Year 1 Pre-Analysis Plan: Impact of Alternate Wetting and Drying on Farm Incomes and Water Savings in Bangladesh

Ujjayant Chakravorty, Manzoor H. Dar, Kyle Emerick

August 2, 2017

1 Overview and Research Questions

This document describes the analysis to be carried out for the first year of data collected on the project "Impact of Alternate Wetting and Drying on Farm Incomes and Water Savings in Bangladesh". We start by describing the experimental design and main parameters of interest. The document then describes the general specification that will be employed for estimation of the main average treatment effects. Each of the following sections then describes how each of the key impact parameters will be estimated. Finally, the analysis plan closes with a list of blank regression tables that will be populated after data are analyzed. This analysis plan has been written before any of the follow up data have been analyzed.

The treatment being studied is a water-management technique for irrigated rice called Alternate Wetting and Drying (AWD). AWD involves inserting a perforated PVC pipe into the soil to allow the farmer to observe soil moisture below the surface. The AWD guidelines suggest that the farmer let the field dry until the water level reaches 15 cm below the surface — which has a visible marking within the pipe. Once this water level is reached, the farmer should re-irrigate the field up to a level that depends on the current status of the crop. The process of alternatively wetting and drying the field should be practiced up to the time that the crop starts to flower or reproduce. The farmer should keep sufficient water in the field during flowering because the crop water requirements are much higher during flowering relative to the previous vegetative stage of growth. The farmer also drains the field approximately one to two weeks before harvest, regardless of their chosen method of irrigation. AWD is meant to reduced total irrigation withdraw relative to a system where the field is never allowed to dry, i.e. "continuous flooding". In addition, agronomic trials on experiment stations have found that AWD reduces the methane emissions from rice relative to continuous flooding.

The nature of water pricing varies across our sample. This is an important feature of the experiment because we expect the water-pricing regime might be important for the effectiveness of AWD. Water pricing varies mostly — but not entirely — across divisions. Rajshahi has metered and government-run deep tube wells (DTW) where farmers use prepaid cards to pay for each hour of irrigation. Our sample area of Rangpur also has government DTW, but those tube wells are operated based on seasonal contracts where farmers pay per unit of land area, not per unit of water. Mymensingh has a shallower groundwater table and farmers in this region irrigate rice with private shallow tube wells (STW). In this case the pricing regime depends on the contract between the owner and users of the tube well. In many cases the water is paid for on a seasonal basis and the owner of the tube well pays for the diesel fuel or electricity to lift the water. In some other cases there is a component of volumetric pricing because the farmer pays for the fuel or electricity, or because the tube well owner is part of the sample. We will use the baseline data to separate the sample into farmers that pay some form of volumetric prices and farmers that pay entirely area-based charges.

This phase of the experiment focuses on these research questions. All impact parameters will be estimated both for areas with and without volumetric water pricing.

1. What is the impact of AWD on usage of water for irrigation?
2. Does practicing AWD reduce the amount of methane that gets emitted from rice paddies?
3. What is the impact of AWD on rice output, usage of other inputs (particularly herbicides and weeding labor), and overall profitability per acre?

2 Research Strategy

The study is a randomized control trial being carried out in 400 villages spread across 12 upazilas in 3 divisions of Bangladesh: Rajshahi, Rangpur, and Mymensingh. The first year of the study, which this document pertains to, is the boro (dry) season of 2017. Prior to this season, we identified 10 plots cultivated by different farmers that were adjacent to the tube well in each of these villages. This sample of 4,000 farmers serves as our sample for the estimation of impacts. We first randomly allocated treatment at the village level, where the randomization was stratified by upazila. The treatment was administered prior to the planting of boro rice, which occurred from January-March 2017, depending on the area. The treatment involved three components:

1. Delivery of an AWD pipe to each of the 10 farmers.
2. A short-training carried out by a local NGO on how to use the AWD pipe.
3. Assistance by the NGO with installation. The installation was required to be completed on the study plot identified close to the tube well, i.e. the plot identified in the beginning of the experiment.

