducting two different weak instruments tests. First, we report the Anderson and Rubin (1949) test statistic. The AR test, in the case of our exactly identified equations, is a test that the coefficients on all endogenous regressors are zero. This test is robust to weak instruments yet it lacks power in the presence of many instruments, which we have (Kleibergen and Paap, 2006). As an alternative to the AR test, we follow Bazzi and Clemens (2013) in constructing p -values for use in conducting weak instrument tests using the Kleibergen-Paap Wald stat, which is robust to heteroskedasticity. Because critical values have not been tabulated for the Kleibergen-Paap Wald stat, we follow the literature and apply the critical values tabulated by Stock and Yogo (2005) for the Cragg-Donald Wald stat to the Kleibergen-Paap results (Baum et al., 2007).¹³

¹³For recent uses of weak instrument tests in applications where the Stock and Yogo (2005) critical values do not directly apply, see Sanderson and Windmeijer (2016) and Bun and Harrison (2018).

To ensure correct hypothesis testing on our coefficients, we allow the variance structure of the error term to vary by household as well as by crop and cluster our standard errors at the household-crop level. This procedure is not without its critics. Bertrand et al. (2004) suggest that clustering at a single level is preferred to clustering at two levels. This provides two alternatives: cluster only at the crop-level or cluster only at the household-level. Given that we only have five different crops, and a large set of parameters to estimate, we are unable to directly cluster standard errors at just the crop level. Clustering at the household-level provides results qualitatively equivalent to those when we cluster at the household-crop level.¹⁴

5 Results

We present the results from a large complement of estimates in Tables 5 - 8. All models are estimated using the log of yield as the dependent variable and log values of measured inputs as independent variables. Hence, point estimates can be read directly as elasticities.¹⁵ For brevity, we only report coefficients on CA, the rainfall deviations, and the interaction terms. Estimated coefficients on measured inputs are presented in Appendix B.

5.1 Main Results

We first focus on the results presented in Table 5 in which CA is treated as exogenous. While it is unlikely that the decision to adopt CA is uncorrelated with the time-varying error term, these results are informative as they allow us to directly compare our estimates with previous literature on the correlation between CA and yields. Column (1) presents a simple yield function that lacks both our rainfall variable and household fixed effects. Results in column (2) come from the same regression but with fixed effects to control for time-invariant household unobservables. The correlation between CA and yields is positive and significant for both maize and groundnut

¹⁴To address the issue of too few crops for clustering, we estimated crop-specific yield functions with household fixed effects as a system of equations. This approach allows for correlation among crop-specific error terms, though it implicitly imposes the assumption that observations are independent if they are in the same household (Greene, 2011). We also estimated our preferred specification but clustered errors at just the household level.

¹⁵Given the prevalence of zero values in the input data, and to a lesser extent in the output data, we use the inverse hyperbolic sine transformation to convert levels into logarithmic values.

with and without household fixed effects. In all other crop cases CA has no statistically significant association with yields.

Adding deviations from average rainfall and its interaction with CA tells a very different story. Columns (3) and (4) present point estimates of the more flexible yield function with and without household fixed effects. Focusing on the fixed effects results in column (4), CA by itself no longer increases yields for any crop and appears to be correlated with lower yields for sorghum. Exposure to a rainfall shock decreases yields for all crops. When we examine the interaction terms, we find that CA is correlated with higher yields for maize and sorghum during periods of rainfall stress. For millet, groundnut, and cowpea CA has no statistically significant association with yields, regardless of rainfall levels.

The results in columns (1) and (2) of a positive correlation between maize yields and CA are similar to the results presented in much of the previous literature (Kassie et al., 2009; Mazvimavi and Twomlow, 2009; Kassie et al., 2010; Teklewold et al., 2013; Brouder and Gomez-Macpherson, 2014; Ndlovu et al., 2014; Abdulai, 2016; Manda et al., 2016). These positive and statistically significant correlations are often interpreted as demonstrating that CA increases yields compared to TC and are used to justify the continued promotion of CA (Giller et al., 2009). However, similar to Arslan et al. (2015), we find that this result is not robust to the inclusion of rainfall measures. Additionally, we expect this result to be biased due to correlation between the decision to adopt and the error term.

Table 6 presents results from the zero-stage probit regressions. For all four specifications the number of households in the ward that receive NGO support is positive and significantly correlated with the choice to adopt CA. In Table 7 we present results similar to those in Table 5 but controlling for the endogeneity of CA. Our instruments pass both the underidentification and weak instrument tests, but only in regressions with household fixed effects. Because of this, our preferred specification is column (4) because it includes CA-rainfall interaction terms and simultaneously controls for unobserved shocks through the IV and unobserved heterogeneity through household fixed effects. We find that CA, by itself, decreases yields on maize and provides no significant advantage over TC for the other crops. Similar to the results in column (4) of Table 5, we find that rainfall deviations

reduce yields on all crops.

While the lack of impact of CA on yields is discouraging, it is not the full story.¹⁶ When we examine the interaction between CA and rainfall we find that CA increases yields in times of rainfall stress for maize and groundnut. For all crops, except sorghum, the coefficients on the CA terms are of a larger magnitude than in the regressions that treat CA as exogenous. Once we control for the endogeneity of the adoption decision, the coefficients on CA tend to be more negative while the coefficients on the interaction terms tend to be more positive. The bias generated from not controlling for the endogeneity of CA appears to underestimate the impact, either positive or negative, of CA on yields. Having controlled for the bias, we conclude that smallholder farmers in Zimbabwe who cultivate their crops using CA practices receive higher yields compared to conventional farmers but only in times of rainfall stress. Appendix C presents the outcomes from a number of robustness checks on our main results.

5.2 Alternative Rainfall Measures

While our main results provide evidence that CA cultivation during deviations from average rainfall helps mitigate crop loss, our measure of rainfall deviation is agnostic to whether or not the shock is from surplus rainfall or a shortage of rainfall. While proponents of CA claim that CA will have a positive impact on yields in both abnormally wet and abnormally dry conditions, there is reason to believe that CA may be more effective in one situation compared to the other, depending on the crop in question.

The first three columns in Table 8 present results from regressions which treat CA as endogenous and include household fixed effects. Column (1) shows results from the regression with the rainfall shock measured as a shortage as in equation (18). Column (2) replaces the rainfall shortage with a rainfall surplus as in equation (19). Column (3) uses both rainfall shortage and surplus measures. Instruments in all three regressions pass the underidentification test and the AR weak instrument test. Only instruments in the regressions from column (1) and (2) pass the Kleibergen-Paap weak

¹⁶It is prudent to remember that in many cases the positive yield impacts of CA only appear after ten to fifteen years of CA cultivation (Giller et al., 2009, 2011). Given that our panel only covers four years, it may be that households have yet to reap the benefits of CA practices.

instrument test. Despite this, the specification in column (3) is our preferred one because it includes the full range of data and allows the impact of CA to vary based on the type of rainfall event and what crop is being cultivated.

Focusing on column (3), rainfall shortages have a negative and significant impact on yields for all crops, while rainfall surpluses have a negative and significant impact on yields for all crops, except groundnut. In average rainfall periods, the use of CA does not have a significant impact on any crop. Examining the interaction terms, the use of CA improves maize yields both when rainfall is above average and during times of drought. For sorghum and cowpea, CA improves yields during times of drought but not surplus rainfall. CA has no specific impact in mitigating losses from either surpluses or shortfalls of rain for millet and groundnut.

In summary, we find that during periods of average rainfall, CA typically has no impact on yields compared to TC. Furthermore, the coefficient on CA is generally negative, suggesting that if CA has any impact on yields it is to reduce them compared to TC. This is in marked contrast to much of the previous literature, which finds a positive correlation between CA and yields. We believe this difference is due to previous studies failing to control for the multiple sources of endogeneity in the CA adoption decision. Second, during seasons that experience above or below average rainfall, CA mitigates yield losses due to these deviations. Maize and groundnut yields are more resilient under CA. Third, when we allow for rainfall shortages to impact yields differently from surplus rainfall we find that only maize yields are consistently more resilient under CA than under TC. Appendix D tests the robustness of these conclusions using a number of different rainfall specifications.¹⁷ In general, CA either has a positive impact on yields in times of stress or it has no impact at all. We conclude that while CA may not improve yields during average seasons, and may even decrease yields, production using CA is more resilient, especially for maize, when rainfall shocks occur.

5.3 Plot Level Controls

One potential concern with our previous results is that the choice to adopt CA may not be driven by unobserved household characteristics but instead by unobserved plot-level characteristics. Given

¹⁷We also conducted tests using satellite temperature data in place of and in combination with rainfall data. Appendix A briefly discusses these results.

that CA is promoted as a technology to halt and reverse land degradation, households may apply CA on plots where they know soil quality is poor. If this is the case, we would expect yields on CA plots to be systematically lower than yields on other plots. Alternatively, though the reasoning is less clear, households might apply CA on plots where they know soil quality is good.

Our panel covers only four seasons with households initially adopting CA in one to four years prior to data collection. Since it takes anywhere between ten and fifteen seasons for CA to significantly increase soil organic matter (Giller et al., 2009), we assume that plot characteristics, such as soil quality, are time-invariant in our data. If this is not the case, our estimates will be biased downward. We employ panel data methods to control for the correlation between the decision to adopt and unobserved plot characteristics. However, there might also be time-variant shocks that influence the decision to put a specific plot under CA instead of influencing the household-level decision to adopt. To control for potential time-invariant shocks we again use instrumental variables.

Columns (4) in Table 8 presents results from a regression designed to control for both sources of plot-level endogeneity. We restrict the sample to plots that we observe more than once over the study period. This reduces our sample size to 5,004 but allows us to directly control for unobserved heterogeneity at the plot-level. However, due to issues of collinearity we are unable to use plot-level fixed effects and instead implement the Mundlak-Chamberlain device in the zero-stage and 2SLS regressions. We find that our results for all crops, when controlling for endogeneity at the household-level (column (3)), are robust in our plot-level endogenous CA regression. CA by itself has no impact on yields and CA continues to build resilience for maize, sorghum, and cowpea. Similar to our previous results, we find that CA has no specific impact in mitigating losses from either surplus or shortfalls of rain for millet and groundnut.

5.4 Returns to CA

To provide some intuition on the size of impact CA has on yields we calculate predicted returns to adoption at various levels of rainfall as in equation (16). Using the results in column (3) of Table 8 (which controls for endogeneity of adoption and household fixed effects), we multiply the coefficient

on the CA-rainfall interaction terms by the realized values of the rainfall shortages or surpluses. We then sum these values along with the coefficient on the CA-only term. We calculate predicted returns for each crop as well as for an average across crops.

Figure 5 graphs the returns to CA across the realized values of rainfall surpluses and shortages. For maize, the returns to CA are positive only when rainfall is one standard deviation above the average. In seasons where cumulative rainfall is below one standard deviation, the returns to CA practices are negative. The returns to using CA to cultivate sorghum are positive for almost any shortage in rainfall while they are close to zero for above average rainfall. For millet, the returns to CA are rarely ever positive, though unlike sorghum, returns are positive when rainfall is well above average. The returns to CA cultivation of groundnut are near zero, regardless of rainfall. Finally, for cowpea, returns to CA are similar to sorghum, though the negative effect of CA is more pronounced.

Taking a weighted average across crops, where the weights are the number of observations for each crop in the data, we can describe the returns to CA for a typical smallholder household in the multi-cropping environment of Zimbabwe. Households would need to experience rainfall shortages greater than one and a half standard deviations away from the mean or rainfall surpluses greater than one standard deviation away from the mean before the returns to CA would become positive. During seasons where rainfall was within this range, the average returns to CA would be negative and households would be better off using TC practices. Examining seasonal rainfall data from 1997 - 2015 across the 45 wards in our sample, 65 percent of seasons have fallen within this ‘normal’ range, making the returns to CA negative. Thirty-four percent of the time rainfall was either low enough or high enough that the returns to CA would have been positive.

To provide a sense of the economic returns to CA adoption in Zimbabwe, we also calculate predicted revenue to both CA and TC practices. First, we predict \hat{y}_{CA} and \hat{y}_{TC} for CA and TC plots from the regression presented in column (3) of Table 8. Second, we create a counterfactual values for each plot by subtracting the estimated contribution of CA ($\hat{\beta}$ and $\hat{\phi}\theta$) from \hat{y}_{CA} and adding the estimated contribution of CA to \hat{y}_{TC} . This allows us to compare yields on a plot with the predict yields had the plot being cultivated using an alternative method. Third, we assign a

monetary value to yields for each crop using published data from the Government of Zimbabwe’s Agricultural Marketing Authority.¹⁸ Finally, we graph the revenue associated with each cultivation practice on each plot, holding input use constant, at different realizations of rainfall surpluses and shortages using kernel-weighted local polynomial smoothing, with confidence intervals drawn at 90 percent.

The results of this “back of the envelope” approach to calculating revenue are encouraging for smallholder agricultural production in Zimbabwe. At market prices, farming generates significant revenue for all crops and both cultivation methods. Whether this is sufficient revenue to allow households to make a profit is difficult to determine because of the difficulty in valuing non-market purchased inputs, such as land, household labor, or recycled seeds. With these caveats in mind, the curves in Figure 6 tell a fairly consistent story, particularly when interpreted in combination with Figure 5. On average, and for maize and cowpea, CA cultivation generates less revenue than TC over a large portion of the rainfall distribution. For sorghum, millet, and groundnut, CA also tends to generate less revenue, though these differences are frequently not statistically significant. Where CA produces more revenue than TC is at the far ends of the distribution, when rainfall is more than one and a half standard deviations away from the mean. Research on other ‘climate-smart’ agricultural technologies has demonstrated the importance of immediate and consistent positive returns in order to sustain adoption (Chabé-Ferret and Subervie, 2013; Damon et al., 2015). The low levels of adoption, and high disadoption rates, of CA in Zimbabwe are likely a result of the technology providing low returns over much of the rainfall distribution. We conclude that CA may not be an appropriate technology for all of Zimbabwe, or any region of Sub-Saharan Africa where rainfall is not highly variable. Rather, policy should target CA for households living in areas prone to frequent severe drought or flooding.

¹⁸A strong caveat to this process is that we use price data from May 2018 in order to value production from the survey, which occurred between 2008 and 2011. The primary reason for this is that the surveys took place during and immediately after the period of hyperinflation in Zimbabwe and therefore did not collect price data but focused on the physical measurement of inputs and output. Because of these data limitations we are unable to calculate returns to adoption in terms of profit using prices from the survey period.

6 Conclusions and Policy Implications

Conservation agriculture has been widely promoted as a way for smallholder farmers in Sub-Saharan Africa to increase yields while also making yields more resilient to changing climate conditions. Using four years of panel data from Zimbabwe, we find evidence that contradicts the first claim but supports the later claim. We estimate a large complement of yield functions that include rainfall shocks, household fixed effects, and control for the endogeneity of the choice to adopt CA practices. In all these cases we find that, where CA has a significant impact on yields, it is always to reduce yields compared to TC practices during periods of average rainfall. When we consider yields during rainfall shocks, we find that yields tend to be more resilient under CA cultivation than under TC practices. Overall, returns to CA are only positive for those households that experience high variability in rainfall patterns. We conclude that previous econometric analysis that found a positive correlation between CA and yields was most likely due to a failure to control for unobserved heterogeneity among households and selection bias in the choice to adopt CA. This conclusion comes with the caveat that our data cover only four years. In the long run, CA may indeed have a positive impact on yield. To our knowledge, no long-run observational data set exists on which this hypothesis can be tested.

