

Does Micro-irrigation Save Energy?

An Investigation in Gujarat, India

Nick Hagerty & Ariel Zucker *

July 31, 2019

Abstract

Energy efficiency is a global priority, but investments in energy efficiency do not always deliver the expected benefits. This paper studies micro-irrigation systems (MIS), a technology thought to reduce the energy required for irrigation by as much as 70 percent. We installed individual meters to directly measure the energy consumption of several hundred farmers in Gujarat, India, and linked this meter data with survey data to yield a comprehensive view into energy use patterns in smallholder agriculture. We document two facts. One, energy use varies widely across farmers, and this variation is unexplained by factors such as farm area or village geography. Two, MIS users in our sample consume 30 to 40 percent more energy than non-users of MIS. This difference does not appear to be explained by observable differences across farmers nor by rebound effects, suggesting that the energy impacts of MIS under real-world conditions may be disappointing. While these findings are not causal, they highlight a need for increased attention to details of implementation and further research into the actual benefits of resource-conserving technologies.

*University of California, Berkeley, nhagerty@berkeley.edu and adzucker@berkeley.edu. We thank Mustufa Patel, Gokul Sampath, and Neha Doshi for excellent and tireless project management and research assistance; Isha Fuletra and Aditya Madhusudan for superb assistance; and Suresh Bharda along with many survey enumerators and village para-workers for data collection. Thanks to the Aga Khan Rural Support Programme (India) for a productive partnership. We are grateful for financial support from the International Growth Centre, the Weiss Fund for Research in Development Economics, and the Abdul Latif Jameel Water and Food Systems Lab (J-WAFS) at MIT.

1 Introduction

Energy efficiency is a global priority, as it has the potential to slow carbon emissions and help expand energy access. However, investments in energy efficiency do not always deliver the expected benefits. One reason is that new technologies do not always deliver promised efficiency benefits under real-world conditions. Many technologies require complementary investments or behavioural changes to operate as designed, and users may not always dedicate the needed funds and efforts. Another reason may be the rebound effect, in which efficiency improvements from new technologies can encourage greater energy consumption, which offsets some of the gains. Evidence from the United States suggests that many technological energy efficiency investments deliver less benefit than expected (Fowlie et al., 2018), and also that the rebound effect is relatively small in most situations (Gillingham et al., 2015). However, less evidence is available from developing countries, where additional constraints could make suboptimal operation more common and rebound effects larger.

This paper explores the energy efficiency gains from micro-irrigation systems (MIS) in smallholder agriculture in Gujarat, India. MIS, referring to both drip and sprinkler irrigation, is a globally significant technology. Many governments heavily subsidize MIS, including Australia, Nigeria, Rwanda, and our setting of Gujarat, and several other states in India have considered large expansions in MIS subsidies. Gujarat alone disburses approximately 70 million USD in MIS subsidies annually (Venot et al., 2017) - nearly 0.05% of state GDP. Reasons for these subsidies are not always clear and tend to include concerns about energy along with groundwater and productivity. Understanding whether, and why, the returns to MIS justify their expense is critical for governments to make informed policy decisions.

MIS is often thought to reduce the energy consumption required for irrigation by as much as 70 percent (Kumar and Palanisami, 2010; Raman, 2009) by efficiently delivering water in small doses directly to crop root zones, thereby reducing groundwater pumping. However, if less energy or water is required to irrigate the same area of land, farmers may choose to respond by irrigating more land in the same season, or farming in additional seasons. This behavioural response, an example of a rebound effect, may benefit the farmer but would reduce the energy savings.

To study the energy efficiency gains from MIS, we use individual meters to directly measure energy consumption during the growing season for a group of several hundred farmers. We then link this novel dataset of pump-level energy consumption to detailed survey data, producing a uniquely detailed window into energy use patterns in smallholder agriculture. Direct measurements of energy or water use are rare in smallholder agriculture; most previous studies rely on self-reported irrigation intensity or engineering estimates from ideal environments (Kumar and Palanisami, 2010; Raman, 2009; Sinha et al., 2017; Surendran et al., 2016). Our data is also likely more complete than administrative electricity data, since we are able to cover some of the many electricity connections that are unlawful or otherwise unmetered.

One year of these measurements has yielded two striking facts. First, there is enormous variation in energy use across farmers. This variation covers three orders of magnitude and does not appear to be explained by major observable factors such as farm area or village geography. Second, MIS users in our sample consume 30 to 40 percent *more* energy than non-users of MIS. This basic fact would seem to counter the common belief that MIS conserves energy. We propose and explore three potential explanations for this fact: selection bias, rebound effects, and non-conservation.

We find little evidence for selection bias, at least from observable factors. In the absence of a

controlled experiment, we use Mahalanobis distance matching methods to try to create groups of MIS users and non-users who are observationally as similar as possible. We find that MIS users consume substantially more energy than non-users even after adjusting for a wide range of farm, socioeconomic, and demographic characteristics, and even when comparing within the same village. The same result persists across all specifications, including linear regression, nearest-neighbor matching, and kernel matching algorithms. There may still be selection bias from factors we were not able to measure in our survey, but the factors that seem likely to be most important are accounted for.

We also find little evidence for a rebound effect in MIS. We investigate the rebound effect by comparing how MIS affects energy use in two ways: per-hectare per-season (i.e., the direct savings), and in total (i.e., the net energy savings, including any rebound effect). This first measure shows the true energy saving achieved by micro-irrigation in a real-world setting, while the difference between the two measures reveals the impact of behaviour changes on energy consumption. We generally find that MIS users also consume 30 to 40 percent more energy on a per-hectare basis, though this outcome is noisier and coefficients are sometimes smaller than for total energy. This result suggests that increased cropped area does not explain the greater energy use under MIS irrigation.

To try to further rule out selection bias from unobserved factors, we also explore a natural experiment that generates a discontinuity in the price of MIS systems available to farmers. Government subsidy levels are based on discrete categories in landholding size, which creates two similar groups of farmers above and below the discontinuity who face different prices for MIS purchase. However, despite the higher prices faced for farmers just above the landholding size cut-off, we do not see a discontinuous decrease in the probability of MIS use at the cut-off in our data. A much larger dataset would be needed to take advantage of this potential research design.