Farmers in the control group were not provided with any of the above items 1-3.

2.1 Sampling

2.1.1 Sampling Frame

The eligible population is all of the dry-season rice-farming households in Bangladesh, i.e. the areas where AWD could potentially be useful. This population is characterized by small farm sizes, yet high rice productivity during the boro season. Our sample size will be 4,000 farmers — 10 farmers in each of 400 villages. This is of course a very small share of the dry-season rice-farming households in the country.

The 12 upazilas were selected based on local knowledge of IRRI Bangladesh researchers. The selection was based on two primary criteria. First, the upazilas had to be in areas where rice is the dominant crop in the boro season. This was validated using satellite imagery of boro rice areas from IRRI. Second, the upazilas in Rajshahi and Rangpur districts needed to have some area covered by government deep tube wells.

Within these upazilas, eligible villages were next identified. Villages in Rajshahi and Rangpur were only considered to be eligible if they have deep tube well (DTW) irrigation for boro rice. This effectively means that every village in the sample in these two districts has a government tube well run by either the Barind Multipurpose Development Authority (BMDA) or the Bangladesh Agricultural Development Corporation (BADC). We selected all villages in Mymensingh based on the census. The 400 villages in the sample were drawn randomly from the set of eligible villages. Within each village, enumerators went to a village leader and identified the 10 farmers that were cultivating plots close to the tube well used for irrigation. In the event that there was more than one tube well, enumerators were instructed to focus on the tube well with the largest command area. In most cases that tube well services more than 10 farmers. Enumerators asked the village leader for the names of the 10 *closest* farmers. These farmers were then asked for consent to participate and in rare cases farmers that did not consent were replaced with the next closest farmer. The resulting 10 farmers (and the plots close to the tube well) constitute the main sample for measuring the effectiveness of AWD. We have checked using the Bangladesh census that the sample is roughly similar to all households in the districts of Rajshahi, Rangpur, and Mymensingh.

2.1.2 Statistical Power

The power calculations were made for a two-sided test assuming a significance level of 0.05 and a power of 0.8. We based our calculations on the information we had when setting the sample before the baseline survey. The minimum detectable effects are for a sample of 400 villages with 10 farmers per village. We have a 50-50 split between treatment and control villages. The power calculations are as follows:

Crop yield: We used the crop yield measure from a baseline survey for a separate evaluation we are carrying out in Rajshahi. This survey had been completed in 163 villages when the calculations were made. Average yields were 2,220 kg per acre with a standard deviation of 2,595. We use normalized yields (mean 0 and SD of 1) for the power calculation. The intra cluster correlation coefficient (ICC) of yields is 0.029. Partly due to this low ICC, the minimum detectable effect is a modest 0.115 standard deviations. This represents 298 kg or a 13.4 percent effect. Importantly, existing agronomic evidence suggests that AWD has zero or no effect on rice yields. Therefore, if we find the same result, then sufficient power will allow for a precisely estimated null effect — rather than a small point estimate with a large standard error.

Water use: We used a pilot survey in 35 villages to measure the number of irrigations during the aman (wet) season. The average number of irrigations was 3 with a standard deviation of 3.6. The intraclass correlation coefficient is 0.307 - much higher than the ICC for yield. This is likely conservative since aman irrigation is supplemental to rainfall and thus much of the variation is determined by rainfall - which varies little within villages. Our experiment is being carried out in the boro (dry) season, which will likely have a lower ICC. Nonetheless, we obtain a minimum detectable effect of 0.199 standard deviations or a 24 percent decrease in the number of irrigations.

The power calculations were done with the clustersampsi command in Stata.

2.1.3 Assignment to Treatment

Villages were randomly assigned to treatment and control conditions using a random number generator in Stata. The randomization has been stratified by the upazila (sub-district).