Two policy recommendations can be drawn from our analysis. First, our results help address the empirical puzzle of low CA adoption rates in Sub-Saharan Africa. We find that, over the four year study period, CA had either a negative impact or no impact at all on yields during periods of average rainfall. Smallholder farmers tend to be risk averse, and CA is often associated with increased labor demand and the need for purchased fertilizer inputs. Given that returns to CA can be negative, especially for maize, it makes sense that smallholders have been hesitant to undertake the added risk and cost of CA practices. We find no evidence at the plot-level that CA is associated with yield increases and therefore conclude that households' decision to not adopt, or disadopt as in the case of Zimbabwe, is most likely rational in the short-term. Policy to promote CA among smallholders should acknowledge this point and take steps to manage expectations regarding the short-term and long-term benefits of CA.

Second, CA can be effective in mitigating yield loss in environments with increased weather

risk. Climate change threatens to disrupt normal rainfall patterns by reducing the duration and frequency of rainfall (prolonged droughts) and also by increasing the intensity of rainfall. We find that in both cases (abnormally high and abnormally low rainfall) yields are often more resilient under CA than under TC. This insight provides a way forward for the promotion of CA practices among smallholders. Policy should be designed to focus on CA's potential benefits in mitigating risk due to changing rainfall patterns.

We conclude that CA is indeed an example of 'climate smart' agriculture to the extent that a changing climate will result in more abnormal rainfall patterns and CA appears effective in mitigating yield loss due to deviations in rainfall. Such a conclusion does not imply that CA is a sustainable approach to agriculture for all farmers, or even for most farmers, living in Sub-Saharan Africa. This is because in periods of normal rainfall the returns to CA are negative, at least in the short run. In order to test the long-term benefits of CA on yields through improved soil fertility, future research should focus on establishing long-run observational datasets on CA practices. The challenge here is to convince enough farmers to consistently adopt costly agricultural practices that may in the short-term cost them, absent the realization of extreme rainfall events.

References

- Abdulai, A. N. (2016). Impact of conservation agriculture technology on household welfare in Zambia. *Agricultural Economics* 47(6), 1–13.
- Anderson, T. and H. Rubin (1949). Estimation of the parameters of single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics* 20, 46–63.
- Andersson, J. A. and S. D’Souza (2014). From adoption claims to understanding farmers and contexts: A literature review of conservation agriculture (CA) adoption among smallholder farmers in Southern Africa. *Agriculture, Ecosystems & Environment* 187, 116–32.
- Arslan, A., N. McCarthy, L. Lipper, S. Asfaw, and A. Cattaneo (2014). Adoption and intensity of adoption of conservation farming practices in Zambia. *Agriculture, Ecosystems & Environment* 187, 72–86.
- Arslan, A., N. McCarthy, L. Lipper, S. Asfaw, A. Cattaneo, and M. Kokwe (2015). Climate smart agriculture? Assessing the adaptation implications in Zambia. *Journal of Agricultural Economics* 66(3), 753–80.
- Baker, J. M., T. E. Ochsner, R. T. Venterea, and T. J. Griffis (2007). Tillage and soil carbon sequestration — What do we really know? *Agriculture, Ecosystems & Environment* 118, 1–5.
- Barrett, C. B., C. M. Moser, O. V. McHugh, and J. Barison (2004). Better technology, better plots, or better farmers? Identifying changes in productivity and risk among Malagasy rice farmers. *American Journal of Agricultural Economics* 86(4), 2869–88.
- Baum, C. F., M. E. Schaffer, and S. Stillman (2007). Enhanced routines for instrumental variables/GMM estimation and testing. *Stata Journal* 7(4), 465–506.
- Bazzi, S. and M. A. Clemens (2013). Blunt instruments: Avoiding common pitfalls in identifying the causes of economic growth. *American Economic Journal: Macroeconomics* 5(2), 152–86.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119(1), 249–275.
- Bezu, S., G. T. Kassie, B. Shiferaw, and J. Ricker-Gilbert (2014). Impact of improved maize adoption on welfare of farm households in Malawi: A panel data analysis. *World Development* 59, 120–131.
- Brouder, S. M. and H. Gomez-Macpherson (2014). The impact of conservation agriculture on smallholder agricultural yields: A scoping review of the evidence. *Agriculture, Ecosystems & Environment* 187, 11–32.
- Bun, M. J. and T. D. Harrison (2018). OLS and IV estimation of regression models including endogenous interaction terms. *Forthcoming, Econometric Reviews*.
- Chabé-Ferret, S. and J. Subervie (2013). How much green for the buck? Estimating additional and windfall effects of French agro-environmental schemes by DID-matching. *Journal of Environmental Economics and Management* 65(1), 12–27.

- Chamberlain, G. (1984). Panel data. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics, Vol. 2*, pp. 1247–318. Amsterdam: North Holland.
- Damon, M., J. Graff Zivin, and H. Thirumurthy (2015). Health shocks and natural resource management: Evidence from Western Kenya. *Journal of Environmental Economics and Management* 69(1), 36–52.
- Di Falco, S. and J.-P. Chavas (2008). Rainfall shocks, resilience, and the effects of crop biodiversity on agroecosystem productivity. *Land Economics* 84(1), 83–96.
- FAO (2013). Climate-smart agriculture sourcebook. Technical report, Food and Agriculture Organization of the United Nations, Rome.
- FAO, IFAD, UNICEF, WFP, and WHO (2017). The state of food security and nutrition in the world 2017: Building resilience for peace and food security. Technical report, Food and Agriculture Organization of the United Nations, Rome.
- Funk, C., P. Peterson, M. Landseld, D. Pedreros, J. Verdin, S. Shukla, G. Husak, J. Roland, L. Harrison, A. Hoell, and J. Michaelsen (2015). The climate hazards infrared precipitation with stations – a new environmental record for monitoring extremes. *Scientific Data* 2(150066).
- Giller, K. E., M. Corbeels, J. Nyamangara, B. Triomphe, F. Affholder, E. Scopel, and P. Titttonell (2011). A research agenda to explore the role of conservation agriculture in African smallholder farming systems. *Field Crops Research* 124, 468–72.
- Giller, K. E., E. Witter, M. Corbeels, and P. Titttonell (2009). Conservation agriculture and smallholder farming in Africa: The heretics’ view. *Field Crops Research* 114, 23–34.
- Greene, W. H. (2011). *Econometric Analysis* (7th ed.). Upper Saddle River: Pearson.
- Jennings, M., A. Kayondo, J. Kagoro, K. Nicholson, N. Blight, and J. Gayfer (2013). Impact evaluation of the Protracted Relief Programme II, Zimbabwe. Technical report, International Organisation Development Ltd., Sheffield.
- Just, R. E. and R. D. Pope (1978). Stochastic specification of production functions and economic implications. *Journal of Econometrics* 7(1), 67–86.
- Kassie, M., H. Teklewold, M. Jaletab, P. Marenayab, and O. Erenstein (2015). Understanding the adoption of a portfolio of sustainable intensification practices in eastern and southern Africa. *Land Use Policy* 42, 400–11.
- Kassie, M., H. Teklewold, P. Marenya, M. Jaleta, and O. Erenstein (2015). Production risks and food security under alternative technology choices in Malawi: Application of a multinomial endogenous switching regression. *Journal of Agricultural Economics* 66(3), 640–59.
- Kassie, M., P. Zikhali, K. Manjur, and S. Edwards (2009). Adoption of sustainable agricultural practices: Evidence from a semi-arid region of Ethiopia. *Natural Resources Forum* 33, 189–98.
- Kassie, M., P. Zikhali, J. Pender, and G. Köhlin (2010). The economics of sustainable land management practices in the Ethiopian highlands. *Journal of Agricultural Economics* 61(3), 605–27.

- Kleibergen, F. and R. Paap (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics* 133(1), 97–126.
- Lal, R. and B. Stewart (2010). *Food Security and Soil Quality*. Boca Raton: CRC Press.
- Lipper, L., P. Thornton, B. M. Campbell, T. Baedeker, A. Braimoh, M. Bwalya, P. Caron, A. Cattaneo, D. Garrity, K. Henry, R. Hottle, L. Jackson, A. Jarvis, F. Kossam, W. Mann, N. McCarthy, A. Meybeck, H. Neufeldt, T. Remington, P. T. Sen, R. Sessa, R. Shula, A. Tibu, and E. F. Torquebiau (2014). Climate-smart agriculture for food security. *Nature Climate Change* 4, 1068–72.
- Lobell, D. B., W. Schlenker, and J. Costa-Roberts (2011). Climate trends and global crop production since 1980. *Science* 333, 616–20.
- Manda, J., A. D. Alene, C. Gardebroke, M. Kassie, and G. Tembo (2016). Adoption and impacts of sustainable agricultural practices on maize yields and incomes: Evidence from rural Zambia. *Journal of Agricultural Economics* 67(1), 130–53.
- Mazvimavi, K. and S. Twomlow (2009). Socioeconomic and institutional factors influencing adoption of conservation farming by vulnerable households in Zimbabwe. *Agricultural Systems* 101, 20–29.
- Mazvimavi, K., S. Twomlow, P. Belder, and L. Hove (2008). An assessment of the sustainable adoption of conservation farming in Zimbabwe. Global Theme on Agroecosystems, Report number 39, ICRISAT, Bulawayo, Zimbabwe.
- Michler, J. D. and G. E. Shively (2015). Land tenure, tenure security and farm efficiency: Panel evidence from the Philippines. *Journal of Agricultural Economics* 66(1), 155–69.
- Morton, J. F. (2007). The impact of climate change on smallholder and subsistence agriculture. *Proceedings of the National Academy of Science* 104(50), 19,680–5.
- Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica* 46(1), 69–85.
- Ndlovu, P. V., K. Mazvimavi, H. An, and C. Murendo (2014). Productivity and efficiency analysis of maize under conservation agriculture in Zimbabwe. *Agricultural Systems* 124, 21–31.
- Pannell, D. J., R. S. Llewellyn, and M. Corbeels (2014). The farm-level economics of conservation agriculture for resource-poor farmers. *Agriculture, Ecosystems & Environment* 187, 52–64.
- Pedzisa, T., L. Rugube, A. Winter-Nelson, K. Baylis, and K. Mazvimavi (2015a). Abandonment of conservation agriculture by smallholder farmers in Zimbabwe. *Journal of Sustainable Development* 8(1).
- Pedzisa, T., L. Rugube, A. Winter-Nelson, K. Baylis, and K. Mazvimavi (2015b). The intensity of adoption of conservation agriculture by smallholder farmers in Zimbabwe. *Agrekon* 54(3), 1–22.
- Piccoli, I., F. Chiarini, P. Carletti, L. Furlan, B. Lazzaro, S. Nardi, A. Berti, L. Sartori, M. Dalconi, and F. Morari (2016). Disentangling the effects of conservation agriculture practices on the vertical distribution of soil organic carbon. evidence of poor carbon sequestration in North-Eastern Italy. *Agriculture, Ecosystems & Environment* 230, 68–78.

- Pittelkow, C. M., X. Liang, B. A. Linquist, K. J. van Groenigen, J. Lee, M. E. Lundy, N. van Gestel, J. Six, R. T. Ventura, and C. van Kessel (2015). Productivity limits and potentials of the principles of conservation agriculture. *Nature* 517, 365–8.
- Powlson, D. S., C. M. Stirling, C. Thierfelder, R. P. White, and M. Jat (2016). Does conservation agriculture deliver climate change mitigation through soil carbon sequestration in tropical agro-ecosystems? *Agriculture, Ecosystems & Environment* 220, 164–74.
- Ricker-Gilbert, J., T. S. Jayne, and E. Chirwa (2011). Subsidies and crowding out: A double-hurdle model of fertilizer demand in Malawi. *American Journal of Agricultural Economics* 93(1), 26–42.
- Sanderson, E. and F. Windmeijer (2016). A weak instrument F -test in linear IV models with multiple endogenous variables. *Journal of Econometrics* 190, 212–21.
- Schlenker, W. and D. B. Lobell (2010). Robust negative impacts of climate change on African agriculture. *Environmental Resource Letters* 5(1), 014010. (pp. 8).
- Schuller, P., D. E. Walling, A. Sepúlveda, A. Castillo, and I. Pino (2007). Changes in soil erosion associated with the shift from conventional tillage to a no-tillage system, documented using ^{137}Cs measurements. *Soil & Tillage Research* 94(1), 183–92.
- Sheffield, J., E. Wood, N. Chaney, K. Guan, S. Sadri, X. Yuan, L. Olang, A. Amani, A. Ali, S. Demuth, and L. Ogallo (2014). A drought monitoring and forecasting system for Sub-Saharan African water resources and food security. *Bulletin of the American Meteorological Society* 95(6), 861–82.
- Smith, P., D. S. Powlson, M. J. Glendining, and J. U. Smith (1998). Preliminary estimates of the potential for carbon mitigation in European soils through no-till farming. *Global Change Biology* 4(6), 679–85.
- Stock, J. H. and M. Yogo (2005). Testing for weak instruments in linear IV regression. In D. W. Andrews and J. H. Stock (Eds.), *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pp. 80–108. New York: Cambridge University Press.
- Suri, T. (2011). Selection and comparative advantage in technology adoption. *Econometrica* 79(1), 159–209.
- Teklewold, H., M. Kassie, and B. Shiferaw (2013). Adoption of multiple sustainable agricultural practices in rural Ethiopia. *Journal of Agricultural Economics* 64(3), 597–623.
- Teklewold, H., M. Kassie, B. Shiferaw, and G. Köhlin (2013). Cropping system diversification, conservation tillage, and modern seed adoption in Ethiopia: Impacts on household income, agro-chemical use, and demand for labor. *Ecological Economics* 93, 85–93.
- Thierfelder, C. and P. C. Wall (2009). Effects of conservation agriculture techniques on infiltration and soil water content in Zambia and Zimbabwe. *Soil & Tillage Research* 105, 217–27.
- Verkaart, S., B. G. Munyua, K. Mausch, and J. D. Michler (2017). Welfare impacts of improved chickpea adoption: A pathway for rural development in Ethiopia? *Food Policy* 66, 50–61.

- Ward, P. S. and G. E. Shively (2015). Migration and land rental as responses to income shocks in rural China. *Pacific Economic Review* 20(4), 511–43.
- Wheeler, T. and J. von Braun (2013). Climate change impacts on global food security. *Science* 341, 508–13.
- Wooldridge, J. M. (2003). Further results on instrumental variables estimation of average treatment effects in the correlated random coefficient model. *Economic Letters* 79(2), 185–191.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. Cambridge: MIT Press.