Limited evidence for either selection bias or rebound effects points toward the third explanation: non-conservation. It may indeed be true that MIS users consume more energy than traditional irrigators. MIS very well may conserve water, but it is possible that these water savings are not translating into energy savings in the typical farmers installation. This situation could arise if, for example, farmers do not invest in complementary technologies such as downsizing their pumps or installing gravity-based storage tanks. It could also arise if farmers use water with high mineral content or do not perform optimal maintenance, resulting in blockage and over-pressurization.

It is important to keep in mind that these findings are merely suggestive and not causal. Because there is no useable natural experiment available in our setting, we cannot be sure that MIS users and non-users are systematically different in unobserved ways. So far we only have one year of data, so we cannot track farmers over time and instead are limited to a cross-sectional analysis. Still, our data may be useful given the dearth of direct measurements of energy consumption either specifically for MIS in India, or more broadly for agriculture in developing countries.

The remainder of the paper proceeds as follows. Section 2 describes the setting and dataset. Section 3 describes the basic facts in the data. Sections 4 and 5 describe the results from the matching and regression discontinuity designs, respectively. Section 6 concludes.

2 Sample and Data

The study sample consists of farming households in the water-scarce region of Saurashtra in Gujarat, India. An initial group of farmers were recruited from lists compiled by partner organization, the Agha Khan Rural Support Programme (AKRSP), while implementing two agricultural development initiatives in Saurashtra: Farmer Interest Groups (FIGs) and Drip Pools (DPs). FIGs are village-level groups formed with AKRSP support and trained in best practices for cotton cultivation and natural resource management. Drip Pools are revolving zero-interest loan mechanisms for the purchase of Drip Irrigation systems administered by AKRSP. Farmers on the lists had expressed interest or participated in one of the two initiatives. In order to participate in the study, farmers were required to irrigate with groundwater¹ using an electric pump,² and had to be willing to install a meter on their groundwater pumpset. When farmers had multiple pumpsets, we selected the well with an electric pump that was used to irrigate their primary farm. This yielded an initial sample of farming households who represent a group of farmers who are broadly interested in micro-irrigation systems and would be likely to voluntarily adopt MIS under the current subsidy regime.

In order to understand how micro-irrigation impacts energy use among this sample, we collect two types of data. First, we directly measure irrigation intensity - our primary outcome of interest - using hours-of-use meters from the full initial sample. Second, in order to identify plausibly exogenous variation in micro-irrigation adoption, we surveyed a subset of these farmers regarding their MIS use, as well as agricultural, demographic, and socioeconomic characteristics.

We measure energy use for irrigation with hours-of-use meters installed on the electric pump starter of farmers pumpsets. The meters measure the total hours of irrigation done by the farmer, and offer several advantages over electricity meters and water meters. First, they are inexpensive (approximately one tenth of the cost of water meters, for example). Second, they are easy and safe to install: in contrast to water meters which must be fit to an irrigation pipe and can cause water blockages, a single hours-of-use meter can be used on nearly all pump starters in the region. Third, while farmers tend to be suspicious of electricity metering (which they view as a potential threat to existing unmetered and subsidized electricity supplies), hours-of-use meters are widely accepted by farmers. Finally, hours-of-use can be converted into energy and water consumption using field measurements. Meters were installed in October and read once per month from November to March by AKRSP field staff using a tablet-based survey. This yields a five-month panel of hours-of-irrigation for each farmer.

We complement the hours-of-irrigation dataset with a survey of 400 of the metered farmers conducted at the conclusion of the growing season. The survey serves three key purposes. First, we collect information that allows us to better interpret hours-of-use and convert it into energy consumption. Second, the survey identifies which farmers use MIS. Finally, we collect information on observable characteristics that may confound estimates of how MIS impacts energy use. In particular, the survey data include whether or not a farmer has used MIS, fixed characteristics that predict MIS use, such as landholding size and household education levels, as well as potential behavioral outcomes such as area cropped and crops chosen. We record the horsepower of

¹We defined farmers as irrigating using groundwater if they met three criteria: they had irrigated their primary farm the previous winter season, they planned to irrigate their primary farm the next winter season, and they irrigated their primary farm from a groundwater source.

²Farmers were required to use an electric pump as hours-of-use meters cannot be installed on diesel pumps.

the metered pump in order to convert hours of use into energy consumption using the formula

$$E = \frac{P}{\eta}t$$

where E is energy consumed, t is duration of pump operation, P is the power rating of the pumps motor, and η is the motor efficiency.

We use these data to construct two outcomes for measuring the impacts of MIS using this dataset. First, we measure total impact of MIS on energy consumption using the natural log of monthly kWh consumed. This outcome is undefined for the seven farmers who did not irrigate, three of whom were MIS users. Second, we measure the mechanical effect of MIS on energy consumption per cropped area using the natural log of monthly kWh consumed per hectare cultivated. This outcome is undefined for the the 30 farmers for whom we are missing data on cultivated area.

3 Basic facts

3.1 Summary statistics

As a first look at our survey data, Table 1 shows the basic characteristics of the farmers in our sample. They are predominantly smallholder farmers, with a mean area of their primary farm of just 1.5 hectares. (For context, the Indian government typically defines farms holding less than 2 hectares as “small and marginal farmers”).³ Cultivated area is a bit larger than farm size on average, reflecting that some farmers are able to harvest one crop and grow another on the same land within the year. Nearly all irrigate, about two-thirds used micro-irrigation (MIS), and the vast majority grow cotton (which does not exclude also growing other crops). Most farms have one well, but some have two or three;⁴ the average depth of water in these wells was 23 meters. About half of farmers provide water to other farms from their wells, while only one in five receive water from other farms.

Turning to socioeconomic and demographic characteristics, the vast majority of farmers’ houses are made of high-quality materials (*pucca* floors and roofs) and are electrified, though only about one-third have a private water tap. On average, the farms have about four cows or buffalo, one bullock, and two pieces of mechanized farm equipment. Nearly all own their land rather than rent; about two-thirds earned income from businesses besides the farm, and about one-third earned income from working outside the farm. Most households in our sample identify as Hindu and one of the “scheduled caste/scheduled tribe/other backward caste” designations. Most are literate and have at least primary education, while about one-third have post-secondary education.

