

UNIVERSITY OF WAIKATO

**Hamilton
New Zealand**

**Can Smallholders Benefit from Complex Information?
Digital Water Quality Monitoring in Vietnamese Aquaculture**

Aya Suzuki, Vu Hoang Nam, Susan Olivia, Guenwoo Lee and Kohei Tanaka

Working Paper in Economics 6/26

October 2026

Corresponding Author

Aya Suzuki

Department of International Studies, Graduate School of Frontier Sciences
University of Tokyo

Email: ayazsk@k.u-tokyo.ac.jp

Vu Hoang Nam

Faculty of International Economics
Foreign Trade University

Susan Olivia

Department of Economics, Waikato Management School
University of Waikato

Guenwoo Lee

Japan International Research Center for Agricultural Sciences

Kohei Tanaka

Graduate School of Frontier Sciences
University of Tokyo

Abstract

Maintaining water quality is critical for aquaculture, yet smallholder farmers in developing countries often lack the tools to monitor it. Although earlier digital farming that provides “simple” information to farmers consistently yielded positive effects, when more detailed and personalized information are provided, the results have been mixed. Modern “aquatech” provides complex biophysical data that directly impacts productivity, and the effects likely depend on the capacity of farmers to process the “complex” information. We examine the impact of digital water-quality monitors and a farm-management app through a ten-month experiment involving shrimp farmers in Southern Vietnam. Using primary data from 304 farmers across two survey rounds, we find significant ITT and LATE treatment effects on reducing the disease outbreak as well as increasing revenue and profit, though the financial estimates are weaker evidence than the disease estimates: they sit at the margin of what our sample can detect. While these effects do not differ by human capital indicators of education, experience, and technical knowledge, the recorded data reveals strong association between education and stability of dissolved oxygen (DO) over crop cycles but not experience nor technical knowledge. Further, DO instability is associated with smaller shrimp size, and greater water stress is associated with lower harvest volume and smaller shrimp size. Our findings show that average farmers indeed benefit from even “complex” high-frequency digital information, and while the mechanism is not confirmed, the important role of education in water quality maintenance is suggested.

Keywords

Digital agriculture

Aquaculture

Smallholder farmers

JEL Codes

O12

O33

Q22

Acknowledgements

We are grateful to seminar audiences at SIPEC IE Seminar at Aoyama Gakuin University, Australasian Development Economics Workshop, Annual Conference of the Japan Society of International Development, Waikato University Economics Brownbag Seminar, and Jobs and Development Workshop for their helpful comments. Funding for this research was supported by Asian Development Bank Institute and Japan Society for the Promotion of Science (JSPS) Kakenhi 19H01476, 19KK0331, and 24K0248. Olivia acknowledges support from the JSPS International Research Fellowship.

1 Introduction

Maintaining water quality is the most critical factor for success in aquaculture (Asche et al., 2022). Poor water management increases the probability of disease outbreaks—the primary concern for fish farmers globally. As Lafferty et al. (2015) aptly note: “Aquaculture’s history is one of victories over diseases followed by new challenges (p476).” These outbreaks are rarely isolated; disease in a single pond can generate significant regional externalities as pathogens spread through shared waterways to neighboring farms (Suzuki et al., 2025). In response, there is growing optimism regarding the role of “Aquaculture 4.0”—the integration of digital devices and AI to stabilize production. Promising “aquatech” startups are now promoting digital tools to mitigate risk (Kruk et al., 2024). While the private sector has deployed tools ranging from automated feeders to acoustic health monitoring and digital traceability, the impact of these technologies on smallholder performance remains under-examined in the empirical literature.

This paper contributes to the burgeoning literature on digital farming in development economics. Earlier studies of digital technology in agriculture found mostly positive results. For example, diffusion of mobile phones and better access to price information improved farmers’ selling behavior and financial outcomes by reducing information asymmetries (e.g., Jensen, 2007; Aker, 2011; Aker and Fafchamps, 2015; Nakasone et al., 2014). Further, digital delivery of agricultural advices by such means as text messages, interactive video responses, or video clips also became popular (e.g., Fu and Akter, 2016; Larochelle et al., 2019; Van Campenhout et al., 2021), and in conducting the meta-analysis of six studies, Fabregas et al. (2025) calculates that digital farming improves crop yield by up to 4%.