Table 1: Descriptive Statistics by Crop

	Maize	Sorghum	Millet	Groundnut	Cowpea	Total
Yield (kg/ha)	1,217 (1,366)	827.7 (1,137)	641.0 (946.7)	1,065 (1,251)	662.4 (1,137)	1,040 (1,284)
CA (= 1)	0.350 (0.477)	0.263 (0.440)	0.137 (0.345)	0.193 (0.395)	0.313 (0.464)	0.290 (0.454)
Basal applied fertilizer (kg)	13.41 (30.39)	2.281 (10.35)	0.923 (5.952)	1.429 (7.653)	3.457 (12.19)	7.716 (23.21)
Top applied fertilizer (kg)	17.58 (32.56)	3.366 (10.69)	1.636 (8.786)	1.843 (9.264)	4.112 (13.11)	10.16 (25.32)
Seed (kg)	8.143 (9.169)	4.896 (6.460)	5.974 (11.40)	10.96 (14.80)	3.099 (4.047)	7.543 (10.20)
Area (m ²)	3,466 (3,782)	3,225 (3,697)	3,978 (4,138)	2,060 (2,100)	1,553 (1,894)	3,035 (3,486)
Rainfall shock	0.469 (0.274)	0.482 (0.267)	0.551 (0.301)	0.469 (0.273)	0.463 (0.258)	0.476 (0.274)
Number of HH in ward with NGO support	20.62 (13.09)	22.01 (15.51)	24.75 (17.62)	21.77 (14.12)	23.29 (15.15)	21.56 (14.25)
number of observations	3,827	1,264	488	1,397	667	7,643
number of plots	2,643	1,007	405	1,220	615	4,171
number of households	715	415	177	598	388	728
number of wards	45	41	26	45	43	45

Note: The first five columns of the table display means of the data by crop with standard deviations in parenthesis. The final column displays means and standard deviations for the pooled data.

Table 2: Descriptive Statistics by Year

	2008	2009	2010	2011
Yield (kg/ha)	760.6 (1,210)	1,278 (1,355)	1,151 (1,415)	936.0 (1,094)
CA (= 1)	0.435 (0.496)	0.385 (0.487)	0.247 (0.431)	0.170 (0.376)
Basal applied fertilizer (kg)	7.044 (20.92)	3.898 (14.73)	4.335 (15.26)	14.00 (32.48)
Top applied fertilizer (kg)	8.679 (20.16)	6.776 (16.82)	7.331 (20.15)	16.14 (34.80)
Seed (kg)	7.961 (10.02)	6.399 (7.737)	7.133 (12.29)	8.499 (9.729)
Area planted (m ²)	3,473 (4,172)	2,870 (3,896)	2,891 (3,348)	3,015 (2,714)
Rainfall shock	0.682 (0.289)	0.487 (0.267)	0.351 (0.225)	0.453 (0.231)
Number of HH in ward with NGO support	27.87 (19.20)	18.00 (10.38)	21.63 (13.34)	20.21 (12.58)
number of observations	1,452	1,732	2,116	2,343
number of plots	1,403	1,677	2,015	2,312
number of households	388	401	432	584
number of wards	29	30	31	43

Note: Columns in the table display means of the data by year with standard deviations in parenthesis. Inputs are measured at the plot-level.

Table 3: Descriptive Statistics by Crop and Cultivation Method

	Maize			Sorghum			Millet			Groundnut			Cowpea		
	TC	CA	MW-test												
Yield(kg/ha)	932.6 (1,127)	1,745 (1,595)	***	736.8 (1,073)	1,082 (1,267)	***	627.7 (979.8)	724.1 (703.2)	**	978.4 (1,117)	1,428 (1,665)	***	626.5 (1,153)	740.9 (1,101)	***
Basal applied fertilizer (kg)	11.46 (31.01)	17.03 (28.86)	***	1.149 (8.582)	5.457 (13.73)	***	0.705 (6.151)	2.291 (4.284)	***	0.547 (4.424)	5.106 (14.32)	***	1.417 (7.236)	7.925 (18.22)	***
Top applied fertilizer (kg)	15.11 (32.24)	22.18 (32.65)	***	0.996 (4.346)	10.01 (17.96)	***	0.724 (5.712)	7.365 (17.98)	***	0.591 (4.137)	7.065 (18.43)	***	1.298 (6.539)	10.27 (20.01)	***
Seed planted (kg)	9.287 (10.61)	6.021 (4.931)	***	5.206 (5.776)	4.023 (8.023)	***	6.477 (12.15)	2.812 (2.750)	***	12.38 (16.01)	5.027 (4.570)	***	2.957 (4.432)	3.408 (3.023)	***
Area planted (m ²)	3,988 (4,299)	2,496 (2,257)	***	3,809 (4,038)	1,582 (1,619)	***	4,465 (4,227)	915.1 (1,249)	***	2,260 (2,191)	1,223 (1,386)	***	1,782 (2,119)	1,051 (1,116)	***
Rainfall shock	0.465 (0.273)	0.473 (0.275)		0.473 (0.262)	0.503 (0.275)		0.547 (0.309)	0.568 (0.241)		0.455 (0.266)	0.527 (0.293)	***	0.466 (0.260)	0.453 (0.252)	
Number of HH in ward with NGO support	19.50 (12.85)	22.68 (13.26)	***	21.23 (14.69)	24.18 (17.42)	**	23.26 (17.01)	34.13 (18.52)	***	20.41 (13.32)	27.40 (15.88)	***	21.54 (14.31)	27.11 (16.21)	***
number of observations	2,486	1,341		932	332		421	67		1,127	270		458	209	
number of plots	2,001	966		792	277		357	58		1,015	254		437	198	
number of households	670	537		369	206		168	44		562	188		306	153	
number of wards	45	44		40	30		26	10		45	24		43	30	

Note: Columns in the table display means of the data by crop with standard deviations in parenthesis. Columns headed TC are output and inputs used under traditional cultivation practices while columns headed CA are output and inputs used under conservation agriculture. The final column for each crop presents the results of Mann-Whitney two-sample tests for differences in distribution. Results are similar if a Kolmogorov-Smirnov test is used. Significance of MW-tests are reported as *p<0.1; **p<0.05; ***p<0.01.

Table 4: Input Change Over Time Conditional on Ward

	2008			2011		
	Always adopter	Future disadopter	MW-test	Always adopter	Future disadopter	MW-test
ln(Yield) (kg/ha)	0.469 (1.190)	0.386 (1.270)		0.512 (0.939)	-0.028 (1.113)	***
ln(Seed) (kg/ha)	0.089 (0.802)	0.206 (0.883)	*	0.029 (0.717)	-0.054 (0.674)	*
ln(Basal fertilizer) (kg/ha)	0.122 (0.799)	0.009 (0.851)		0.188 (0.912)	-0.312 (0.835)	***
ln(Top fertilizer) (kg/ha)	0.201 (0.846)	0.052 (0.857)		0.363 (0.974)	-0.242 (0.940)	***
Observations	162	299		86	346	
	Never adopter	Future adopter	MW-test	Never adopter	Future adopter	MW-test
ln(Yield) (kg/ha)	-0.394 (1.281)	-0.239 (1.366)		-0.154 (1.188)	0.455 (1.064)	***
ln(Seed) (kg/ha)	-0.172 (0.944)	-0.001 (0.831)		0.012 (0.769)	0.132 (0.857)	*
ln(Basal fertilizer) (kg/ha)	-0.298 (1.198)	0.054 (0.597)		-0.127 (0.813)	0.125 (1.109)	
ln(Top fertilizer) (kg/ha)	-0.415 (1.057)	0.068 (0.778)	**	-0.190 (0.873)	0.285 (1.157)	***
Observations	446	84		656	75	

Note: Table displays the mean residuals of output and inputs by adoption type and year. Residuals are calculated from either a regression of output on ward indicators or the input on ward indicators. In the upper panel, “Always adopters” are those who in every year use CA. They are compared to “Future disadopters,” those households who use CA in 2008 but disadopt in subsequent years. In the lower panel, “Never adopters” are those who in every year do not use CA. They are compared to “Future adopters,” those households who do not use in 2008 but adopt in subsequent years. The final column for each year presents the results of Mann-Whitney two-sample tests for differences in distribution. Results are similar if a Kolmogorov-Smirnov test is used. Significance of MW-tests are reported as *p<0.1; **p<0.05; ***p<0.01.

Table 5: Yield Function with CA as Exogenous

	(1)	(2)	(3)	(4)
<i>Maize</i>				
CA (= 1)	0.631*** (0.081)	0.573*** (0.081)	0.222 (0.139)	0.207 (0.137)
rainfall shock			-0.667*** (0.215)	-0.904*** (0.227)
CA × rainfall shock			0.872*** (0.247)	0.744*** (0.237)
<i>Sorghum</i>				
CA (= 1)	0.041 (0.180)	-0.051 (0.190)	-0.756*** (0.271)	-0.596** (0.288)
rainfall shock			-1.285*** (0.303)	-1.432*** (0.317)
CA × rainfall shock			1.652*** (0.458)	1.130** (0.499)
<i>Millet</i>				
CA (= 1)	-0.145 (0.359)	0.118 (0.374)	-0.994 (0.709)	-0.711 (0.736)
rainfall shock			-1.094*** (0.306)	-1.521*** (0.319)
CA × rainfall shock			1.492 (0.940)	1.475 (1.171)
<i>Groundnut</i>				
CA (= 1)	0.299** (0.142)	0.323** (0.144)	-0.326 (0.287)	0.130 (0.279)
rainfall shock			-0.403* (0.208)	-0.489** (0.228)
CA × rainfall shock			1.106** (0.493)	0.331 (0.456)
<i>Cowpea</i>				
CA (= 1)	-0.035 (0.241)	0.194 (0.250)	-0.844* (0.466)	-0.362 (0.460)
rainfall shock			-1.053** (0.452)	-1.216*** (0.448)
CA × rainfall shock			1.688* (0.905)	1.114 (0.878)
Household FE	No	Yes	No	Yes
Observations	7,643	7,643	7,643	7,643
R^2	0.899	0.922	0.900	0.923

Note: Dependent variable is log of yield. Though not reported, all specifications include crop-specific inputs and intercept terms, and year dummies. See Table B1 in the Appendix for coefficient estimates of crop-specific inputs. Column (1) excludes the rainfall variable as well as household fixed effects. Column (2) excludes the rainfall variable but includes household fixed effects. Column (3) includes the rainfall variable and its interaction with CA but excludes household fixed effects. Column (4) includes both the rainfall variable, its interaction with CA, and household fixed effects. Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table 6: Zero-Stage Probit

	(1)	(2)	(3)	(4)
Number of HH in ward with NGO support	0.010*** (0.001)	0.009*** (0.001)	0.010*** (0.001)	0.009*** (0.001)
Household MCD	No	Yes	No	Yes
Observations	7,643	7,643	7,643	7,643
Log Likelihood	-3,153	-3,110	-3,143	-3,100

Note: Dependent variable is an indicator for whether or not CA was used on the plot. Though not reported, all probit regressions include crop-specific inputs and intercept terms, and year dummies. Column (1) excludes the rainfall variable as well as the Mundlak-Chamberlain device (MCD). Column (2) excludes the rainfall variable but includes the MCD. Column (3) includes the rainfall variable but excludes the MCD. Column (4) includes both the rainfall variable and the MCD. Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table 7: Yield Function with CA as Endogenous

	(1)	(2)	(3)	(4)
<i>Maize</i>				
CA (= 1)	13.404 (10.916)	-1.939* (1.122)	15.697 (18.025)	-2.853** (1.141)
rainfall shock			0.569 (2.198)	-1.675*** (0.353)
CA × rainfall shock			1.335 (2.896)	2.628*** (0.704)
<i>Sorghum</i>				
CA (= 1)	19.230 (13.886)	-0.275 (1.498)	24.937 (24.052)	-0.171 (1.603)
rainfall shock			-0.940 (0.823)	-1.476*** (0.353)
CA × rainfall shock			-2.555 (4.717)	0.899 (0.904)
<i>Millet</i>				
CA (= 1)	32.388 (30.713)	-4.551 (3.402)	49.502 (64.115)	-3.784 (3.808)
rainfall shock			0.654 (2.163)	-1.715*** (0.341)
CA × rainfall shock			-12.938 (30.199)	3.357 (2.279)
<i>Groundnut</i>				
CA (= 1)	14.997 (10.282)	0.050 (1.199)	16.999 (16.357)	-0.272 (1.324)
rainfall shock			-1.659* (0.886)	-0.950*** (0.265)
CA × rainfall shock			2.423 (2.895)	1.433* (0.832)
<i>Cowpea</i>				
CA (= 1)	10.530 (8.312)	-0.938 (1.147)	12.451 (14.121)	-1.564 (1.349)
rainfall shock			-0.061 (1.643)	-1.606*** (0.485)
CA × rainfall shock			0.775 (3.789)	2.123 (1.404)
Household FE	No	Yes	No	Yes
Observations	7,643	7,643	7,643	7,643
Log Likelihood	-24,462	-16,378	-25,914	-16,192
Kleibergen-Paap LM stat	1.773	26.89***	1.028	33.86***
Anderson-Rubin Wald stat	37.38***	19.47***	75.09***	35.88***
Kleibergen-Paap Wald stat	0.355	4.593***	0.102	2.955***

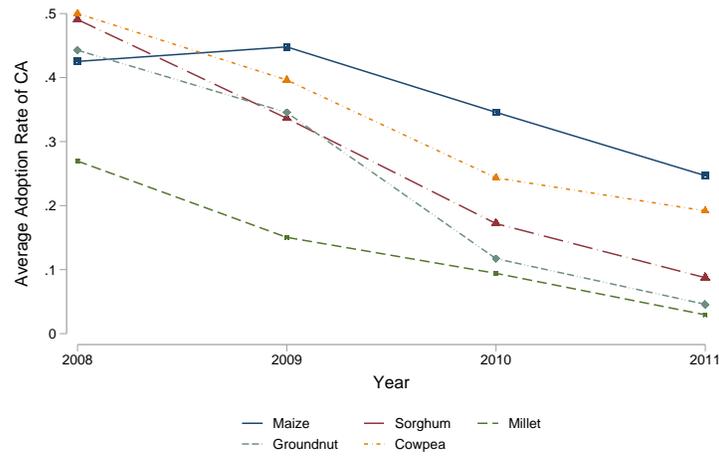
Note: Dependent variable is log of yield. Though not reported, all specifications include crop-specific inputs and intercept terms, and year dummies. See Table B2 in the Appendix for coefficient estimates of crop-specific inputs. In each regression the adoption of CA is treated as endogenous and is instrumented with the Inverse Mills Ratio (IMR) calculated from the predicted values of the zero-stage probits reported in Table 6. The CA × rainfall shock term is also treated as endogenous and instrumented using the interaction of the IMR and the rainfall shock term. The null hypothesis of the Kleibergen-Paap LM test is that the rank condition fails (i.e., the first-stage equation is underidentified). The null hypothesis of the Anderson-Rubin Wald test is that the coefficients on the endogenous regressors in the structural equation are jointly equal to zero (i.e., the instruments in the first-stage equation are weak). The null hypothesis of the Kleibergen-Paap Wald test is that a *t*-test at the 5% significance level on the coefficients of the endogenous regressors rejects no more than 25% of the time (i.e., the instruments in the first-stage equation are weak). Standard errors clustered by household and crop are reported in parentheses (**p*<0.1; ***p*<0.05; ****p*<0.01).