2.1.4 Attrition from the Sample

We don't expect attrition to be problematic. While individuals within households may move occasionally, land markets are very thin and therefore the agricultural land is almost always retained and cultivated by a different household member. Our surveys will be conducted with the household member that cultivates the study plot in the event that one household member migrates away. We also don't expect a large number of households to refuse our surveys. We base these conclusions on recent work for a different project in Bangladesh. In that study of 256 villages in Rajshahi, attrition from the study occurred for less than 2% of households and therefore was nowhere near large enough to impact the treatment effect estimates.

2.2 Fieldwork

2.2.1 Instruments

All of our instruments are developed and the data collection takes place using the ODK platform. Our baseline and follow up survey instruments will focus on agricultural production for two plots: the study plot (where AWD is being used by treatment farmers) and one other randomly selected plot for each farmer. These surveys will include sufficient detail to measure profitability per acre for each of these plots. This includes information on yield, price received (if sold), wage labor expenditures by activity, expenditures on material inputs such as fertilizer, herbicides, and pesticides. We will collect family labor for planting, harvesting, and weeding. The family labor will be valued at the village wage rate for the profitability calculation.

In addition to information on profitability, the instruments collect information on water usage — one of the key outcome variables discussed below. There are three components to this: 1) How many times was the field irrigated, 2) How many times was the field drained / dried, and 3) What was the cost per unit of land paid for irrigation water.

We have used a similar instrument before for an ongoing study in Rajshahi Bangladesh and therefore do not need to do extensive pre-testing.

Water levels in fields are recorded using a simple instrument where the enumerator simply enters the farmer's ID information and the amount of water that they physically measure to be in that field. A similar instrument is used for methane measurements. In this case the field staff is taking the samples and not entering the value for the methane flux (since that needs to be calculated in the lab). Instead, the enumerator enters information such as the time of day for the sample, and the temperature for the base which is an input into the calculation of the methane flux done in the lab.

3 Empirical Analysis

3.1 Variables

The outcome variables for input use will be:

1. Number of fertilizer applications, i.e. the number of times fertilizer was applied

2. KG urea per acre

3. KG TSP per acre

4. KG potash per acre
5. KG of other fertilizer per acre
6. Pesticide expenditure per acre
7. Herbicide expenditure per acre
8. Hired labor expenditure for planting
9. Hired labor expenditure for weeding
10. Hired labor expenditure for harvesting
11. Imputed family labor cost for planting (opportunity cost imputed with village average wage)
12. Imputed family labor cost for weeding (opportunity cost imputed with village average wage)
13. Imputed family labor cost for harvesting (opportunity cost imputed with village average wage)

The follow-up survey will also contain self-reported measures of irrigation usage,

14. Number of times irrigation water was applied
15. Number of times field was drained
16. Cost per acre of water

The measures of revenue and profits from the follow-up will be

17. Yield (kilograms of output per acre)
18. Revenue per acre (If the farmer did not sell, then price is imputed with village-average)
19. Profit per acre (revenue per acre - cost for acre of items 5-12 and 15 above)
20. Log yield (kilograms of output per acre)
21. Log revenue per acre
22. Log profit per acre
23. Profit per acre, trim top and bottom 1.5%

24. Profit per acre, robust regression

25. Log profit per acre, trim top and bottom 1.5%

26. Log profit per acre, robust regression

3.2 Balancing Checks

Balance between treatment and control groups will be checked with a regression of the baseline characteristic on a constant, the village-level treatment indicator, and a strata fixed effect. The regression is identical to the main specification for estimating impacts, given in Equation 1 below.

The variables for the balancing check are age of the farmer, years of education, household size, livestock ownership, landholdings, television ownership, refrigerator ownership, tube well ownership, baseline knowledge of AWD, indicator if the study plot was rented or sharecropped, area of the study plot, indicator for volumetric water pricing, number of crops being grown on the study plot, indicator for a rice-rice cropping system (two rice crops being grown), number of irrigations in the dry season, revenue per acre in the dry season, total cost per acre in the dry season, water cost per acre in the dry season, and revenue per acre in the wet season.