3.2 Energy consumption varies wildly

Next, we show a first look at our meter data. Panel (a) of Figure 1 plots the distribution of pumping time among farmers in our sample as a histogram.⁵ The horizontal axis in this graph is shown on a logarithmic scale. Pumping time varies enormously across farmers: The mode is

³A caveat here is that we only gathered data on each farmer’s self-defined primary farm. Some farmers may have multiple non-contiguous farms, for a greater landholding total.

⁴Farmers with more than three wells on their primary farm were excluded from the survey collection.

⁵Pumping time is shown in hours per month; to construct this we sum the hours of use measured across all four monthly meter readings, divide by the number of days elapsed between meter installation and the final meter reading, and scale to month.

	Mean	SD	Min	Max	N
A. Agricultural statistics					
Farm size (ha)	1.45	0.99	0.16	6.80	400
Cultivated area (ha)	1.76	1.54	0.00	10.20	400
Irrigated (share)	0.98	0.13	0.00	1.00	400
Used MIS (share)	0.69	0.46	0.00	1.00	400
Cotton grown (share)	0.91	0.29	0.00	1.00	400
Active wells on primary farm	1.30	0.59	0.00	3.00	400
Groundwater depth (approximate, meters)	23.15	35.65	0.00	213.36	338
Provides water to other farm(s)	0.47	0.50	0.00	1.00	400
Receives water from other farm(s)	0.20	0.40	0.00	1.00	400
B. Socioeconomics					
Pucca floor (share)	0.80	0.40	0.00	1.00	400
Pucca roof (share)	0.95	0.21	0.00	1.00	400
Household electrified (share)	0.95	0.21	0.00	1.00	400
Household has private water tap (share)	0.34	0.48	0.00	1.00	400
Cows or buffalo	3.79	3.60	0.00	35.00	400
Bullocks	0.91	0.89	0.00	4.00	400
Mechanized farm equipment	1.77	1.43	0.00	12.00	400
Earned agricultural income from own land (share)	0.97	0.17	0.00	1.00	400
Earned agricultural income from rented land (share)	0.03	0.18	0.00	1.00	400
Earned sharecropping income (share)	0.15	0.36	0.00	1.00	400
Earned labor income (share)	0.34	0.48	0.00	1.00	400
Earned business income (share)	0.70	0.46	0.00	1.00	400
C. Demographics					
Household size	6.07	2.81	1.00	24.00	399
Religion: Hindu (share)	0.95	0.23	0.00	1.00	400
Religion: Muslim (share)	0.00	0.00	0.00	0.00	400
Caste: SC/ST/OBC (share)	0.81	0.39	0.00	1.00	400
Head of household literate (share)	0.91	0.29	0.00	1.00	400
No education	0.01	0.10	0.00	1.00	400
Primary or secondary education (share)	0.71	0.52	0.00	2.00	400
Post-secondary education (share)	0.31	0.46	0.00	1.00	400

Table 1: Summary statistics.

Note: This table displays summary statistics for the sample of farmers with both baseline data and hours-of-use data.

around 20 hours per month, but many farmers pump more than 100 or fewer than 3 hours per month. Means and standard deviations are listed in Table 2; the mean of pumping time is 33.3 hours per month.

Perhaps surprisingly, this wide variance cannot be explained by cultivated area. Panel (b) of Figure 1 plots the distribution of pumping time per hectare cultivated. (Cultivated area is defined as the sum of the areas of all crops planted, so it can be greater than one.) This histogram is equally wide, and in fact the variance of this ratio is larger.

Our main outcome of interest, however, is not pumping time but rather energy use. To calculate energy consumption, we use the formula in Section 2 along with survey data on each pump’s rated brake horsepower and assuming a 74% motor efficiency.⁶ Panels (c) and (d) of Figure 1 plot the distribution of energy consumption, and energy consumption per cultivated area, across farmers. The distribution remains quite wide and are not explained by cultivated area. As Table 2 shows, the average farmer in our sample uses 172 kilowatt-hours (kWh) of electricity per month.

	Full Sample		No MIS	Use MIS	Difference=0 p-value
	Mean	SD	Mean	Mean	
A. Pumping time					
Pumping time (hours/month)	33.27	47.31	31.02	34.28	0.57
Ln(Pumping time)	2.73	1.38	2.47	2.84	0.02
Pumping time per area cropped (hrs/mo/ha)	28.29	74.20	30.33	27.33	0.78
Ln(Pumping time per area cropped)	2.34	1.50	2.17	2.42	0.15
B. Energy consumption					
Energy used (kWh/month)	172.00	265.53	163.13	175.98	0.68
Ln(Energy used)	4.27	1.45	4.01	4.39	0.02
Energy per area cropped (kWh/month/ha)	135.74	291.49	137.73	134.81	0.94
Ln(Energy per hectare cultivated)	3.89	1.54	3.71	3.98	0.13
Sample size					
Number of individuals	400		124	276	

Table 2: Hours of electricity used in full sample, and by MIS-usage.

3.3 MIS users consume more energy

To complete our roundup of basic facts, we break down energy consumption by whether farmers use MIS. Panel (a) of Figure 2 plots the kernel density of energy consumption for farmers who use MIS (thick dashed line in red) and those who do not (thin solid line in blue). While both distributions have high variance, the distribution of energy use for MIS users is shifted noticeably to the right, indicating that they consume more energy than non-users of MIS. Table 2 confirms numerically that MIS users consume more energy on average: 176 kWh per month vs. 163 for non-users of MIS. In natural logs, the difference in means is 0.38, which can be interpreted as approximately 38 percent greater energy consumption. A *t*-test rejects the idea that these log means are the same at a 95 percent confidence level ($p = 0.02$).