Table 8: Yield Function with Rain Shortage or Surplus

	(1)	(2)	(3)	(4)
<i>Maize</i>				
CA (= 1)	-0.466 (1.379)	-2.657* (1.378)	-1.893 (1.472)	3.393 (3.156)
rainfall shortage	0.089 (0.400)		-1.076** (0.481)	-1.411** (0.695)
CA × rainfall shortage	0.147 (0.744)		2.098** (0.987)	3.522** (1.597)
rainfall surplus		-1.619*** (0.304)	-1.884*** (0.358)	-0.444 (0.637)
CA × rainfall surplus		2.012*** (0.513)	2.931*** (0.657)	1.953** (0.845)
<i>Sorghum</i>				
CA (= 1)	-0.262 (1.464)	-1.211 (2.042)	-0.512 (1.780)	5.808 (4.791)
rainfall shortage	-1.079*** (0.389)		-1.702*** (0.454)	-2.235*** (0.769)
CA × rainfall shortage	2.437* (1.263)		2.356* (1.326)	3.654* (1.977)
rainfall surplus		-0.550 (0.342)	-1.214*** (0.377)	-0.531 (0.533)
CA × rainfall surplus		-0.369 (1.008)	0.464 (0.987)	-0.144 (1.507)
<i>Millet</i>				
CA (= 1)	-2.913 (2.890)	-3.914 (3.404)	-2.671 (3.293)	12.779 (12.606)
rainfall shortage	-0.508 (0.610)		-1.604** (0.658)	-1.438 (0.927)
CA × rainfall shortage	-0.136 (1.576)		1.930 (2.344)	-4.318 (8.524)
rainfall surplus		-1.287*** (0.356)	-1.689*** (0.320)	-0.427 (0.584)
CA × rainfall surplus		2.726* (1.548)	3.170 (2.202)	-0.775 (7.605)
<i>Groundnut</i>				
CA (= 1)	0.359 (1.159)	-1.519 (1.550)	-0.266 (1.499)	4.764 (3.664)
rainfall shortage	-1.125*** (0.318)		-1.229*** (0.338)	-2.298** (0.934)
CA × rainfall shortage	0.240 (0.772)		0.405 (0.984)	1.499 (1.570)
rainfall surplus		0.036 (0.297)	-0.170 (0.281)	0.506 (0.465)
CA × rainfall surplus		0.056 (0.899)	-0.205 (0.985)	-0.607 (1.422)
<i>Cowpea</i>				
CA (= 1)	-0.742 (1.141)	-0.929 (1.415)	-1.188 (1.439)	1.829 (3.818)
rainfall shortage	-0.628 (0.477)		-1.493*** (0.553)	-2.468** (1.221)
CA × rainfall shortage	3.171*** (1.008)		3.763*** (1.272)	6.006** (2.558)
rainfall surplus		-1.103** (0.482)	-1.492*** (0.557)	-0.970 (1.002)
CA × rainfall surplus		-0.286 (1.309)	1.234 (1.624)	3.148 (2.543)
Type of Effect	Household FE	Household FE	Household FE	Plot MCD
Observations	7,643	7,643	7,643	5,004
Log Likelihood	-15,908	-16,406	-15,902	-12,268
Kleibergen-Paap LM stat	22.96***	23.02***	25.23***	3.790**
Anderson-Rubin Wald stat	22.10**	29.68***	44.79***	38.53***
Kleibergen-Paap Wald stat	1.930*	1.989*	1.444	0.251

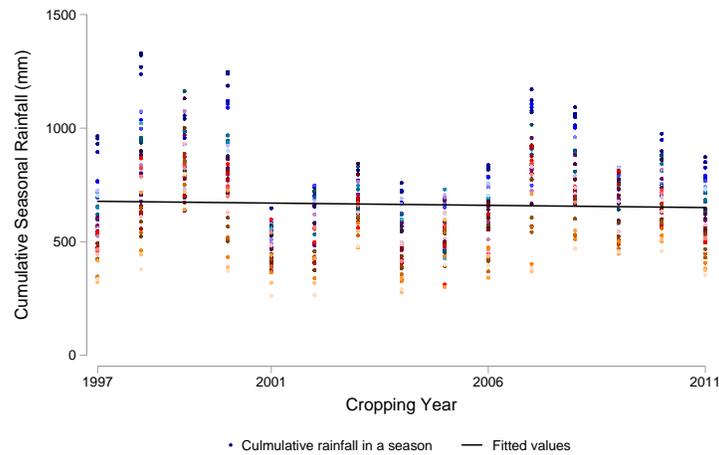
Note: Dependent variable is log of yield. Though not reported, all specifications include crop-specific inputs and intercept terms, and year dummies. See Table B4 in the Appendix for coefficient estimates of crop-specific inputs. In each regression the adoption of CA is treated as endogenous and is instrumented with the Inverse Mills Ratio (IMR) calculated from the predicted values of zero-stage probits which are presented in the Appendix, Table B3. The CA × rainfall shortage and CA × rainfall surplus terms are also treated as endogenous and instrumented using the interaction of the IMR and the rainfall terms. The null hypothesis of the Kleibergen-Paap LM test is that the rank condition fails (i.e., the first-stage equation is underidentified). The null hypothesis of the Anderson-Rubin Wald test is that the coefficients on the endogenous regressors in the structural equation are jointly equal to zero (i.e., the instruments in the first-stage equation are weak). The null hypothesis of the Kleibergen-Paap Wald test is that a *t*-test at the 5% significance level on the coefficients of the endogenous regressors rejects no more than 25% of the time (i.e., the instruments in the first-stage equation are weak). Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Figure 1: Average Annual Level of CA Adoption by Crop



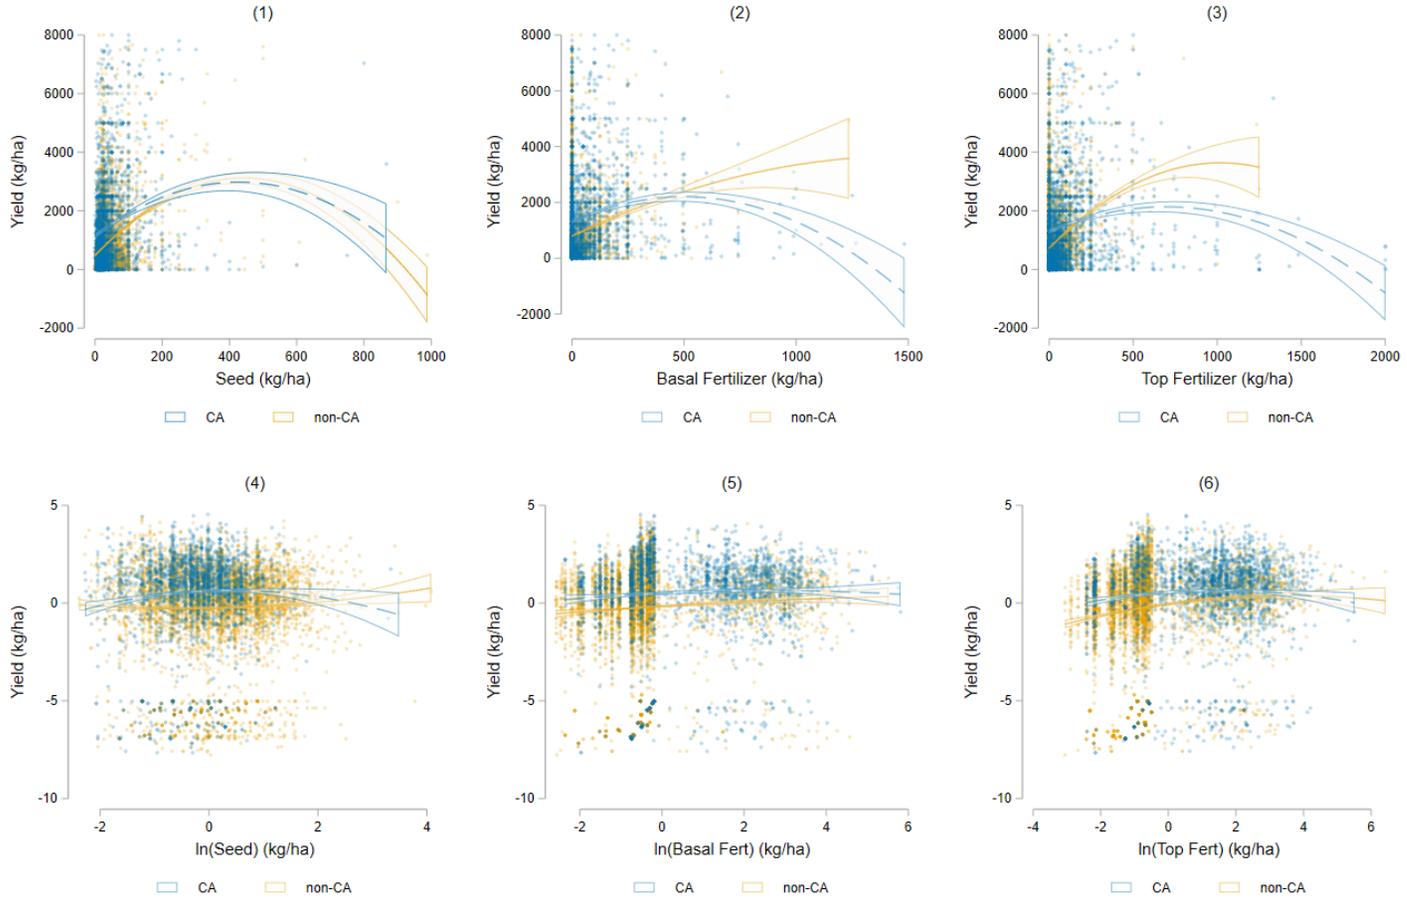
Note: Figure displays the average number of plots on which CA techniques are used by crop per year.

Figure 2: Historic Seasonal Rainfall by Ward



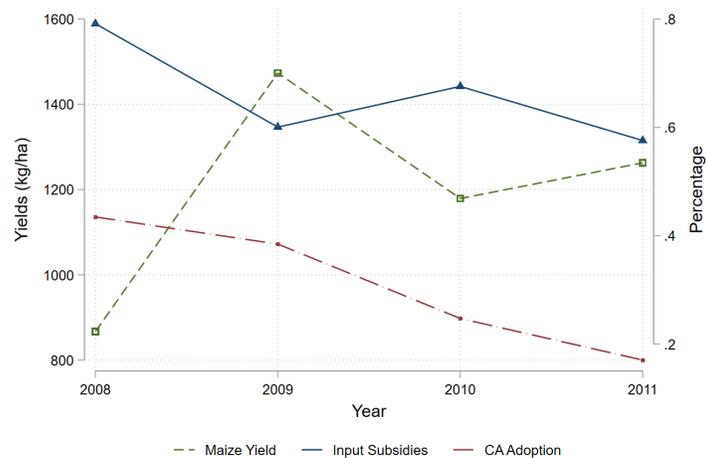
Note: Figure displays the cumulative rainfall for each cropping season in each ward for the 15 year period from 1997-2011. Colors represent wards and each dot is cumulative rainfall in a ward in a season. A linear trendline is fitted to the data with a slope of -3.71 which is significantly different from zero at the 99 percent level. Data comes from the CHIRPS database (Funk et al., 2015).

Figure 3: Yield Response to Inputs by CA Adoption



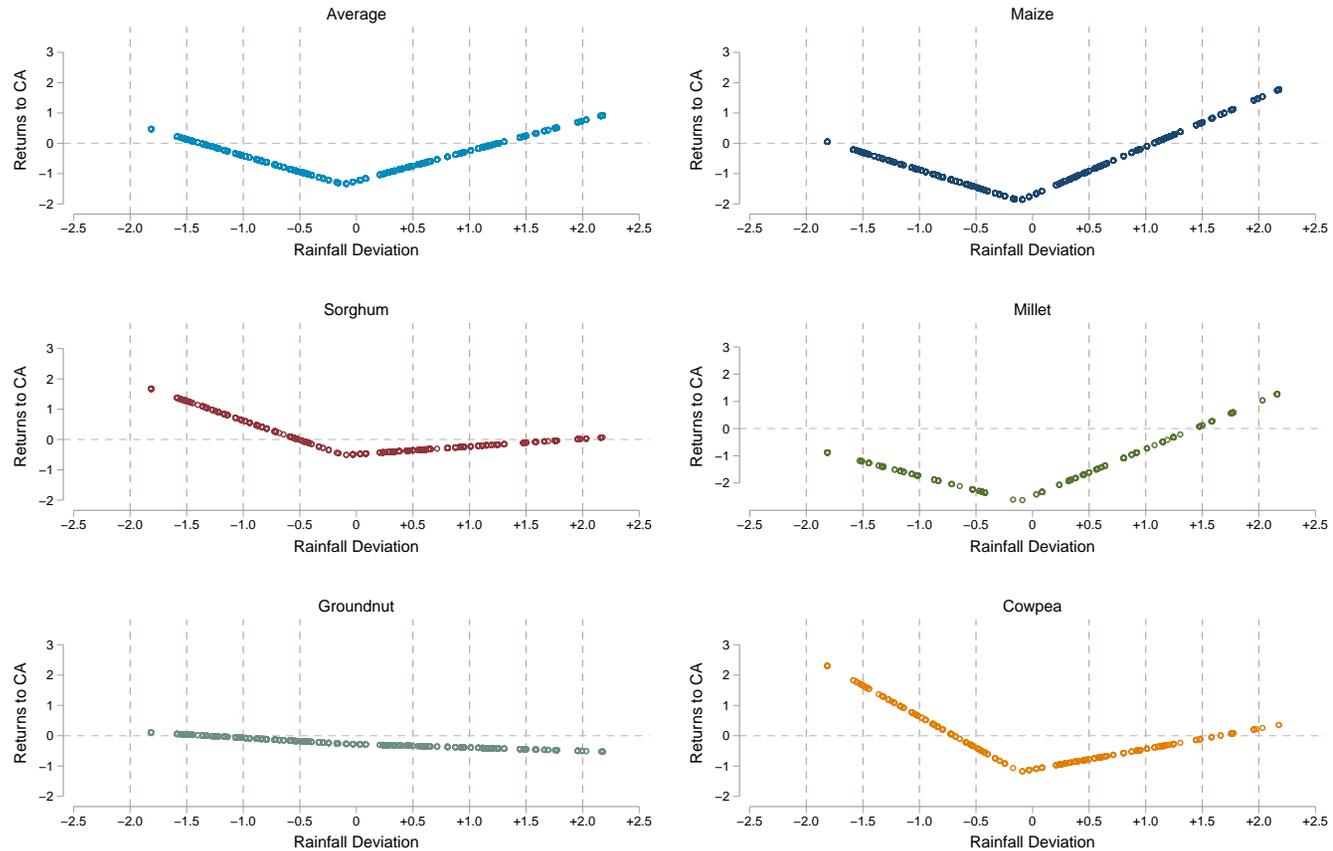
Note: Graphs in the top row display scatter plots of yields and input use per hectare (seeds, basal applied fertilizer, and top applied fertilizer). Blue diamonds represent observations from CA plots and yellow circles represent observations from non-CA plots. To these graphs are added a quadratic predicted fit line with 90% confidence intervals. Blue dashed lines are the fit of output and input use for CA plots and yellow solid lines are the fit from for non-CA plots. Graphs in the bottom row display scatter plots using residuals from a regression of output on ward indicators and each input on ward indicators.