We will first check whether attrition is balanced across treatment arms by regression the attrition indicator on the treatment indicator and strata fixed effects. We don't expect balanced attrition to be problematic. Nonetheless, we will check balance for the non-attriting sample using the exact same procedure as above.

3.3 Treatment Effects

3.3.1 Intent to Treat

All of our estimates will be Intention to Treat (ITT). Our main specification will be

$$y_{ivj} = \alpha_j + \beta Treat_{vj} + \varepsilon_{ivj}, \tag{1}$$

where y_{ivj} is the outcome for farmer i located in village v and upazila j, $Treat_{vj}$ is an indicator for the 200 random treatment villages, and ε_{ivj} is the random error term. Standard errors will be clustered at the village level in all regressions. The term α_j is the upazila (randomization strata) fixed effects. Our main parameter of interest will be β which gives the intention to treat (ITT) effect of AWD on the various outcomes of interest. The sample will consist of the 4000 study plots, 2000 of which had AWD installed and 2,000 of which were identified at the start of the experiment, but were part of the control group and thus did not have AWD installed. We have baseline data for all 4,000 plots.

While our main specification will be that in (1), we will also include appendix tables where control variables are introduced. More specifically, the specification will be

$$y_{ivj} = \alpha_j + \beta Treat_{vj} + \delta X_{ivj} + \varepsilon_{ivj}, \tag{2}$$

where everything remains the same as in (1), except for the inclusion of the vector of controls X_{ivj}. This vector will include age of the farmer, years of education, household size, livestock ownership, landholdings, television ownership, refrigerator ownership, tube well ownership, baseline knowledge of AWD, indicator if the study plot was rented or sharecropped, area of the study plot, indicator for volumetric water pricing, number of crops being grown on the study plot, baseline number of irrigations in the dry season, baseline water cost per acre in the dry season, and baseline revenue per acre in the dry season.

Focusing first on water usage, each of the 4,000 fields was visited on two randomly chosen days: one during the first half of the growing season and the other during the second. This means we will have 8,000 observations where an enumerator visited the study plot on an unannounced and random day in order to observe the amount of water in the field.[1] Using these data, the main estimate will be akin to that in (1), except for we will have two observations per farmer instead of one. There will be two dependent variables of interest: the amount of water (cm) in the field on that day, and and indicator variable for whether the field is dry (no water). The regression tables to be estimated for water usage are included in Tables 1 to 7. The heterogeneity is discussed below.

We will also evaluate the change in the entire distribution of water levels using quantile regressions. We don't expect effects of AWD throughout the whole distribution. If AWD farmers irrigate their fields to the same level (when they choose to irrigate) then the distributions will slowly converge at the upper deciles. We will estimate a separate quantile regression for each decile, and report results in a graph where the decile is on the horizontal axis and the quantile regression coefficient (and 95% confidence interval) is on the vertical axis. This analysis will be done both for the entire sample, and separately for farmers with and without volumetric pricing.

We selected 104 random villages for measurements of methane fluxes. Within each village, one farmer was further randomly selected for measurement of methane. Trained individuals have taken 10 methane readings for 24 of these farmers (2 farmers per upazila). Three readings have been taken for each of the remaining 80 farmers. All readings were taken on

[1] The days were randomized at the village level so that the enumerator observed all 10 plots on the same day.

random days. In combination we will have 480 observations to estimate the effect of AWD on methane emissions. The specification will be

$$y_{ivjd} = \alpha_j + \beta Treat_{vj} + \varepsilon_{ivjd}, \tag{3}$$

where in this case y is the methane flux (measured in mg per square meter per hour). We expect some skewness in the distribution of the methane readings given the difficulty of taking gas samples in the field. Our main estimates will trim the top 1.5% of readings. We will also report results on cumulative season emissions where we multiple the hourly flux by the number of hours in a season and the size of the plot (also trimming the top 1.5% of outliers). Standard errors will again be clustered at the village level since we have multiple observations per village.