Is this just because MIS users grow more crops? Panel (b) of Figure 2 shows the same comparison for energy consumption per area cropped. The distribution of energy consumption for MIS users is still shifted right relative to non-users of MIS, though the difference is not as noticeable. Quantitatively, the difference in log means is no longer statistically significant ($p = 0.13$), but the

⁶Motor efficiency is unknown without intensive physical testing. Absent this information, we simply assign all pumps a central value from the literature. Because the assumed value is a multiplicative factor, different choices will not affect the results when outcome variables are in logarithms. Results could be biased if actual motor efficiency is correlated with either brake horsepower or pumping time.




Figure 1: Histograms of average energy consumption in the experimental sample

Note: Figure displays histograms of energy consumption over the winter 2018-19 cropping season. All histograms are shown on a log scale. Panel (a) shows the distribution of average monthly pumping hours, calculated as total hours at the final meter reading divided by the number of months the meter was read. Panel (b) shows the distribution of average monthly pumping hours normalized by the total cultivated area in the winter 2018-19 cropping season. Panel (c) shows the distribution of average monthly energy consumed by the metered pump, and Panel (d) shows the distribution of average monthly energy consumption normalized by cultivated area. The figures show that the distribution of energy consumed for irrigation purposes is very dispersed, even after controlling for the total area cropped.

difference is still fairly large at 0.28. It seems that cultivated area might explain some, but not all, of the difference in energy use among MIS users.

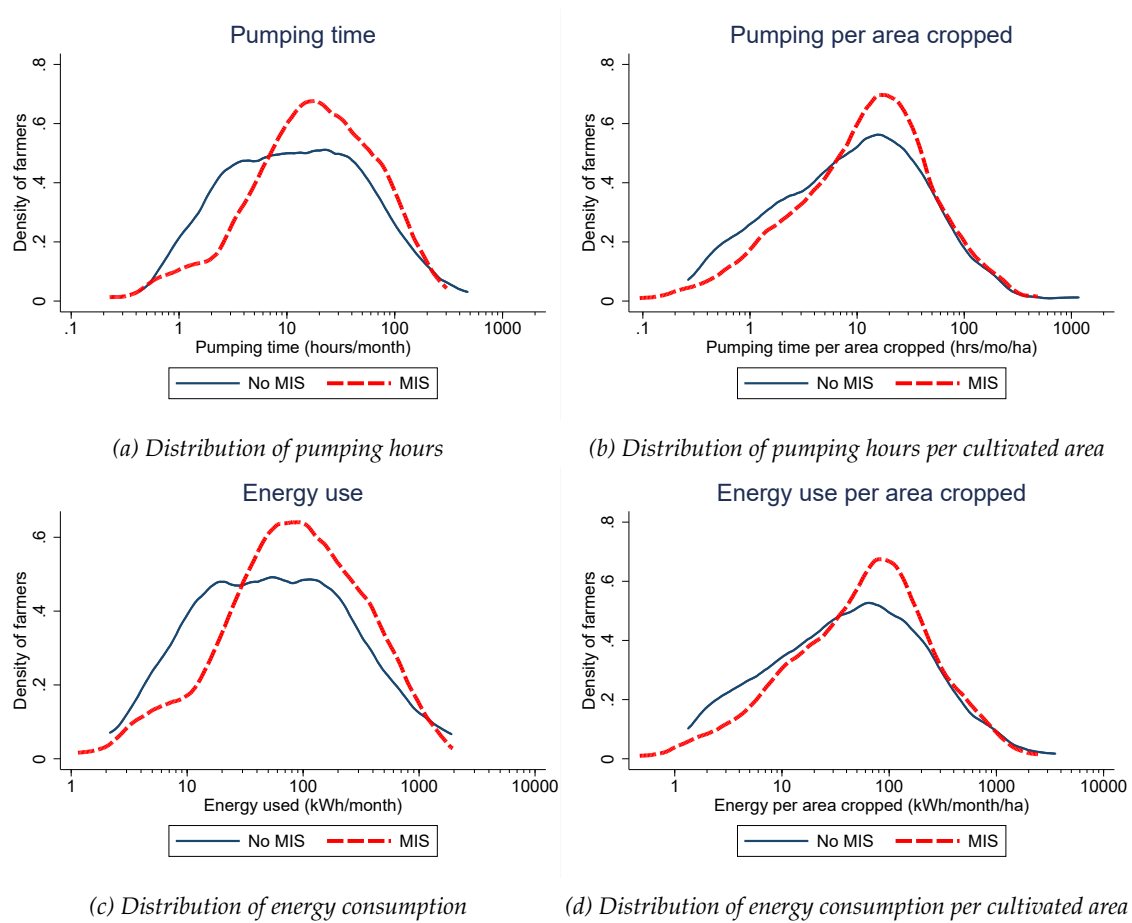


Figure 2: Histograms of average energy consumption in the experimental sample

Note: Figure displays kernel density plots of energy consumption for those with and without micro-irrigation systems (MIS) over the winter 2018-19 cropping season. The x-axes of all plots are drawn on a log scale. Panel (a) shows the distribution of average monthly pumping hours, calculated as total hours at the final meter reading divided by the number of months the meter was read. Panel (b) shows the distribution of average monthly pumping hours normalized by the total cultivated area in the winter 2018-19 cropping season. Panel (c) shows the distribution of average monthly energy consumed by the metered pump, and Panel (d) shows the distribution of average monthly energy consumption normalized by cultivated area. The figures show an overall shift to the right of energy use and energy use per hectare by MIS users.

3.4 Why do MIS users consume more energy?

This basic descriptive fact in our data seems to contradict the conventional wisdom that MIS is a resource-conserving technology. How can this be? We propose three candidate explanations:

1. **Non-conservation:** Regardless of whether it saves water, MIS actually takes more energy than traditional irrigation to irrigate the same crops under real-world conditions.
2. **Rebound effects:** MIS conserves water and/or energy, allowing its users to grow more crops and irrigate more area than they otherwise would.
3. **Selection bias:** MIS users are fundamentally different from non-users in any number of ways – perhaps they use different farming methods or have better access to water – and so

comparing their raw data is not useful.

The rest of our analysis attempts to distinguish between these three potential explanations.

4 Effects of MIS via regression and matching

4.1 Methods

To learn how MIS affects energy use and whether there are any rebound effects, we would like to know what MIS users would have done had they not adopted MIS. This is impossible, so instead we need to try to construct a group of non-users of MIS that are very similar to the MIS users in all other ways, so that they form a plausible comparison group.