Figure 4: Maize Yields, Rates of CA Adoption, and Input Subsidies



Note: Figure displays average maize yield on the left axis. On the right axis is the average rate of CA adoption at the plot-level along with the average percentage of households receiving input subsidies within a ward.

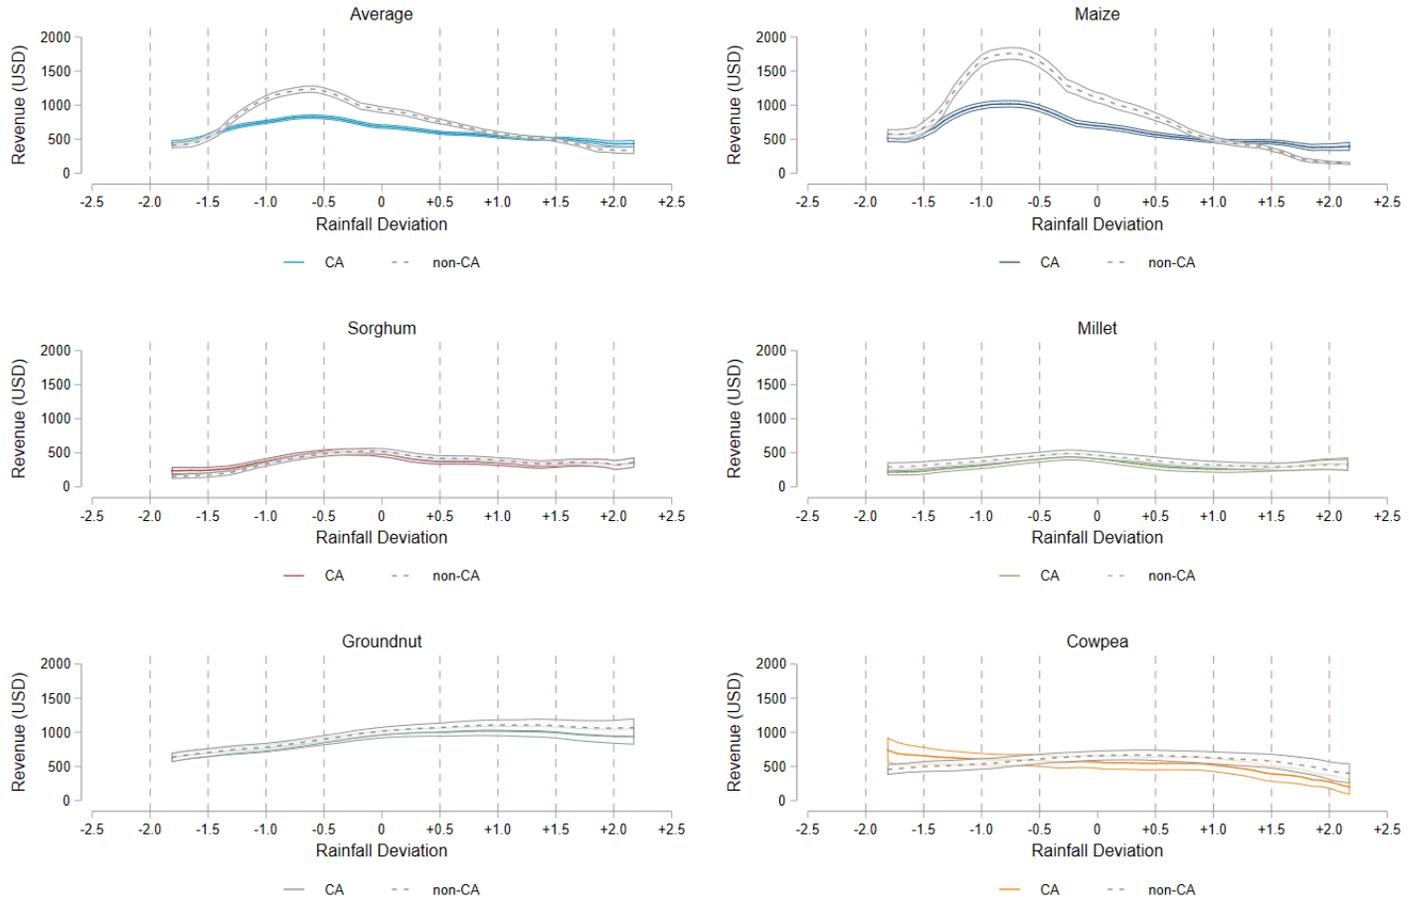
Figure 5: Predicted Returns to CA by Crop



77

Note: Figure displays the predicted returns to CA averaged over all crops and also by crop. Returns are calculated using the coefficients from column (3) in Table 8. By crop, we multiply the coefficient on the $CA \times$ rainfall shortage by realized shortages and multiply the coefficient on the $CA \times$ rainfall surplus by realized surpluses. We then sum these values along with the coefficient on the CA term. For the average return we take a weighted average of the crop specific returns where the weights are the number of observations for each crop in the data. The x-axis marks the realizations of rainfall surpluses and shortages where shortages are to the left of zero. Dots represent individual predicted returns where a lower density of dots represents fewer observations, and thus less confidence. Vertical lines are drawn at \pm half, one, one half, and two standard deviations from the mean, which is 0.062.

Figure 6: Predicted Revenue to CA and non-CA by Crop



45

Note: Figure displays the predicted revenue to CA and non-CA averaged over all crops and also by crop. First, we predicted \hat{y} for CA and non-CA plots from the regression presented in column (3) in Table 8. Next, we assign a monetary value to yields for each crop using prices published by Zimbabwe’s grain marketing board. We then graph the revenue at different realizations of rainfall surpluses and shortages using kernel-weighted local polynomial smoothing. The x-axis marks the realizations of rainfall surpluses and shortages where shortages are to the left of zero. Vertical lines are drawn at \pm half, one, one half, and two standard deviations from the mean, which is 0.062.

Online-Only Appendix to “Conservation Agriculture and Climate Resilience”

A Crop Production & Rainfall Histories in Zimbabwe

This appendix contains information regarding the distribution of observations in our unbalanced panel, historical rainfall patterns in Zimbabwe, and a discussion of potential concerns about the correlation between weather variables and output.

Our analysis relies on unbalanced panel data. While in most observational data sets, a lack of balance in a panel is a result of attrition, in our case the lack of balance is a result of adding households to the survey over time. Because of this, we have observations from more households (and more plots) in later years. Since our sample is random, we do not expect there to be any differences in the types of households added to the survey in later years. However, weather events are not evenly distributed across time. Thus, to some extent, our overall results are driven by the weather that was occurring in 2010 and 2011 because that is when a majority of the data comes from. Figures [A1-A3](#) provide the reader with an idea to what extent our results are driven by the unbalanced distribution across years.

Figure [A1](#) presents a bar graph of the frequency of plot-level observations for each crop in each year. The number of observations of plots with maize growing on them increases consistently from year-to-year. A similar pattern exists for observations of plots with groundnut and cowpea. While the entire sample size increases each year, the number of plots with sorghum and millet does not always increase. Figure [A2](#) graphs the percentage of observations for each crop in each year. Here we see that while the total number of observations changed from year-to-year, the percentage breakdown for each crop was fairly stable. Maize makes up between 47 and 54 percent of plot-level observations each year. Sorghum makes up between 12 and 19 percent of plot-level observations each year. Similarly narrow bands of variation exist for the other three crops. Taken together, Figures [A1](#) and [A2](#) demonstrate that while the number of observations increases over time, the make-up of our sample, at least with regards to crops, does not change much over time.

Despite the fair amount of consistency in regards to the make-up of crops in our sample over time, we do see variation in crop yields. In Figure [A3](#) we draw kernel densities for all crops by year and each individual crop by year. Here we see that 2009 tended to be the best year in terms of yields. Yields tended to be lower (and more variable) in 2011. Obviously, a number of factors effect yields, and our econometric analysis is designed to identify the impact of CA adoption, rainfall, and their joint impacts, while controlling for inputs and household characteristics. What Figure [A3](#) makes clear, though, is that yields in 2011 were lower than in previous years and since there are more observations from 2011 what occurred in that year will play a large roll in determining our results.

Beyond understanding when our data comes from, it is important to know what sort of weather occurred in each year and if that weather was in line with historical weather trends. An initial concern with using rainfall data to predict crop yield is that one has chosen the correct window within which to measure rainfall. As discussed in the paper, rainfall in Zimbabwe is often of high intensity but low duration and frequency, meaning that monthly averages may be more appropriate than our seasonal averages used in the paper. To test for this, we examine the impact of both seasonal and monthly deviations in rainfall on crop output (see Table [A1](#)). We find little empirical evidence that late onset of rainfall or mid-season dry spells reduce yields. Rainfall in the first month

of the growing season only impacts maize and groundnut yields while deviations in the mid-season months tend to only affect maize, groundnut, and cowpea. By contrast, for all five crops, deviations in seasonal rainfall have a significant and negative impact on crop yields. Based on these results we conclude that our use of seasonal deviations is a strong proxy for rainfall induced stresses to agricultural production.

Having determined that our seasonal measure of rainfall is a strong predictor of crop yield, a second concern is that the data contain sufficient variation in realized rainfall over the four year study period. To determine this we draw the distribution of cumulative seasonal rainfall (see Figure A4). The distribution forms a fairly broad band, with half the observations (49 percent) falling outside half a standard deviation of the mean and nearly a third of observations (29 percent) falling outside one standard deviations of the mean.

Incidentally, we also explored using satellite temperature data from Sheffield et al. (2014) in our analysis. However, unlike the rainfall data there was insufficient variation in the temperature data. We calculated deviations in average seasonal temperature from long-run averages along with growing degree days (GDDs). For both measures there was almost no variation over space or over time. In the areas of Zimbabwe where the data come from, temperature throughout the growing season did not fluctuate greatly during our time frame, making it difficult to identify the impact of temperature on crops without a longer panel.

A final concern is that because we have only four years of data, one outlier year might play an out-sized role in our results. To determine how rainfall differed across the four years, Table A2 presents summary statistics of the rainfall shocks by year. 2008 and 2010 were wet years relative to the 15 year average while 2011 was a dry year. 2009 was variable, with some regions experiencing surplus rainfall while others experienced rainfall shortage. From Figure 2 in the paper, it is clear that our data come from a period of above average rainfall, though the final year of 2011 was drier than average. 2011 is the sixth driest season in the past 15 years while 2008 was the fifth rainiest season in the past 15 years. Based on the evidence in Table A2 and Figure 2 we conclude that the rainfall draws from the four years of household data are typical of historic rainfall patterns in Zimbabwe, though on the slightly rainier side.

B Full Results from Yield Function

This appendix contains full specifications and zero-stage probit results for the models presented in an abbreviated form in the paper. In all specifications we include controls for crop-specific inputs. These inputs include logs of basal fertilizer application, of top application of fertilizer, of seed, and of land area. Since our primary focus is the impact of CA and rainfall shocks on yields, we treat these variables as exogenous controls and refrained from reporting the relevant point estimates in the paper. For those interested in the full specifications, the following tables will be of interest.

C Additional Robustness Checks

In this appendix we verify the robustness of our main results to changes in our underlying specification.

A first concern is that our results rely on the functional form of our zero-stage regression. To check this, we replace our zero-stage probit using the Mundlak-Chamberlain device with a zero-stage Tobit using fixed effects. Results are consistent regardless of how we specify the zero-stage

selection decision (see Tables C1 and C2).

A second concern is that our indicator for CA is capturing something else and thereby creating spurious correlation between CA and yields. We can check to see if this is the case by constructing placebo tests or pseudo-treatments.¹⁹ To do this, we divide the sample using a similar logic behind our sample split in Table 4 from the paper. We begin by constructing two different comparison groups. The first defines a pseudo-treatment in which “treated” plots are those that will, sometime in the future (2009, 10, or 11), have CA but at the moment are cultivated using conventional methods. As an example, if a plot was cultivated using traditional methods in 2008 and 2009 but was a CA plot in 2010 and 2011, we mark the plot as “treated” and keep only the observations from 2008 and 2009. The “control” plots are those that never have CA and are always cultivated using conventional methods.

The second comparison group defines a pseudo-treatment in which “treated” plots are those that were, sometime in the past (2008, 09, or 10), CA plots but at the moment are cultivated using conventional methods. As an example, if a plot was a CA plot in 2008 but was cultivated using traditional methods in 2009-11, we mark the plot as “treated” and keep only the observations from 2009-11. These are compared to “control” plots which never have CA and are always cultivated using conventional methods.

By constructing our pseudo-treatments in this way, we are comparing plots that are always conventionally cultivated with plots that are currently being conventionally cultivated but where in the past, or will be in the future, CA plots. We expect that CA will have no significant effect on yields, nor will the CA-rainfall interaction have a significant effect. We do, however, expect the rainfall shock to reduce yields.

Results from these regressions are presented in Table C3. Column (1) presents results from our first comparison group, where “treatment” is being a future CA plot and “control” is always being a conventional plot. CA and the CA-rainfall interaction has no impact on yields for maize, millet, groundnut, and cowpea. Only in the case of sorghum does our placebo test generate significant results. Column (2) presents results from our second comparison group, where “treatment” is being a CA plot in the past and “control” is always being a conventional plot. Here CA and CA-rainfall are never significant. Across the two regressions, the rainfall shock term reduces yields in 7 of 10 cases.

Based on the regression results in Tables C2 and C3 we conclude that our main results are robust to changes in the underlying specification and that our estimates of the impact of CA are casually identified.

D Alternative Rainfall Measures

In this appendix we test the robustness of our main conclusions in the paper to changes in our rainfall data.

First, one might be concerned that, because we have only four years of data (i.e., four rainfall draws), one outlier year might be driving the results.²⁰ To see if this is the case, we estimate the yield function for all years, dropping one year at a time. Table D1 reports results when we treat CA as endogenous and use our single rainfall shock variable. For purposes of comparison, column (1) reproduces results from Table 7 in the paper. Column (2) removes data from 2011, column

¹⁹We thank an anonymous reviewer for suggesting this approach.

²⁰We thank an anonymous reviewer for bringing this concern to our attention.

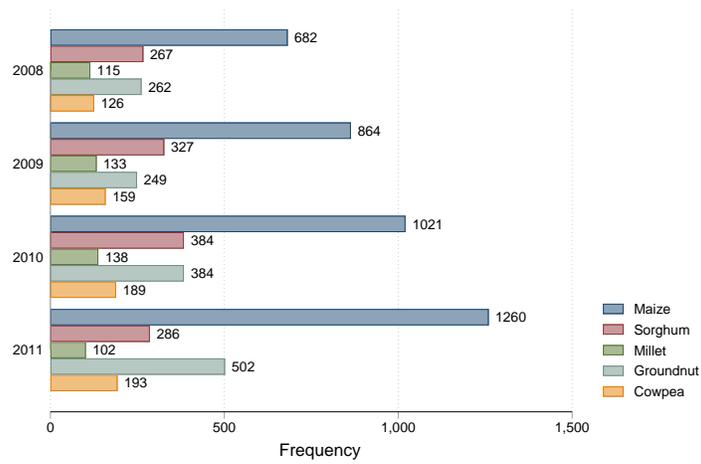
(3) removes data from 2010, column (4) removes data from 2009, and column (5) removes data from 2008. Because we have an unbalanced panel, the number of observations vary between 5,300 and 7,643. Despite this, the signs and significance levels on coefficients in columns (3) - (5) are very similar to those in column (1), which uses all four years. Additionally, the coefficients from sorghum, millet, groundnut, and cowpea are very similar across all columns. Where we do see differences is for maize in column (2), where we drop the relatively dry year of 2011.