Adding certain control variables is likely to reduce residual variance and decrease standard errors for methane fluxes. While our main estimates will use only the experimental variation, we will also include an appendix table where we include certain controls in the interest of improving precision. These are:

1. Days after transplanting (DAT), included as a cubic function

2. Temperature

3. Precipitation on the day and day before the reading

4. Use of organic manure

5. Baseline yield

6. Baseline number of irrigations

In the robustness analysis we will also 1) exclude observations where no methane was detected 2) replace the level of the methane flux with the log as the dependent variable 3) check sensitivity to using robust regressions (which place less weight on outlier observations). The regression tables to be estimated are included in Tables 8 to 10.

The follow up survey will be carried out with all 4,000 farmers and it will track production and input use on two plots: the study plot and another randomly selected plot where we collected baseline outcome data. Our main estimates will focus on the study plot and will be generated from a specification like (1). The dependent variables are listed in Section 3.1 above. Unlike the other analysis, this analysis includes multiple outcomes that measure different things. We will adjust the p-values using the Family-wise Error Rate in Anderson (2008 Journal of the American Statistical Association p.1486). The families of outcomes will be grouped as follows: chemical inputs (1-7 above), labor inputs (8-13), irrigation and methane (14-16 plus objectively-measured water levels and methane measurements), and profitability (17-19). The regression tables to be estimated are included in Tables 11 to 16.

Our enumerators took the GPS locations for each of the 4,000 study plots during the baseline survey. We will match these locations to remote sensing data from the SENTINEL satellite of the European Space Agency in addition to Landsat using Google Earth Engine. There are two objectives of this exercise.

First, we want to obtain a measure of the "greenness" of each of these plots (throughout time) to measure whether treatment plots appear any more water stressed as a result of AWD. These data will directly complement the yield analysis described above in. In this case we will have multiple images for each plot, depending on cloudiness. We will take all cloud-free mages in between planting and harvesting to estimate the basic specification

$$y_{ivjd} = \alpha_j + \beta Treat_{vj} + \varepsilon_{ivjd}, \tag{4}$$

where y is the log of either Normalized Difference Vegtation Index (NDVI), Enhanced Vegetation Index (EVI), or Normalized Difference Moisture Index (NDMI) all of which are vegetation indices which are plausibly correlated with plant health. We will also estimate equation 4 more flexibly by estimating the non-parametric relationship between y and days after planting, separately for treatment and control groups. As a final and more parametric approach, we will bin the data into 10-day bins and estimate separate values of β for each 10 day bin, from transplanting to harvesting.

Second, we will test for differences in soil moisture between treatment and control plots, to the extent this is measurable with the imagery. The analysis will proceed identically to the analysis for greenness, except for that we will use either the Normalized Difference Water Index or the soil wetness index.

3.4 Heterogeneous Effects

3.4.1 Intent to Treat

In addition to the average effects across the entire sample, it is important to estimate the effect of AWD separately for farmers with and without volumetric pricing. We will introduce this interaction term in the regression

$$y_{ivjd} = \alpha_j + \beta_1 Treat_{vj} + \beta_2 Volume_{ivj} + \beta_3 Treat_{vj} * Volume_{ivj} + \varepsilon_{ivjd}. \tag{5}$$

In this specification $Volume_{ivj}$ is an indicator for farmers with volumetric pricing and we have added a d subscript to note the day of the water-level observation. We are testing for the same heterogeneity in the methane analysis as well.

In addition to the variation in water pricing, we would expect the difference between treatment and control plots to depend on the time of the growing season. AWD is not supposed to be practiced during flowering and farmers always drain their field during the end of the growing season. We would therefore expect treatment effects to be largest in the pre-flowering period of the growing season. Yet, we do not know the exact time of flowering, and this varies by crop variety and weather conditions. An approximate time for flowering is around 60-80 days after transplanting. We will estimate (5) separately for the period from 0-70 days after transplanting and the period beyond 70 days after transplanting. Since 70 days is an approximation, we will include appendix tables with both 60 and 80 day thresholds.