To reduce selection bias, we apply regression and matching techniques. These methods can help adjust for observed differences between the groups of MIS users and non-users, re-weighting group members in order to construct two groups that are as similar as possible except for the fact that one uses MIS and the other does not. However, no matter how comprehensive the set of observed control variables, the possibility remains that there are additional unobserved factors that are different between MIS users and non-users. After all, there was a reason that some farmers chose to adopt MIS and others did not. Because we cannot fully eliminate selection bias, the results here can be interpreted only as correlational rather than causal.

Regression. Our regression specifications take the form:

$$Energy_i = \alpha + \beta MIS_i + \mathbf{X}'_i \boldsymbol{\Pi} + \varepsilon_i$$

where $Energy_i$ is the average per-month energy consumption measured for the metered well of farmer i , MIS_i is a binary variable indicating whether farmer i used MIS on their primary farm, and \mathbf{X}_i is a set of covariates. Standard errors are calculated using the Huber-White heteroskedasticity-consistent estimator.

Although our baseline survey provides us with a large set of possible covariates, selecting them is not completely straightforward. Many characteristics of farming, cropping, and irrigation patterns are likely determined after the decision of whether to adopt MIS. These may actually be outcome variables – channels through which the effects of MIS operate. Therefore, the best controls are *pre-treatment* variables – those that are unlikely to be affected by MIS adoption. We form several groups of control variables that we refer to throughout the analysis. They are:

- **Village fixed effects:** Binary indicators for each of the 44 villages in our sample.
- **Agricultural controls:** Total area of primary farm; whether the farm is larger than two hectares; whether the metered well is a borewell; number of active wells on the primary farm; whether wells on primary farm also irrigate any other farms; whether farm uses water from wells on other farms; rated brake horsepower of the electric pump on the primary well.
- **Economic controls:** Whether household has a *pucca* floor; whether household has a *pucca* roof, whether household is electrified; whether household has a private water tap; number of cattle or buffalo; number of bullocks; number of pieces of mechanized farm equipment; whether household earned agricultural income from own land; whether household earned agricultural income from rented land; whether household earned sharecropping income;

whether household earned labor income; whether household earned business income outside of the farm.

- **Demographic controls:** Number of people in the household; whether household identifies as Hindu; whether household identifies as SC/ST/OBC, whether head of household is literate, whether head of household has any formal education; whether head of household has primary or secondary education; whether head of household has post-secondary education.

In addition to these, we form one group of controls that may violate the principle of choosing only pre-treatment variables. We include these variables in some specifications despite this because they might be especially crucial in explaining the differences in pumping amounts.

- **Behavioral controls:** Whether grew cotton; depth to water in metered well; inverse depth to water in metered well.

Results using this group of control variables should be interpreted with the understanding that these variables may block off some channels of the effects of MIS.

Matching. We also apply matching methods that estimate treatment effects by forming explicit matches between observations in our data. For each farmer with MIS, these methods attempt to locate the farmer or farmers without MIS who are otherwise most similar, according to our survey variables. Regression implicitly makes the same sorts of comparisons but also relies on linear extrapolation; matching makes the comparisons explicit and better enforces that they actually take similar values between MIS and non-MIS farmers (i.e., common support).

The variables we use for matching are the same sets of variables as listed above for regression covariates. We use two matching methods: nearest neighbor and kernel matching. Both are based on Mahalanobis distance matching, which calculates the pairwise similarity of observations across all matching variables in a way that takes into account the variance and covariance of each of the variables (Elizabeth A. Stuart, 2010). The difference between the two methods is that nearest neighbor matching compares each MIS farmer to the single non-MIS farmer with the closest Mahalanobis distance, while kernel matching compares each MIS farmer to all non-MIS farmers within a fixed Mahalanobis distance, called a kernel. The number of farmers within this kernel may be one, zero, or multiple.

4.2 Results

Table 3, Panel A reports coefficients from regressing the natural log of energy consumption on MIS use, along with different sets of control variables. Column 1 shows that the coefficient with no control variables is 0.375, indicating that the raw mean of energy consumption is approximately 38 percent higher for MIS users than for non-users of MIS. This coefficient is numerically equal to the difference in log means shown in Table 2. It is significantly different from zero at a 95 percent confidence level.

Columns 2-5 show the coefficients from including each group of control variables listed above (village fixed effects, agricultural, economic, and demographic controls), while column 6 shows the coefficient from a regression that includes all four groups of controls. Column 7 implements the post-double-selection Lasso methodology of Belloni et al. (2017). This method guards against model overfitting in the presence of many control variables by selecting only the subset of controls that are most relevant in predicting either the outcome variable (energy consumption) or the

independent variable of interest (MIS use). Finally, column 8 uses the “behavioral” set of controls which risk blocking some channels of the effects of MIS.

Across all of these specifications, the coefficient on MIS use is stable, ranging only between 0.340 and 0.471 and remaining statistically significant. Because the outcome is in natural logs, these coefficients are semi-elasticities, meaning they can be interpreted as percentage changes. Users of MIS still use 35 to 50 percent more energy than non-users, even after adjusting for a barrage of controls that describe the farm (including land area, water sources, and water availability as measured by water depth), a rich set of household characteristics (including many demographic and socioeconomic measures), and any unobserved factors that are common to a particular village (captured by the village fixed effects).

MIS users appear to consume more energy in total, but what about energy per hectare? Panel B of Table 3 reports estimated coefficients from the same regression specifications described above, except that the dependent variable is energy use per hectare cultivated. These results show that, even on a per-hectare basis, MIS users still consume 27 to 50 percent more energy than non-users. In some specifications (columns 1-5 and 8), the estimated effect of MIS is a bit smaller, with slightly larger standard errors. However, in the regressions with the full and Lasso-selected sets of controls (columns 6-7), the effect is at least as large as in Panel A and statistically significant. None of the estimates for log energy per hectare are statistically different from their counterparts for log energy.