Using our preferred specification from the paper (which includes both shortages and surplus) we estimate the yield function using all three year combinations. The results are presented in Table D2. Again, the only place where we see differences in estimates are for maize in column (2), particularly for CA, which now has a positive impact on maize yields, and CA interacted with rainfall shortage, which now reduces maize yields.

It is clear that our results for maize are driven by rainfall shortages in the 2011 year. However, we are okay with this fact because 2011 is the only year where we see significant rainfall shortage in maize growing regions. Over 60 percent of our observations for maize during rainfall shortage come from 2011. Estimating the yield function without data from 2011 amounts to asking our data to do out of sample predictions for maize during dry spells using a sample skewed towards wet years.

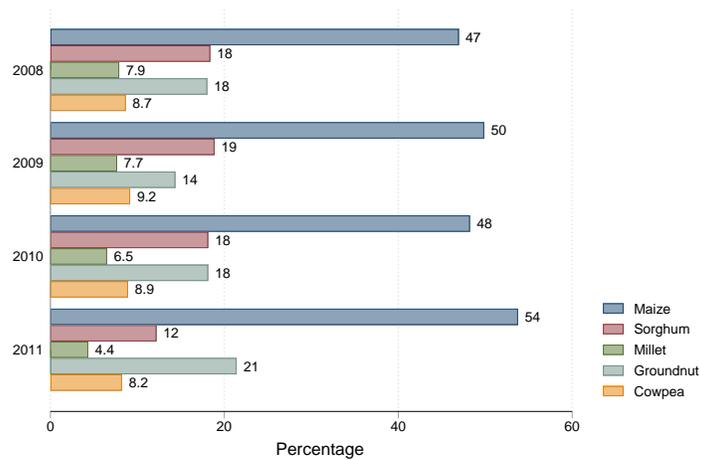
A second concern might be that our results rely on our chosen measure of rainfall shock. To test this, we develop two new, non-linear measures of rainfall shocks. In the first we replace any deviation in rainfall that is within \pm one standard deviation of the mean (0.476) with a zero. Thus, any realized value that is $0.202 \leq R_{wt} \leq 0.749$ is set to zero. In the second we replace any deviation in rainfall that is within \pm one half of a standard deviation of the mean (0.476) with a zero. Thus, any realized value that is $0.339 \leq R_{wt} \leq 0.612$ is set to zero. These changes to our rainfall shock term do not have a material effect on our results (see Table D3)

Figure A1: Plot Observations by Crop and Year



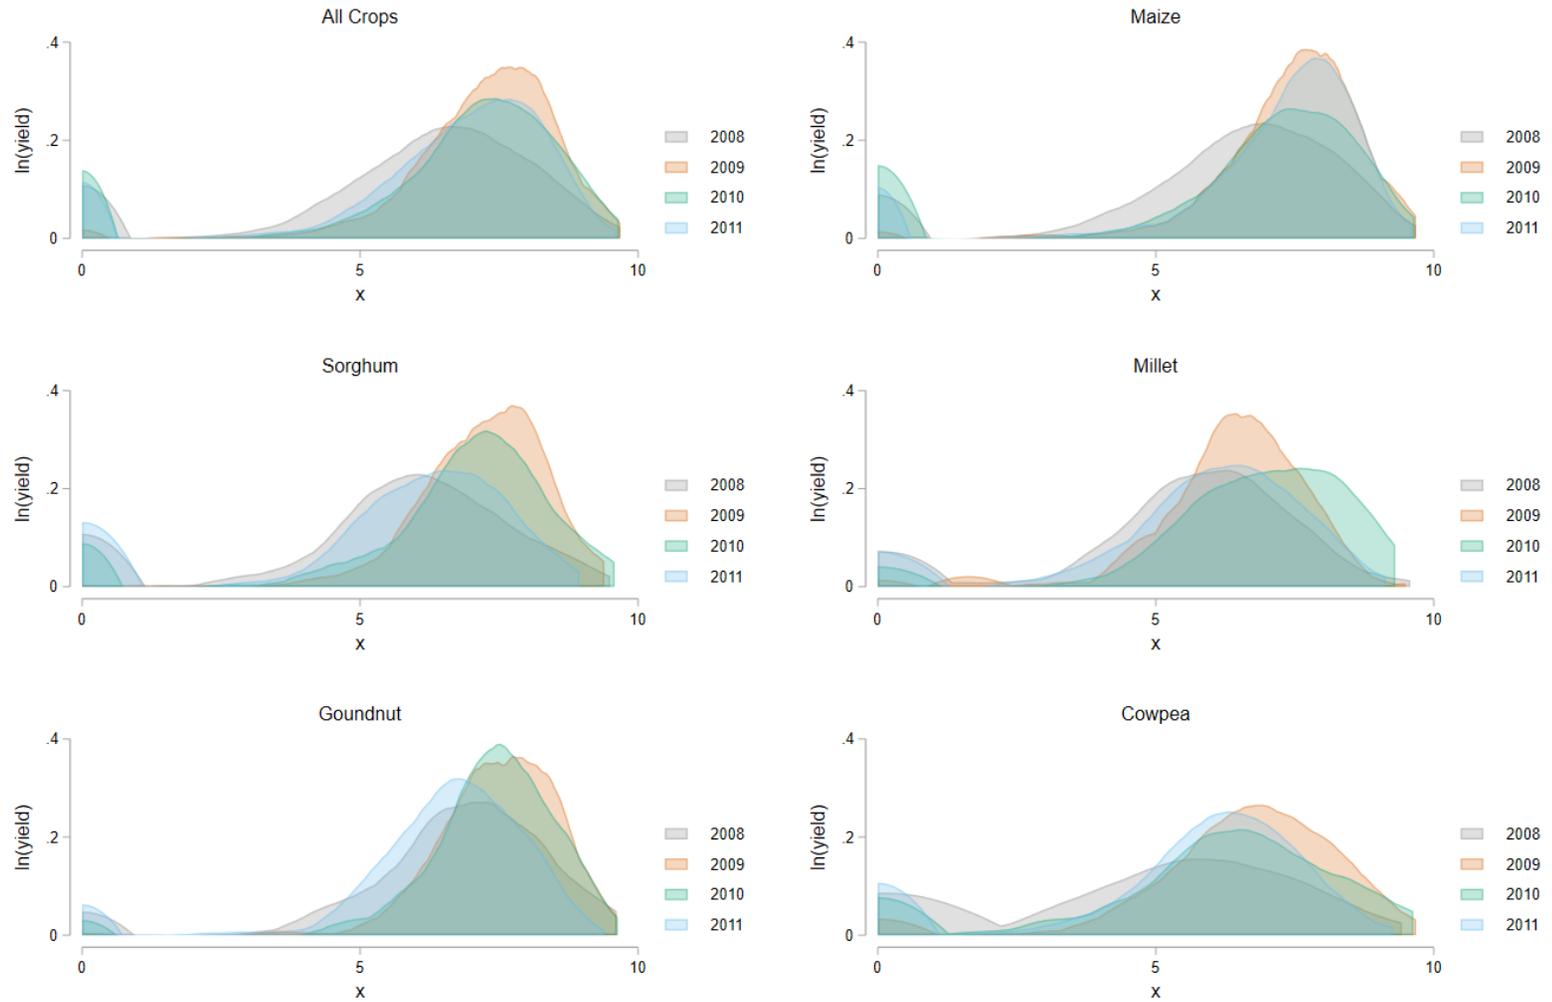
Note: Figure displays the frequency of plot-level observations for each crop in each year.

Figure A2: Percentage of Crops in Each Year



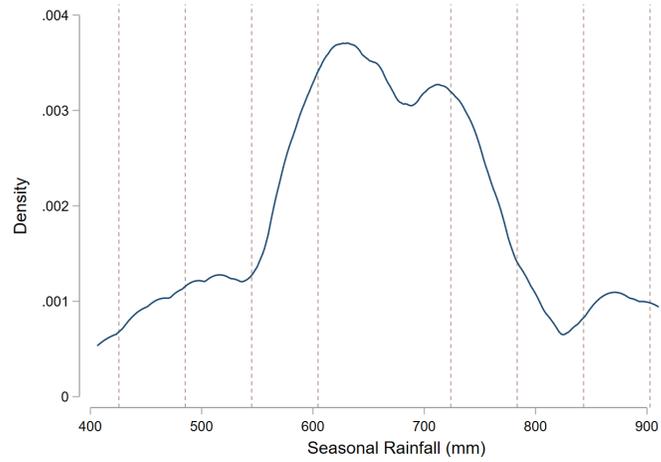
Note: Figure displays the percentage of plot-level observations growing each crop in each year.

Figure A3: Yield Distributions by Crop and Year



Note: Figure displays kernel densities of yields for each crop in each year.

Figure A4: Distribution of Seasonal Rainfall by Ward



Note: Figure displays the kernel density of cumulative seasonal rainfall from each ward for the four year study period (2008-2011). Vertical lines are drawn at \pm half, one, one half, and two standard deviations from the mean, which is 664mm. Data comes from the CHIRPS database (Funk et al., 2015).

Table A1: Yield Function with Rainfall Deviations

	Maize (1)	Sorghum (2)	Millet (3)	Groundnut (4)	Cowpea (5)
<i>Panel A: Deviation in Seasonal Rainfall</i>					
Seasonal deviation	-0.631*** (0.192)	-1.125*** (0.269)	-1.389*** (0.312)	-0.407* (0.212)	-0.913** (0.399)
<i>Panel B: Deviation in Monthly Rainfall</i>					
Oct deviation	0.324* (0.196)	-0.262 (0.212)	0.023 (0.302)	0.534*** (0.199)	0.129 (0.314)
Nov deviation	-0.213* (0.111)	-0.056 (0.176)	-0.076 (0.194)	0.154 (0.127)	-0.028 (0.238)
Dec deviation	-0.199* (0.114)	-0.122 (0.222)	-0.535* (0.309)	0.226* (0.136)	0.338 (0.265)
Jan deviation	-0.571*** (0.167)	0.119 (0.176)	-0.267 (0.261)	0.251 (0.153)	0.418* (0.235)
Feb deviation	-1.054*** (0.213)	-0.320 (0.274)	-0.254 (0.358)	-0.683*** (0.259)	-2.011*** (0.443)
Mar deviation	-1.001*** (0.187)	-0.701** (0.322)	-0.587 (0.506)	-0.764*** (0.242)	0.112 (0.386)
Apr deviation	-0.295*** (0.092)	0.104 (0.108)	-0.372*** (0.140)	-0.154 (0.105)	-0.015 (0.155)

Note: Coefficients are presented in columns based on crop type in order to minimize space. Panel A reports results from a single regression with log yield as the dependent variable and deviations in cumulative seasonal rainfall as the variable of interest (Observations = 7,643; $R^2 = 0.923$). Panel B reports results from a single regression with log yield as the dependent variable and deviations in cumulative monthly rainfall as the variables of interest (Observations = 7,643; $R^2 = 0.926$). Though not reported, both specifications include crop-specific CA adoption terms, crop-specific inputs, crop-specific intercept terms, year dummies, and household fixed effects. Standard errors clustered by household and crop are reported in parentheses (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

Table A2: Rainfall Shocks by Year

	2008	2009	2010	2011
Rainfall shock	0.682 (0.289)	0.487 (0.267)	0.351 (0.225)	0.453 (0.231)
Rainfall shortage	0.000 (0.000)	0.282 (0.257)	0.028 (0.065)	0.442 (0.246)
Rainfall surplus	0.682 (0.289)	0.205 (0.347)	0.323 (0.253)	0.011 (0.042)
number of observations	1,452	1,732	2,116	2,343

Note: Columns in the table display means of the data by year with standard deviations in parenthesis.

Table B1: Yield Function with CA as Exogenous

	(1)	(2)	(3)	(4)
<i>Maize</i>				
CA (= 1)	0.631*** (0.081)	0.573*** (0.081)	0.222 (0.139)	0.207 (0.137)
rainfall shock			-0.667*** (0.215)	-0.904*** (0.227)
CA × rainfall shock			0.872*** (0.247)	0.744*** (0.237)
ln(basal)	0.118*** (0.025)	0.079*** (0.024)	0.120*** (0.025)	0.085*** (0.023)
ln(top)	0.321*** (0.028)	0.270*** (0.026)	0.312*** (0.028)	0.261*** (0.025)
ln(seed)	0.250*** (0.075)	0.248*** (0.074)	0.236*** (0.075)	0.233*** (0.073)
ln(area)	-0.449*** (0.061)	-0.496*** (0.064)	-0.442*** (0.061)	-0.483*** (0.063)
<i>Sorghum</i>				
CA (= 1)	0.041 (0.180)	-0.051 (0.190)	-0.756*** (0.271)	-0.596** (0.288)
rainfall shock			-1.285*** (0.303)	-1.432*** (0.317)
CA × rainfall shock			1.652*** (0.458)	1.130** (0.499)
ln(basal)	0.038 (0.089)	0.015 (0.070)	0.037 (0.091)	0.020 (0.070)
ln(top)	0.107 (0.066)	0.130** (0.060)	0.106 (0.065)	0.124** (0.059)
ln(seed)	0.395*** (0.112)	0.428*** (0.119)	0.396*** (0.111)	0.426*** (0.118)
ln(area)	-0.527*** (0.085)	-0.560*** (0.098)	-0.500*** (0.085)	-0.530*** (0.098)
<i>Millet</i>				
CA (= 1)	-0.145 (0.359)	0.118 (0.374)	-0.994 (0.709)	-0.711 (0.736)
rainfall shock			-1.094*** (0.306)	-1.521*** (0.319)
CA × rainfall shock			1.492 (0.940)	1.475 (1.171)
ln(basal)	0.230** (0.105)	0.180 (0.127)	0.191* (0.102)	0.136 (0.127)
ln(top)	-0.233* (0.123)	-0.126 (0.130)	-0.216* (0.122)	-0.109 (0.125)
ln(seed)	0.221 (0.156)	0.264* (0.148)	0.258* (0.153)	0.298** (0.144)
ln(area)	-0.451***	-0.584***	-0.465***	-0.601***

Continued on Next Page...