A simpler way to get at heterogeneity according to time of the growing season is to use the fact that we have randomized the day of the measurements. This allows us to estimate the non-parametric relationship between the two measures of water usage (levels in centimeters and an indicator for dry fields) and days after planting separately for treatment and control groups. This approach has the advantage of not requiring us to impose a specific threshold. Rather, if the treatment effect at all varies by the time of the season, we will observe this by estimating non-parametric Fan regressions of the outcome on days after transplanting. These regressions are estimated separately for the treatment and control plots.

Some farmers in our sample are the "deep drivers" for the village tube well. The driver is the person responsible for collecting payments and planning the allocation of water. These individuals have a greater degree of control over water allocation and planning, as has been observed in qualitative field work prior to the study. While Rajshahi and Rangpur have deep drivers for the government-run tube wells, water allocation in Mymensingh is controlled by tube well owners. A plausible hypothesis is that these individuals will respond better to AWD because they have a greater degree of control over when and how long to irrigate. We will estimate heterogeneous effects where an indicator for drivers or tube well owners is interacted with treatment.

Tables

Table 1: Main effects on water usage

	(1) Level in CM	(2) Dry Field (0/1)	(3) Level in CM	(4) Dry Field (0/1)
AWD Treatment				
AWD Treatment * Volumetric Pricing				
Volumetric Pricing				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
p-Value: Treat+Treat*Volumetric				
Number of Observations	8000	8000	8000	8000
R squared				

Columns 3 and 4 will include the interaction term and the level effect of volumetric pricing.

Table 2: Effects on water usage, by growth stage of the crop

	0-60 Days After Planting		60+ Days After Planting	
	(1)	(2)	(3)	(4)
	Level	Dry	Level	Dry
AWD Treatment				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
Number of Observations				
R squared				

Table 3: Effects on water usage, by growth stage of the crop

	0-70 Days After Planting		70+ Days After Planting	
	(1)	(2)	(3)	(4)
	Level	Dry	Level	Dry
AWD Treatment				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
Number of Observations				
R squared				

Table 4: Effects on water usage, by growth stage of the crop

	0-80 Days After Planting		80+ Days After Planting	
	(1)	(2)	(3)	(4)
	Level	Dry	Level	Dry
AWD Treatment				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
Number of Observations				
R squared				

Table 5: Effects on water usage, by growth stage of the crop

	0-60 Days After Planting		60+ Days After Planting	
	(1)	(2)	(3)	(4)
	Level	Dry	Level	Dry
AWD Treatment				
AWD Treatment * Volumetric Pricing				
Volumetric Pricing				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
p-Value: Treat+Treat*Volumetric				
Number of Observations				
R squared				

All columns will include the interaction term and the level effect of volumetric pricing.

Table 6: Effects on water usage, by growth stage of the crop

	0-70 Days After Planting		70+ Days After Planting	
	(1)	(2)	(3)	(4)
	Level	Dry	Level	Dry
AWD Treatment				
AWD Treatment * Volumetric Pricing				
Volumetric Pricing				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
p-Value: Treat+Treat*Volumetric				
Number of Observations				
R squared				

All columns will include the interaction term and the level effect of volumetric pricing.

Table 7: Effects on water usage, by growth stage of the crop

	0-80 Days After Planting		80+ Days After Planting	
	(1)	(2)	(3)	(4)
	Level	Dry	Level	Dry
AWD Treatment				
AWD Treatment * Volumetric Pricing				
Volumetric Pricing				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
p-Value: Treat+Treat*Volumetric				
Number of Observations				
R squared				

All columns will include the interaction term and the level effect of volumetric pricing.