Turning to the matching approach, Table 4 reports average treatment effects estimated using matching methods. Columns 1-3 show the estimates from nearest-neighbor matching, while columns 4-6 show the estimates from kernel matching. Within each of these groups, the first column includes the agricultural, economic, and demographic controls, the second column adds village fixed effects, and the third column matches on the possibly endogenous behavioral controls. The effect of MIS on log energy consumption (Panel A) is approximately 40 to 50 percent across all of these specifications. These effects are larger than the estimates produced using regression, and they are all statistically significant. The effect on MIS on log energy consumption per cultivated area (Panel B) appear slightly smaller and noisier, just as in the regression results, but again they are not statistically different from the Panel A results.

The fact that the effects are of a similar magnitude when examining energy consumption on a per-hectare basis suggests that there is little evidence for a rebound effect. The rebound effect hypothesis holds that MIS use conserve resources on a per-hectare basis, encouraging farmers to expand production. If this were true, the effects of MIS on energy consumed per unit area would be smaller than the effect on total energy consumed. Our results suggest that increased cropping area does not explain the observed fact that MIS users consume more energy.

4.2.1 Water availability as an unobserved factor

The evidence from both regression and matching methods suggests that selection, at least on observed characteristics, is not responsible for the difference between users and non-users of MIS. If anything, the estimated difference is larger after adjusting for a large number of farm and household characteristics. However, it remains possible that users and non-users of MIS are different in ways that our survey variables do not capture.

What are these unobserved characteristics that might explain the large differences across farmers in irrigation amounts? The farmers in our sample are nearly all growing the same crop (cotton), share social networks, and have access to very similar input and output markets, so they are unlikely to be taking dramatically different approaches to agricultural production. Informa-

OLS Regressions of Energy Consumption on MIS Use and Controls								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Log Energy Consumption								
Used MIS	0.375** [0.164]	0.372** [0.159]	0.352** [0.163]	0.378** [0.163]	0.389** [0.166]	0.471*** [0.179]	0.436*** [0.169]	0.340* [0.173]
R ²	0.014	0.248	0.243	0.039	0.057	0.416		0.161
Observations	391	387	330	391	390	325	330	329
B. Log Energy Consumption per Hectare Cultivated								
Used MIS	0.272 [0.180]	0.329* [0.185]	0.299* [0.178]	0.290 [0.182]	0.300* [0.181]	0.500** [0.201]	0.440** [0.193]	0.301 [0.190]
R ²	0.007	0.219	0.198	0.032	0.047	0.419		0.134
Observations	361	358	304	361	361	300	304	303
Controls: Village FEs		X				X	X	
Agricultural			X			X	X	
Economic				X		X	X	
Demographic					X	X	X	
Behavioral								X
Lasso selection							X	

Table 3: Usage of micro-irrigation systems (MIS) and energy consumption.

Note: This table displays regressions of energy consumption and energy consumption per hectare cultivated on a dummy for whether the farmer uses MIS. Robust standard errors are in brackets. Significance levels: * 10%, ** 5%, *** 1%.

tion and education about optimal irrigation practices could be a factor, but anecdotally, NGOs and extension services are widespread in our study region, and farmers often insist they know how much water they should be applying to their crops.

Instead, the most obvious factor is water availability: some farms simply have better groundwater availability than others. The hydrogeology of our study region is complex and it is often difficult to predict where drilling a well will yield abundant water, or which areas have interconnected aquifers. Our list of “behavioral controls” includes measured depth to groundwater level, but this is likely an imperfect proxy since water availability is not just about water level but also how quickly the well fills back up after pumping.

Another way we can try to adjust for water availability is by making comparisons only among farmers whose well did not go dry during the period of meter reading. This is also an imperfect proxy, but it at least allows us to exclude the most egregious cases, in which water availability was so poor that the well went completely dry at some point. A caveat is that a well going dry may itself be endogenous: if MIS conserves water, MIS users will be less likely to find their well dry. Farmers remaining in the sample are either MIS users or MIS non-users who pumped less, introducing an artificial positive correlation between using MIS and pumping more. Still, we think this comparison is worth considering, since it might help to reduce selection bias in water availability.

Tables 5-6 present the same results from regression and matching as Tables 3-4, but for the subsample of farmers whose wells did not go dry during the study period. Regression coefficients reported in Table 5 tend to be smaller than for the full sample, for both log energy consumption and log energy consumption per hectare. Matching estimates reported in Table 6 are also smaller for kernel matching but volatile for nearest-neighbor matching. Because this sample is

Mahalanobis Distance Matching Regressions of Energy Consumption on MIS Use						
	(1)	(2)	(3)	(4)	(5)	(6)
A. Log Energy Consumption						
main						
Used MIS	0.507** [0.201]	0.519** [0.231]	0.494** [0.211]	0.461*** [0.157]	0.453** [0.195]	0.328** [0.163]
R^2						
Observations	390	390	329	390	390	329
B. Log Energy Consumption per Hectare Cultivated						
main						
Used MIS	0.406* [0.215]	0.467* [0.267]	0.410* [0.234]	0.337* [0.189]	0.356 [0.246]	0.278 [0.178]
R^2						
Observations	361	361	303	361	361	303
Match variables: Village FEs		X			X	
Agricultural	X	X		X	X	
Economic	X	X		X	X	
Demographic	X	X		X	X	
Behavioral			X			X
Matching algorithm	nearest-neighbor	nearest-neighbor	nearest-neighbor	kernel	kernel	kernel

Table 4: Impact of micro-irrigation systems (MIS) on energy consumption: Matching estimates.

Note: This table displays matching-estimates of the impact of MIS use on energy consumption and energy consumption per hectare cultivated. Robust standard errors are in brackets. Significance levels: * 10%, ** 5%, *** 1%.

smaller than the full sample, the standard errors are larger for all of these estimates. This means that we cannot statistically distinguish most of these estimates from a zero effect, nor can we distinguish them from the results for the full sample. Evidence from this sub-sample is therefore inconclusive.