Table B1 – Continued

	(1)	(2)	(3)	(4)
<i>Groundnut</i>	(0.135)	(0.129)	(0.134)	(0.127)
CA (= 1)	0.299** (0.142)	0.323** (0.144)	-0.326 (0.287)	0.130 (0.279)
rainfall shock			-0.403* (0.208)	-0.489** (0.228)
CA × rainfall shock			1.106** (0.493)	0.331 (0.456)
ln(basal)	0.075 (0.076)	-0.007 (0.067)	0.081 (0.075)	-0.008 (0.067)
ln(top)	-0.085 (0.076)	-0.039 (0.065)	-0.052 (0.075)	-0.030 (0.066)
ln(seed)	0.560*** (0.076)	0.470*** (0.076)	0.558*** (0.076)	0.470*** (0.076)
ln(area)	-0.490*** (0.077)	-0.614*** (0.074)	-0.487*** (0.077)	-0.608*** (0.074)
<i>Cowpea</i>				
CA (= 1)	-0.035 (0.241)	0.194 (0.250)	-0.844* (0.466)	-0.362 (0.460)
rainfall shock			-1.053** (0.452)	-1.216*** (0.448)
CA × rainfall shock			1.688* (0.905)	1.114 (0.878)
ln(basal)	0.306*** (0.112)	0.106 (0.095)	0.298*** (0.112)	0.108 (0.095)
ln(top)	-0.134 (0.112)	-0.085 (0.099)	-0.127 (0.115)	-0.082 (0.102)
ln(seed)	0.440*** (0.164)	0.385** (0.163)	0.459*** (0.165)	0.395** (0.163)
ln(area)	-0.541*** (0.108)	-0.652*** (0.114)	-0.550*** (0.109)	-0.660*** (0.115)
Household FE	No	Yes	No	Yes
Observations	7,643	7,643	7,643	7,643
R ²	0.899	0.922	0.900	0.923

Note: Dependent variable is log of yield. All specifications include crop-specific intercept terms and year dummies. Column (1) excludes rainfall variables as well as household fixed effects. Column (2) excludes rainfall variables but includes household fixed effects. Column (3) includes the rainfall shock and its interaction with CA but excludes household fixed effects. Column (4) includes both the rainfall shock, its interaction with CA, and household fixed effects. Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table B2: Yield Function with CA as Endogenous

	(1)	(2)	(3)	(4)
<i>Maize</i>				
CA (= 1)	13.404 (10.916)	-1.939* (1.122)	15.697 (18.025)	-2.853** (1.141)
rainfall shock			0.569 (2.198)	-1.675*** (0.353)
CA × rainfall shock			1.335 (2.896)	2.628*** (0.704)
ln(basal)	-0.632 (0.639)	0.219*** (0.069)	-0.789 (0.993)	0.219*** (0.063)
ln(top)	-0.283 (0.527)	0.404*** (0.065)	-0.433 (0.795)	0.360*** (0.057)
ln(seed)	1.819 (1.347)	-0.052 (0.164)	2.189 (2.105)	-0.017 (0.150)
ln(area)	-0.006 (0.429)	-0.608*** (0.090)	0.101 (0.630)	-0.568*** (0.082)
<i>Sorghum</i>				
CA (= 1)	19.230 (13.886)	-0.275 (1.498)	24.937 (24.052)	-0.171 (1.603)
rainfall shock			-0.940 (0.823)	-1.476*** (0.353)
CA × rainfall shock			-2.555 (4.717)	0.899 (0.904)
ln(basal)	-0.840 (0.678)	0.043 (0.106)	-0.978 (1.006)	0.023 (0.096)
ln(top)	-3.028 (2.260)	0.150 (0.236)	-3.788 (3.640)	0.060 (0.220)
ln(seed)	0.246 (0.317)	0.395*** (0.113)	0.236 (0.387)	0.400*** (0.113)
ln(area)	1.480 (1.453)	-0.551*** (0.180)	1.936 (2.290)	-0.467*** (0.166)
<i>Millet</i>				
CA (= 1)	32.388 (30.713)	-4.551 (3.402)	49.502 (64.115)	-3.784 (3.808)
rainfall shock			0.654 (2.163)	-1.715*** (0.341)
CA × rainfall shock			-12.938 (30.199)	3.357 (2.279)
ln(basal)	-1.721 (2.080)	0.453* (0.237)	-2.112 (2.943)	0.228 (0.196)
ln(top)	-4.778 (4.456)	0.547 (0.530)	-6.297 (7.398)	0.209 (0.482)
ln(seed)	-1.215 (1.547)	0.424* (0.225)	-1.737 (2.612)	0.371* (0.198)
ln(area)	3.501 (3.877)	-1.163*** (0.437)	4.771 (6.379)	-0.858** (0.402)
<i>Groundnut</i>				
CA (= 1)	14.997 (10.282)	0.050 (1.199)	16.999 (16.357)	-0.272 (1.324)
rainfall shock			-1.659* (0.886)	-0.950*** (0.265)
CA × rainfall shock			2.423 (2.895)	1.433* (0.832)
ln(basal)	-1.137 (0.893)	0.001 (0.114)	-1.370 (1.369)	-0.036 (0.106)
ln(top)	-1.966 (1.336)	-0.008 (0.159)	-2.362 (2.133)	-0.031 (0.158)
ln(seed)	1.090*** (0.398)	0.566*** (0.084)	1.197** (0.591)	0.561*** (0.086)
ln(area)	0.405 (0.676)	-0.654*** (0.103)	0.603 (1.040)	-0.610*** (0.096)
<i>Cowpea</i>				
CA (= 1)	10.530 (8.312)	-0.938 (1.147)	12.451 (14.121)	-1.564 (1.349)
rainfall shock			-0.061 (1.643)	-1.606*** (0.485)
CA × rainfall shock			0.775 (3.789)	2.123 (1.404)
ln(basal)	0.049 (0.296)	0.146 (0.099)	-0.014 (0.399)	0.134 (0.096)
ln(top)	-1.999	0.100	-2.385	0.047

Continued on Next Page...

Table B2 – Continued

	(1)	(2)	(3)	(4)
ln(seed)	(1.477) -0.231	(0.205) 0.433***	(2.297) -0.362	(0.197) 0.446***
ln(area)	(0.592) 0.517	(0.167) -0.776***	(0.901) 0.745	(0.168) -0.741***
	(0.852)	(0.163)	(1.319)	(0.159)
Household FE	No	Yes	No	Yes
Observations	7,643	7,643	7,643	7,643
Log Likelihood	-24,462	-16,378	-25,914	-16,192
Kleibergen-Paap LM stat	1.773	26.89***	1.028	33.86***
Anderson-Rubin Wald stat	37.38***	19.47***	75.09***	35.88***
Kleibergen-Paap Wald stat	0.355	4.593***	0.102	2.955***

Note: Dependent variable is log of yield. All specifications include crop-specific intercept terms and year dummies. In each regression the adoption of CA is treated as endogenous and is instrumented with the Inverse Mills Ratio (IMR) calculated from the predicted values of the zero-stage probits reported in Table 6. The CA \times rainfall shock term is also treated as endogenous and instrumented using the interaction of the IMR and the rainfall shock term. The null hypothesis of the Kleibergen-Paap LM test is that the rank condition fails (i.e., the first-stage equation is underidentified). The null hypothesis of the Anderson-Rubin Wald test is that the coefficients on the endogenous regressors in the structural equation are jointly equal to zero (i.e., the instruments in the first-stage equation are weak). The null hypothesis of the Kleibergen-Paap Wald test is that a t -test at the 5% significance level on the coefficients of the endogenous regressors rejects no more than 25% of the time (i.e., the instruments in the first-stage equation are weak). Standard errors clustered by household and crop are reported in parentheses (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

Table B3: Zero-Stage Probit with Rain Shortage or Surplus

	(1)	(2)	(3)	(4)
Number of HH in ward with NGO support	0.009*** (0.001)	0.009*** (0.001)	0.009*** (0.001)	0.010*** (0.002)
Type of Effect	Household MCD	Household MCD	Household MCD	Plot MCD
Observations	7,643	7,643	7,643	5,004
Log Likelihood	-3,107	-3,102	-3,096	-2,102

Note: Dependent variable is an indicator for whether or not CA was used on the plot. Though not reported, all probit regressions include crop-specific inputs and intercept terms, and year dummies. Column (1) excludes the rainfall variable as well as the Mundlak-Chamberlain device (MCD). Column (2) excludes the rainfall variable but includes the MCD. Column (3) includes the rainfall variable but excludes the MCD. Column (4) includes both the rainfall variable and the MCD. Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table B4: Yield Function with Rain Shortage or Surplus

	(1)	(2)	(3)	(4)
<i>Maize</i>				
CA (= 1)	-0.466 (1.379)	-2.657* (1.378)	-1.893 (1.472)	3.393 (3.156)
rainfall shortage	0.089 (0.400)		-1.076** (0.481)	-1.411** (0.695)
CA × rainfall shortage	0.147 (0.744)		2.098** (0.987)	3.522** (1.597)
rainfall surplus		-1.619*** (0.304)	-1.884*** (0.358)	-0.444 (0.637)
CA × rainfall surplus		2.012*** (0.513)	2.931*** (0.657)	1.953** (0.845)
ln(basal)	0.131 (0.080)	0.248*** (0.084)	0.166** (0.081)	-0.115 (0.172)
ln(top)	0.319*** (0.074)	0.397*** (0.078)	0.305*** (0.073)	0.138 (0.108)
ln(seed)	0.122 (0.187)	-0.053 (0.193)	0.100 (0.185)	0.699** (0.342)
ln(area)	-0.532*** (0.091)	-0.595*** (0.096)	-0.518*** (0.089)	-0.462*** (0.150)
<i>Sorghum</i>				
CA (= 1)	-0.262 (1.464)	-1.211 (2.042)	-0.512 (1.780)	5.808 (4.791)
rainfall shortage	-1.079*** (0.389)		-1.702*** (0.454)	-2.235*** (0.769)
CA × rainfall shortage	2.437* (1.263)		2.356* (1.326)	3.654* (1.977)
rainfall surplus		-0.550 (0.342)	-1.214*** (0.377)	-0.531 (0.533)
CA × rainfall surplus		-0.369 (1.008)	0.464 (0.987)	-0.144 (1.507)
ln(basal)	0.018 (0.107)	0.114 (0.108)	0.037 (0.101)	-0.190 (0.264)
ln(top)	0.079 (0.249)	0.311 (0.289)	0.093 (0.257)	-0.671 (0.553)
ln(seed)	0.396*** (0.113)	0.407*** (0.112)	0.405*** (0.112)	0.508*** (0.194)
ln(area)	-0.518*** (0.182)	-0.663*** (0.195)	-0.510*** (0.179)	-0.008 (0.420)
<i>Millet</i>				
CA (= 1)	-2.913 (2.890)	-3.914 (3.404)	-2.671 (3.293)	12.779 (12.606)
rainfall shortage	-0.508 (0.610)		-1.604** (0.658)	-1.438 (0.927)
CA × rainfall shortage	-0.136 (1.576)		1.930 (2.344)	-4.318 (8.524)
rainfall surplus		-1.287*** (0.356)	-1.689*** (0.320)	-0.427 (0.584)
CA × rainfall surplus		2.726* (1.548)	3.170 (2.202)	-0.775 (7.605)
ln(basal)	0.342* (0.207)	0.275 (0.198)	0.161 (0.177)	-0.524 (0.671)
ln(top)	0.312 (0.455)	0.364 (0.483)	0.098 (0.424)	-1.148 (1.090)
ln(seed)	0.358* (0.208)	0.380* (0.198)	0.325* (0.186)	0.022 (0.360)
ln(area)	-0.957** (0.383)	-0.935** (0.391)	-0.738** (0.353)	0.809 (1.229)
<i>Groundnut</i>				
CA (= 1)	0.359 (1.159)	-1.519 (1.550)	-0.266 (1.499)	4.764 (3.664)
rainfall shortage	-1.125*** (0.318)		-1.229*** (0.338)	-2.298** (0.934)
CA × rainfall shortage	0.240 (0.772)		0.405 (0.984)	1.499 (1.570)
rainfall surplus		0.036 (0.297)	-0.170 (0.281)	0.506 (0.465)
CA × rainfall surplus		0.056 (0.899)	-0.205 (0.985)	-0.607 (1.422)
ln(basal)	-0.012 (0.110)	0.119 (0.128)	0.027 (0.110)	-0.256 (0.291)
ln(top)	-0.072 (0.158)	0.191 (0.182)	0.014 (0.172)	-0.551 (0.435)

Continued on Next Page...

Table B4 – Continued

	(1)	(2)	(3)	(4)
ln(seed)	0.563*** (0.083)	0.516*** (0.092)	0.543*** (0.086)	0.891*** (0.252)
ln(area)	-0.646*** (0.100)	-0.751*** (0.101)	-0.689*** (0.093)	-0.600*** (0.161)
<i>Cowpea</i>				
CA (= 1)	-0.742 (1.141)	-0.929 (1.415)	-1.188 (1.439)	1.829 (3.818)
rainfall shortage	-0.628 (0.477)		-1.493*** (0.553)	-2.468** (1.221)
CA × rainfall shortage	3.171*** (1.008)		3.763*** (1.272)	6.006** (2.558)
rainfall surplus		-1.103** (0.482)	-1.492*** (0.557)	-0.970 (1.002)
CA × rainfall surplus		-0.286 (1.309)	1.234 (1.624)	3.148 (2.543)
ln(basal)	0.073 (0.098)	0.168 (0.104)	0.090 (0.098)	-0.162 (0.276)
ln(top)	0.003 (0.205)	0.092 (0.218)	-0.008 (0.207)	-0.650 (0.569)
ln(seed)	0.451*** (0.168)	0.454*** (0.171)	0.463*** (0.170)	0.659** (0.287)
ln(area)	-0.720*** (0.162)	-0.799*** (0.169)	-0.723*** (0.162)	-0.491* (0.291)
Type of Effect	Household FE	Household FE	Household FE	Plot MCD
Observations	7,643	7,643	7,643	5,004
Log Likelihood	-15,908	-16,406	-15,902	-12,268
Kleibergen-Paap LM stat	22.96***	23.02***	25.23***	3.790**
Anderson-Rubin Wald stat	22.10**	29.68***	44.79***	38.53***
Kleibergen-Paap Wald stat	1.930*	1.989*	1.444	0.251

Note: Dependent variable is log of yield. All specifications include crop-specific intercept terms and year dummies. In each regression the adoption of CA is treated as endogenous and is instrumented with the Inverse Mills Ratio (IMR) calculated from the predicted values of zero-stage probits which are presented in the Appendix, Table B3. The CA × rainfall shortage and CA × rainfall surplus terms are also treated as endogenous and instrumented using the interaction of the IMR and the rainfall terms. The null hypothesis of the Kleibergen-Paap LM test is that the rank condition fails (i.e., the first-stage equation is underidentified). The null hypothesis of the Anderson-Rubin Wald test is that the coefficients on the endogenous regressors in the structural equation are jointly equal to zero (i.e., the instruments in the first-stage equation are weak). The null hypothesis of the Kleibergen-Paap Wald test is that a *t*-test at the 5% significance level on the coefficients of the endogenous regressors rejects no more than 25% of the time (i.e., the instruments in the first-stage equation are weak). Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table C1: Zero-Stage Tobit with Rain Shortage and Surplus

	(1)	(2)
Number of HH in ward with NGO support	0.005** (0.002)	0.007*** (0.001)
Household Control	FE	MCD
Observations	7,643	7,643
Log Likelihood	-4,310	-4,870