Table 8: Main effects on methane emissions

	Hourly Flux (mg/m^2/hr)		Season total (mg)	
	(1)	(2)	(3)	(4)
AWD Treatment				
AWD Treatment * Volumetric Pricing				
Volumetric Pricing				
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean in Control				
p-Value: Treat+Treat*Volumetric				
Number of Observations	480	480	480	480
R squared				

Columns 2 and 4 will include the interaction term and the level effect of volumetric pricing.

Table 9: Robustness of methane estimates

	(1) Flux	(2) Flux, Robust	(3) Exclude 0's	(4) Log Flux	(5) Flux	(6) Flux, Robust	(7) Exclude 0's	(8) Log Flux
AWD Treatment								
Strata Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Mean in Control								
Number of Observations	480	480	480	480	480	480	480	480
R squared								

Table 10: Robustness of methane estimates

	(1) Flux	(2) Flux, Robust	(3) Exclude 0's	(4) Log Flux	(5) Flux	(6) Flux, Robust	(7) Exclude 0's	(8) Log Flux
AWD Treatment								
AWD Treatment * Volumetric Pricing								
Volumetric Pricing								
Strata Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Mean in Control								
Number of Observations	480	480	480	480	480	480	480	480
R squared								

All columns will include the interaction term and the level effect of volumetric pricing.

Table 11: Effects on input usage

	Fertilizer					Chemical Expenditures		Hired Labor			Family Labor		
	(1) N apps	(2) Urea	(3) TSP	(4) Potash	(5) Other	(6) Pesticide	(7) Herbicide	(8) Plant	(9) Weed	(10) Harvest	(11) Plant	(12) Weed	(12) Harvest
AWD Treatment													
Strata Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean in Control													
Number of Observations	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000
R squared													

Table 12: Effects on self-reported water use

	(1) Number Irrigations	(2) Times Drained	(3) Cost per Acre
AWD Treatment			
Strata Fixed Effects	Yes	Yes	Yes
Mean in Control			
Number of Observations	4000	4000	4000
R squared			

Table 13: Effects on revenues and profits

				Log:		
	(1) Yield	(2) Revenue Per Acre	(3) Profit Per Acre	(4) Yield	(5) Revenue Per Acre	(6) Profit Per Acre
AWD Treatment						
Strata Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Mean in Control						
Number of Observations	4000	4000	4000	4000	4000	4000
R squared						

Table 14: Effects on input usage

	Fertilizer					Chemical Expenditures		Hired Labor			Family Labor		
	(1) N apps	(2) Urea	(3) TSP	(4) Potash	(5) Other	(6) Pesticide	(7) Herbicide	(8) Plant	(9) Weed	(10) Harvest	(11) Plant	(12) Weed	(12) Harvest
AWD Treatment													
AWD Treatment * Volumetric Pricing													
Volumetric Pricing													
Strata Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean in Control													
p-Value: Treat+Treat*Volumetric													
Number of Observations	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000	4000
R squared													

All columns will include the interaction term and the level effect of volumetric pricing.

Table 15: Effects on self-reported water use

	(1) Number Irrigations	(2) Times Drained	(3) Cost per Acre
AWD Treatment			
AWD Treatment * Volumetric Pricing			
Volumetric Pricing			
Strata Fixed Effects	Yes	Yes	Yes
Mean in Control			
p-Value: Treat+Treat*Volumetric			
Number of Observations	4000	4000	4000
R squared			

All columns will include the interaction term and the level effect of volumetric pricing.

Table 16: Effects on revenues and profits

				Log:		
	(1) Yield	(2) Revenue Per Acre	(3) Profit Per Acre	(4) Yield	(5) Revenue Per Acre	(6) Profit Per Acre
AWD Treatment						
AWD Treatment * Volumetric Pricing						
Volumetric Pricing						
Strata Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Mean in Control						
p-Value: Treat+Treat*Volumetric						
Number of Observations	4000	4000	4000	4000	4000	4000
R squared						

All columns will include the interaction term and the level effect of volumetric pricing.