OLS Regressions of Energy Consumption on MIS Use and Controls: Farmers whose wells do not dry								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Log Energy Consumption								
Used MIS	0.218 [0.225]	0.291 [0.229]	0.118 [0.197]	0.242 [0.229]	0.211 [0.233]	0.189 [0.224]	0.276 [0.205]	0.219 [0.202]
R^2	0.004	0.241	0.215	0.033	0.064	0.396		0.153
Observations	248	245	248	248	248	245	248	247
B. Log Energy Consumption per Hectare Cultivated								
Used MIS	0.192 [0.245]	0.283 [0.260]	0.133 [0.220]	0.277 [0.239]	0.205 [0.254]	0.354 [0.233]	0.406* [0.241]	0.208 [0.224]
R^2	0.003	0.248	0.158	0.040	0.028	0.427		0.148
Observations	226	223	226	226	226	223	226	225
Controls: Village FEs		X				X	X	
Agricultural			X			X	X	
Economic				X		X	X	
Demographic					X	X	X	
Behavioral								X
Lasso selection							X	

Table 5: Usage of micro-irrigation systems (MIS) and energy consumption for farmers with plentiful water.

Note: This table displays regressions of energy consumption and energy consumption per hectare cultivated on a dummy for whether the farmer uses MIS. The sample is limited to farmers whose wells do not go dry for the full season, and therefore have access to plentiful water. Robust standard errors are in brackets. Significance levels: * 10%, ** 5%, *** 1%.

5 Effects of MIS via regression discontinuity

Another way we can try to rule out selection bias from unobserved factors is by looking for a natural experiment. In this section, we explore one possible natural experiment for MIS adoption. Farmers in Gujarat have faced different prices for MIS over the last decade, with eligibility for higher subsidies beginning at a sharp discontinuity in land area. Since 2007, the Government of Gujarat has heavily subsidized MIS. For most of this period, the subsidy has been set at 50% of the purchase price for farmers with more than two hectares of land and 60% for farmers with less than two hectares of land; these amounts were later increased to 60% and 70%. If the price difference introduced by this policy generated a discontinuous jump in probability of adoption at the two-hectare cutoff, it would set the stage for a regression discontinuity approach to understanding the impacts of MIS. In particular, because farmers just below and above the two-hectare farm size cutoff would be identical but for an infinitesimal difference in farm size, any discontinuous change in energy use between farmers at the cutoff could be causally attributed to the additional MIS adoption caused by the discontinuous subsidy schedule.

Unfortunately for our analysis, there is no evidence that the higher subsidy amount for farmers below the two-hectare cutoff has led to increased adoption of MIS. Panel (a) of Figure 3 plots the share of farmers adopting MIS at different farm size bins in blue. MIS adoption is not sig-

Mahalanobis Distance Matching Regressions of Energy Consumption on MIS Use: Farmers whose wells do not dry						
	(1)	(2)	(3)	(4)	(5)	(6)
A. Log Energy Consumption						
main						
Used MIS	0.284 [0.271]	0.836 [0.553]	0.255 [0.197]	0.332 [0.223]	0.263 [0.301]	0.280* [0.155]
R^2						
Observations	248	248	247	248	248	247
B. Log Energy Consumption per Hectare Cultivated						
main						
Used MIS	0.179 [0.281]	0.982 [0.608]	0.109 [0.197]	0.320 [0.283]	0.268 [0.293]	0.204 [0.204]
R^2						
Observations	226	226	225	226	226	225
Match variables: Village FEs		X			X	
Agricultural	X	X		X	X	
Economic	X	X		X	X	
Demographic	X	X		X	X	
Behavioral			X			X
Matching algorithm	nearest-neighbor	nearest-neighbor	nearest-neighbor	kernel	kernel	kernel

Table 6: Impact of micro-irrigation systems (MIS) on energy consumption: Matching estimates for farmers with plentiful water.

Note: This table displays matching-estimates of the impact of MIS use on energy consumption and energy consumption per hectare cultivated. The sample is limited to farmers whose wells do not go dry for the full season, and therefore have access to plentiful water. Robust standard errors are in brackets. Significance levels: * 10%, ** 5%, *** 1%.

nificantly higher just below two hectares; in fact, if anything, the share of farmers using MIS is slightly lower below the cut-off. Consistent with the finding in Section 4.2 that increased MIS adoption is associated with increased energy utilization, Panel (b) of Figure 3 shows that energy use also decreases slightly below the two hectare cutoff; however, this decrease is also statistically indistinguishable from zero.

Why might this be? It is likely that the failure to find a discontinuity at the cutoff is simply due to insufficient data. The histogram of farm size in Figure 3 (in orange) shows that very few farmers in our sample have farms less than two hectares. It may also be that the increased subsidy amounts was too small to induce large variation in MIS adoption. More data from farmers near the two-hectare cut-off will be necessary to assess whether an regression discontinuity design can be used to estimate the causal impacts of MIS.

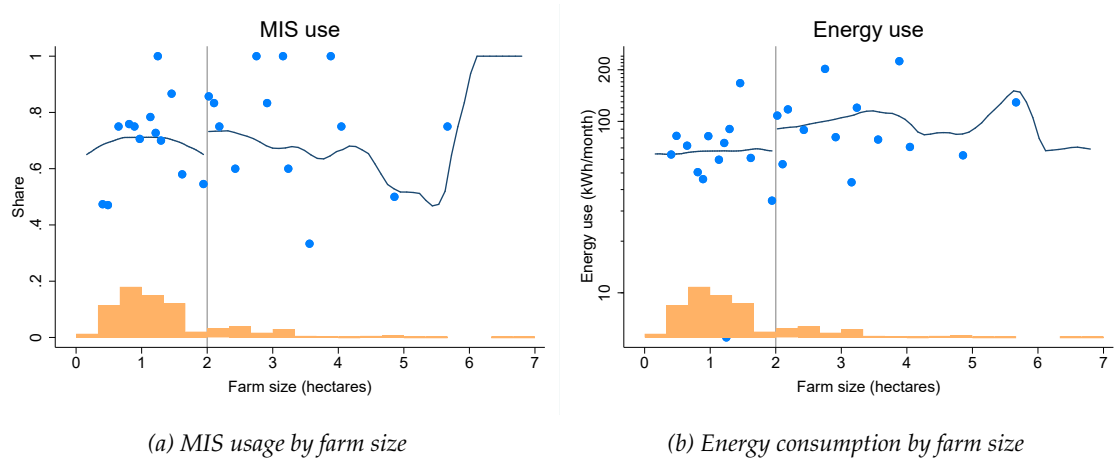


Figure 3: MIS and energy use above and below 2 hectares

Note: Figures overlay binscatter plots of MIS and energy usage according to farm size in blue on farm size histograms in orange. Panel (a) illustrates the failure of the first stage of a regression discontinuity design: the discontinuous increase in MIS price at 2 hectares is not associated with a discontinuous decrease in MIS usage. Panel (b) shows the reduced form impact of the MIS price discontinuity on energy use, and again finds no evidence of a discontinuity at 2 hectares. The histograms show that in our data, the number of farmers with farms near the 2 hectare cutoff is limited.