Note: Dependent variable is an indicator for whether or not CA was used on the plot. Though not reported, all Tobit regressions include crop-specific inputs and intercept terms, and year dummies. Column (1) includes household fixed effects. Column (2) includes the Mundlak-Chamberlain device. Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table C2: Yield Function using Zero-Stage Tobit

	(1)	(2)
<i>Maize</i>		
CA (= 1)	-0.141 (0.448)	-0.611 (0.736)
rainfall shortage	-0.366 (0.365)	-0.639 (0.440)
CA × rainfall shortage	0.216 (0.684)	0.958 (0.927)
rainfall surplus	-1.329*** (0.313)	-1.659*** (0.321)
CA × rainfall surplus	1.544*** (0.533)	2.450*** (0.613)
<i>Sorghum</i>		
CA (= 1)	-0.185 (0.763)	-0.099 (1.002)
rainfall shortage	-1.770*** (0.417)	-1.929*** (0.434)
CA × rainfall shortage	2.460** (1.105)	3.335*** (1.243)
rainfall surplus	-1.207*** (0.365)	-1.410*** (0.364)
CA × rainfall surplus	0.579 (0.787)	1.141 (0.838)
<i>Millet</i>		
CA (= 1)	-1.942 (1.749)	-1.352 (1.767)
rainfall shortage	-1.859*** (0.629)	-1.866*** (0.637)
CA × rainfall shortage	3.248 (2.573)	3.356 (2.629)
rainfall surplus	-1.595*** (0.308)	-1.699*** (0.299)
CA × rainfall surplus	2.778 (1.819)	3.663* (2.030)
<i>Groundnut</i>		
CA (= 1)	0.177 (0.535)	0.333 (0.627)
rainfall shortage	-1.258*** (0.318)	-1.368*** (0.321)
CA × rainfall shortage	0.555 (0.731)	1.049 (0.808)
rainfall surplus	-0.061 (0.248)	-0.088 (0.246)
CA × rainfall surplus	-0.321 (0.601)	-0.339 (0.673)
<i>Coupea</i>		
CA (= 1)	-0.800 (0.734)	-0.606 (0.992)
rainfall shortage	-1.430*** (0.533)	-1.500*** (0.550)
CA × rainfall shortage	3.558*** (1.027)	3.802*** (1.189)
rainfall surplus	-1.572*** (0.536)	-1.440*** (0.537)
CA × rainfall surplus	1.733 (1.421)	1.234 (1.604)
Household FE	Yes	Yes
Observations	7,643	7,643
Kleibergen-Paap LM Stat	279.51***	182.71***
Log Likelihood	-15,737	-15,770

Note: Dependent variable is log of yield. Though not reported, all specifications include crop-specific inputs and intercept terms, and year dummies. Column (1) uses the predicted value and the Inverse Mills Ratio from the zero-stage Tobit with household fixed effects as presented in column (1) of Table C1. Column (2) uses the predicted value and the Inverse Mills Ratio from the zero-stage Tobit with household a Mundlak-Chamberlain device as presented in column (2) of Table C1. All standard errors are clustered by household and crop and are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table C3: Placebo Tests

	Future Adopter v. Conventional (1)	Future Disadopter v. Conventional (2)
<i>Maize</i>		
CA (= 1)	0.326 (0.606)	0.093 (0.293)
rainfall shock	-0.333 (0.290)	-0.492* (0.283)
CA × rainfall shock	-0.924 (0.921)	-0.611 (0.648)
<i>Sorghum</i>		
CA (= 1)	1.329** (0.530)	0.222 (0.464)
rainfall shock	-1.063*** (0.390)	-1.292*** (0.387)
CA × rainfall shock	-1.419* (0.809)	0.391 (0.865)
<i>Millet</i>		
CA (= 1)	-0.928 (0.772)	0.063 (0.382)
rainfall shock	-1.188*** (0.415)	-1.379*** (0.405)
CA × rainfall shock	0.555 (1.270)	0.000 (0.588)
<i>Groundnut</i>		
CA (= 1)	0.048 (0.710)	0.370 (0.257)
rainfall shock	-0.200 (0.287)	-0.206 (0.284)
CA × rainfall shock	-0.423 (1.125)	-0.225 (0.589)
<i>Cowpea</i>		
CA (= 1)	0.052 (1.094)	0.264 (0.849)
rainfall shock	-1.116** (0.553)	-1.274** (0.543)
CA × rainfall shock	0.296 (1.432)	0.106 (1.862)
Observations	4,132	4,629
R^2	0.928	0.924

Note: Dependent variable is log of yield. Though not reported, all specifications include crop-specific inputs and intercept terms, and year dummies. Column (1) uses a pseduo-treatment in which “treated” plots are those that will, sometime in the future (2009, 10, or 11), have CA but at the moment are cultivated using conventional methods. These are compared to “control” plots which never have CA and are always cultivated using conventional methods. Column (2) uses a pseduo-treatment in which “treated” plots are those that were, sometime in the past (2008, 09, or 10), CA plots but at the moment are cultivated using conventional methods. These are compared to “control” plots which never have CA and are always cultivated using conventional methods. All standard errors are clustered by household and crop and are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).

Table D1: Yield Function Using Three out of Four Years

	All Years (1)	W/O 2011 (2)	W/O 2010 (3)	W/O 2009 (4)	W/O 2008 (5)
<i>Maize</i>					
CA (= 1)	-2.853** (1.141)	4.907 (3.036)	-3.570*** (1.195)	-2.764** (1.169)	-5.010*** (1.242)
rainfall shock	-1.675*** (0.353)	-0.382 (0.691)	-1.493*** (0.376)	-1.645*** (0.395)	-2.041*** (0.456)
CA × rainfall shock	2.628*** (0.704)	-0.255 (1.271)	2.679*** (0.670)	2.722*** (0.809)	3.316*** (1.082)
<i>Sorghum</i>					
CA (= 1)	-0.171 (1.603)	-0.592 (2.709)	-0.571 (1.854)	0.134 (1.513)	-0.543 (1.914)
rainfall shock	-1.476*** (0.353)	-1.118** (0.466)	-1.082** (0.440)	-1.555*** (0.505)	-1.323*** (0.386)
CA × rainfall shock	0.899 (0.904)	0.800 (1.284)	0.597 (0.932)	0.937 (1.293)	0.508 (1.092)
<i>Millet</i>					
CA (= 1)	-3.784 (3.808)	-3.881 (3.431)	-2.160 (2.638)	0.308 (4.022)	-10.450 (9.108)
rainfall shock	-1.715*** (0.341)	-1.605*** (0.442)	-0.102 (0.459)	-1.418*** (0.448)	-2.234*** (0.398)
CA × rainfall shock	3.357 (2.279)	3.810* (2.256)	-4.656 (5.458)	3.092 (3.555)	5.284 (6.466)
<i>Groundnut</i>					
CA (= 1)	-0.272 (1.324)	-0.947 (1.526)	-1.726 (1.494)	2.370 (1.603)	-3.185** (1.597)
rainfall shock	-0.950*** (0.265)	-0.864*** (0.287)	-0.756** (0.362)	-1.095*** (0.306)	-0.865*** (0.326)
CA × rainfall shock	1.433* (0.832)	1.498** (0.759)	2.633*** (0.961)	0.245 (1.058)	1.443 (1.296)
<i>Cowpea</i>					
CA (= 1)	-1.564 (1.349)	-0.401 (1.698)	-2.818* (1.624)	-1.022 (1.482)	-2.622* (1.479)
rainfall shock	-1.606*** (0.485)	-1.438** (0.604)	-1.410** (0.639)	-1.666*** (0.547)	-1.274** (0.553)
CA × rainfall shock	2.123 (1.404)	1.690 (1.681)	2.719* (1.652)	1.783 (1.604)	1.887 (1.582)
Household FE	Yes	Yes	Yes	Yes	Yes
Observations	7,643	5,300	5,527	5,911	6,191
Kleibergen-Paap LM Stat	33.8***	11.0***	19.4***	61.7***	16.8***
Log Likelihood	-16,192	-11,684	-11,478	-12,809	-13,624

Note: For purposes of comparison, column (1) reproduces our preferred results from column (4) in Table 7 in the paper. Columns (2) - (5) use the same specification as in column (1) but in each case we remove one year of the data. Standard errors are clustered by household and crop (*p<0.1; **p<0.05; ***p<0.01).

Table D2: Yield Function Using Three out of Four Years with Rain Shortage and Surplus

	All Years (1)	W/O 2011 (2)	W/O 2010 (3)	W/O 2009 (4)	W/O 2008 (5)
<i>Maize</i>					
CA (= 1)	-1.893 (1.472)	6.615** (3.092)	-4.551** (1.876)	-0.502 (1.545)	-4.133*** (1.359)
rainfall shortage	-1.076** (0.481)	2.821** (1.357)	-2.021*** (0.481)	-0.063 (0.562)	-1.605*** (0.502)
CA × rainfall shortage	2.098** (0.987)	-6.690*** (2.459)	1.512 (1.068)	3.436*** (1.209)	3.345*** (1.153)
rainfall surplus	-1.884*** (0.358)	0.097 (0.719)	-1.066** (0.472)	-2.103*** (0.434)	-2.443*** (0.468)
CA × rainfall surplus	2.931*** (0.657)	-1.161 (1.324)	2.070*** (0.745)	3.101*** (0.754)	3.382*** (1.032)
<i>Sorghum</i>					
CA (= 1)	-0.512 (1.780)	0.196 (2.644)	-1.031 (2.707)	-1.548 (1.919)	-0.871 (1.900)
rainfall shortage	-1.702*** (0.454)	-0.783 (0.682)	-1.273** (0.512)	-1.232* (0.677)	-1.570*** (0.508)
CA × rainfall shortage	2.356* (1.326)	2.644 (1.615)	-0.379 (1.912)	-0.374 (2.418)	1.763 (1.417)
rainfall surplus	-1.214*** (0.377)	-1.042** (0.462)	-1.480*** (0.543)	-1.100* (0.564)	-0.823** (0.411)
CA × rainfall surplus	0.464 (0.987)	0.346 (1.300)	1.375 (0.969)	0.581 (1.365)	-0.300 (1.177)
<i>Millet</i>					
CA (= 1)	-2.671 (3.293)	-2.484 (3.601)	-0.528 (2.593)	2.747 (3.679)	-8.180 (8.628)
rainfall shortage	-1.604** (0.658)	-1.857 (1.220)	0.442 (0.985)	-0.624 (0.692)	-1.938*** (0.659)
CA × rainfall shortage	1.930 (2.344)	3.886 (3.689)	-8.037 (7.069)	7.960 (7.937)	6.055 (7.729)
rainfall surplus	-1.689*** (0.320)	-1.426*** (0.433)	-0.305 (0.433)	-1.670*** (0.531)	-2.146*** (0.365)
CA × rainfall surplus	3.170 (2.202)	3.131 (2.151)	-3.628 (5.522)	1.877 (4.286)	2.374 (4.797)
<i>Groundnut</i>					
CA (= 1)	-0.266 (1.499)	0.313 (1.634)	-2.366 (2.214)	2.376 (2.263)	-2.989* (1.626)
rainfall shortage	-1.229*** (0.338)	-2.880*** (0.777)	-1.211*** (0.406)	-0.471 (0.404)	-0.971** (0.385)
CA × rainfall shortage	0.405 (0.984)	3.512*** (1.338)	0.407 (1.104)	2.219 (2.781)	1.814 (1.373)
rainfall surplus	-0.170 (0.281)	-0.930*** (0.299)	-0.388 (0.571)	-0.511 (0.316)	-0.121 (0.339)
CA × rainfall surplus	-0.205 (0.985)	1.182 (0.800)	2.342* (1.384)	-1.186 (1.261)	-2.554 (1.994)
<i>Cowpea</i>					
CA (= 1)	-1.188 (1.439)	-0.171 (1.715)	-3.095 (2.233)	-0.963 (1.661)	-2.171 (1.503)
rainfall shortage	-1.493*** (0.553)	-1.195 (0.848)	-1.182 (0.756)	-0.850 (0.670)	-1.219** (0.591)
CA × rainfall shortage	3.763*** (1.272)	4.783*** (1.743)	3.580** (1.640)	1.827 (1.684)	3.214** (1.401)
rainfall surplus	-1.492*** (0.557)	-1.418** (0.613)	-1.835** (0.814)	-1.610** (0.635)	-0.697 (0.695)
CA × rainfall surplus	1.234 (1.624)	1.241 (1.744)	3.860* (2.112)	1.656 (1.825)	-1.581 (2.609)
Household FE	Yes	Yes	Yes	Yes	Yes
Observations	7,643	5,300	5,527	5,911	6,191
Kleibergen-Paap LM Stat	25.23***	11.35***	9.05***	37.3***	18.2***
Log Likelihood	-15,902	-11,967	-11,949	-12,561	-13,245

Note: For purposes of comparison, column (1) reproduces our preferred results from column (3) in Table 8 in the paper. Columns (2) - (5) use the same specification as in column (1) but in each case we remove one year of the data. Standard errors are clustered by household and crop (*p<0.1; **p<0.05; ***p<0.01).

Table D3: Yield Function with Alternative Rainfall Measures

	(1)	(2)
<i>Maize</i>		
CA (= 1)	-2.411** (1.130)	-3.587*** (1.240)
rainfall shock	-1.050*** (0.309)	-0.879*** (0.279)
CA × rainfall shock	2.095*** (0.546)	2.153*** (0.544)
<i>Sorghum</i>		
CA (= 1)	-0.538 (1.480)	-1.511 (1.628)
rainfall shock	-0.735** (0.286)	-0.753*** (0.286)
CA × rainfall shock	0.874 (0.601)	0.546 (0.605)
<i>Millet</i>		
CA (= 1)	-3.088 (2.841)	-4.552 (3.115)
rainfall shock	-0.602* (0.311)	-0.764** (0.316)
CA × rainfall shock	-1.687 (2.346)	-1.063 (1.663)
<i>Groundnut</i>		
CA (= 1)	-0.178 (1.194)	-1.283 (1.270)
rainfall shock	-0.498** (0.245)	-0.487** (0.219)
CA × rainfall shock	1.095* (0.631)	1.276** (0.608)
<i>Coupea</i>		
CA (= 1)	-1.024 (1.145)	-1.969 (1.281)
rainfall shock	-0.798* (0.433)	-1.207*** (0.419)
CA × rainfall shock	0.411 (1.339)	1.847* (1.094)
Household FE	Yes	Yes
Kleibergen-Paap LM	30.05***	26.68***
Observations	7,643	76,43
Log Likelihood	-16,417	-16,800

Note: Dependent variable is log of yield. All specifications include crop-specific intercept terms and year dummies. In each regression the adoption of CA is treated as endogenous and is instrumented as previously discussed. Column (1) replaces any deviation in rainfall that is with \pm one standard deviation with a zero. Thus, any realized value that is $0.202 \leq R_{wt} \leq 0.749$ is set to zero. Column (2) replaces any deviation in rainfall that is with \pm one half of a standard deviation with a zero. Thus, any realized value that is $0.339 \leq R_{wt} \leq 0.612$ is set to zero. Standard errors clustered by household and crop are reported in parentheses (*p<0.1; **p<0.05; ***p<0.01).