6 Discussion and conclusion

By combining direct meter-based measurements of groundwater pumping with comprehensive survey data, we provide a unique description of energy use patterns among smallholder farmers in a water-scarce region of Gujarat. We find two basic facts that are striking. First, energy use varies widely across farmers, a pattern that does not appear to be explained by other observed factors such as crop area. Second, micro-irrigation (MIS) users in our sample consume 30 to 40 percent more water than non-users of MIS. This contrasts with the conventional wisdom on MIS, which holds that water savings should translate to energy savings. We propose three hypotheses that might explain this basic fact: non-conservation (MIS does not actually reduce energy for the same crop), rebound effects, and selection bias.

We find little evidence for selection bias. The large difference in energy consumption by irrigation technology persists even after adjusting for farm size, well depth, pump power, other farm and water access descriptors, and a wide range of socioeconomic and demographic characteristics. Not only does a difference remain after adjusting for all these factors, its size changes little

across specifications. This stability is an indirect test of the influence of other unobserved factors: if the estimated magnitude is not sensitive to observed factors that are likely to be important, it suggests that it also might be robust to unobserved factors. It remains possible that unobserved factors might be confounding the relationship between MIS use and energy consumption, but it is difficult to think of such factors that could make up for such a large difference.

We also find little evidence for a rebound effect for MIS. If there is a rebound effect, the effect of MIS on total energy consumption should be larger than the effect of MIS on energy consumption per hectare cultivated. This is true in some of our specifications, but not all. We are unable to statistically distinguish the two effects in any specification, but this could mean either that there truly is no difference, or that there is in fact a difference but we do not have enough data to confirm it.

The absence of evidence for either selection bias or rebound effects suggests that there is a real possibility that MIS actually does not save energy in a real-world setting. This could be true if farmers are not operating their systems according to best practices, resulting in over-pressurization, or if they are using pumps that are too powerful for their MIS system. This conclusion would be consistent with at some evidence in the prior literature (Fishman et al., 2015).

Further research is needed to more definitively understand the effects of MIS on energy consumption and whether there is a rebound effect. An improved study design would track farmers over several years so that energy consumption could be compared for the same farmer before and after adoption of MIS. So far this has been difficult to implement since very few farmers adopt MIS in any given year. The best approach would be to conduct a controlled trial that randomly offers free MIS technology to farmers, but this would be expensive and it may still be difficult to reach a high take-up rate. Absent these kinds of studies, our analysis provides some suggestive evidence that the energy impacts of MIS under real-world conditions may be disappointing.

If it is true that MIS adoption actually increases energy consumption, this result would carry several implications for policy. First, government subsidies for MIS adoption may be a less attractive investment than previously thought. MIS likely provides multiple benefits to farmers, from higher yields to conserved groundwater, but to the extent that energy efficiency is part of the rationale for subsidies, this case may need to be reassessed. Second, it may be worthwhile to increase funding for training and extension services. If energy consumption goes up after adopting MIS, it may indicate that the systems are being installed without appropriate complementary technologies or are maintained incorrectly. Providing greater access to ongoing education and maintenance services might improve the chances that farmers enjoy the full benefits of the new technology.

Finally, there is a need for more detailed, quantitative monitoring of energy efficiency and resource consumption in real-world settings. Our relatively small endeavour has yielded one result with potentially unexpected consequences. More widespread and longer-term measurements of energy and water consumption are critical for policymakers and stakeholders to understand the reality of the situation on the ground and to guide appropriate policy responses.

References

- Belloni, A., Chernozhukov, V., Fernandez-Val, I., and Hansen, C. (2017). Program Evaluation and Causal Inference With High-Dimensional Data. *Econometrica*, 85(1):233–298.
- Elizabeth A. Stuart (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21.
- Fishman, R., Devineni, N., and Raman, S. (2015). Can improved agricultural water use efficiency save India's groundwater? *Environmental Research Letters*, 10(8):84022.
- Fowlie, M., Greenstone, M., and Wolfram, C. D. (2018). Do energy efficiency investments deliver? Evidence from the weatherization assistance program. *Becker Friedman Institute for Research in Economics Working Paper No. 2621817*.
- Gillingham, K., Rapson, D., and Wagner, G. (2015). The Rebound Effect and Energy Efficiency Policy. *Review of Environmental Economics and Policy*, 10(1).
- Kumar, D. S. and Palanisami, K. (2010). Impact of Drip Irrigation on Farming System: Evidence from Southern India. *Agricultural Economics Research Review*, 23:265–272.
- Raman, S. (2009). Micro Irrigation for Electricity Saving in Gujarat - A Potentiality Assessment. GGRCL.
- Sinha, I., Buttar, G., and Brar, A. (2017). Drip irrigation and fertigation improve economics, water and energy productivity of spring sunflower (*Helianthus annuus* L.) in Indian Punjab. *Agricultural Water Management*, 185:58–64.
- Surendran, U., Jayakumar, M., and Marimuthu, S. (2016). Low cost drip irrigation: Impact on sugarcane yield, water and energy saving in semiarid tropical agro ecosystem in India. *Science of The Total Environment*, 573:1430–1440.
- Venot, J.-P., Kuper, M., and Zwarteveen, M., editors (2017). *Drip irrigation for agriculture: untold stories of efficiency, innovation and development*. Earthscan from Routledge, Abingdon, UK, 1 edition.