More recent experiments focus on providing more detailed and personalized information to farmers, finding mixed results. These experiments are catered to farmers’ individual situations, such as climatic conditions, soil quality, or crop health. Introduction of personalized extension advice on nutrients management was found to increase yield and profit in Nigeria (Arouna et al., 2021; Oyinbo et al., 2022). The use of voice-based ICT advice improved demand for information and input use but had no impacts on increasing yields nor profitability in India (Cole and Fernando, 2021). Provision of personalized soil health cards did not work either in India, and the possible reason attributed is the complexity of information provided (Cole and Sharma, 2018). While provision of more detailed information seems to benefit farmers, inconsistent results found in the existing studies call for the need to explore this mechanism further.

In digital aquaculture, water quality monitors provide a fundamentally different type of information. Unlike market price information, which is easily interpretable and action-able, biophysical parameters such as pH levels and salinity are technically complex and change constantly. Improving productivity with this kind of information requires farmers to correctly interpret real-time data and adapt their management decisions to their specific pond conditions - a cognitively demanding task that goes beyond simply possessing technical knowledge.

This suggests a natural hypothesis. In labor economics, technological change is skill-biased in a specific sense: capital substitutes for routine tasks while complementing workers engaged in non-routine problem-solving (Autor et al., 2003; Acemoglu and Autor, 2011). Acting on real-time dissolved oxygen or pH readings is a non-routine task in exactly this sense, requiring farmers to diagnose deviations, identify likely causes, and tailor corrective actions to their own ponds. Selective attention theory points the same way, predicting that the binding constraint is not domain-specific knowledge but the general capacity to translate unfamiliar information into novel decisions (Schwartzstein, 2014; Hörner et al., 2022; Naeher and Schündeln, 2022). Together these frameworks predict that the returns to digital aquaculture should accrue disproportionately to farmers with more general human capital, with the corollary that such tools would widen rather than narrow rural inequality.

We examine this dynamic by evaluating the introduction of digital water quality measurement devices in the shrimp farming sector of Southern Vietnam. Shrimp is a globally traded commodity, and major producers—including Vietnam, India, Indonesia, and Ecuador—face a highly competitive market where quality and consistency are paramount. In collaboration with an Indonesian startup JALA Tech and the Vietnam Sustainable Shrimp Alliance (VSSA), we conducted a ten-month intervention targeting shrimp farmers from April 2023 to January 2024. Treated farmers were provided free access to the JALA Tech “Baruno” device, which is a water quality meter linked to their app and a premium app subscription to record water quality parameters as well as farming practices. To ensure data fidelity, field assistants assisted treated farmers with recording of water quality parameters by going around villages with the Baruno. We estimate the impact of this intervention on the incidence of disease, sales revenue, production costs, and profits.

Based on ANCOVA models of intention-to-treat (ITT) and local average treatment effects (LATE) with treatment assignment as instrumental variables for actual take-up, we find that the intervention had significant impact on reducing the probability of disease outbreaks and increasing their revenues and profits and reject the hypothesis that the effect is skill-biased. Neither schooling, nor technical knowledge of water chemistry, nor farming experience moderates the treatment effect; the gains are shared across the education distribution. This result is robust to DID models of ITT and LATE as well as PSM-DID. The high-frequency records from the devices nonetheless suggest that education operates through a narrower channel, namely the day-to-day stability with which farmers maintain dissolved oxygen (DO). Greater water stress is associated with less harvested volume and smaller shrimp size, and DO instability is linked to smaller shrimp size. Overall, we find clear effects of visualizing water quality parameters for shrimp farmers. The effects do not depend on human capital, but the recorded data suggests the importance of education over technical knowledge or experience.

Our contribution is twofold. First, we provide the first experiment on digital technology adoption in smallholder aquaculture. Aquaculture production involves many parameters that are critical for farming outcomes yet unobservable to farmers without specialized equipment—water temperature, dissolved oxygen, pH, and salinity fluctuate continuously and interact in complex ways. With the advancement of affordable sensor technology, these parameters can now be monitored in real time. Our results show that average farmers can benefit from complex high-frequency data and utilize this information to reduce the incidence of disease outbreak

and to achieve higher financial outcomes. Note that our sample farmers do not belong to the most marginalized group of farmers in the region as they practice super-intensive farming, which requires sophisticated aquaculture technologies. These farmers can turn the complex information into positive farming outcomes when given the chances.

Second, we present suggestive evidence on the linkages between water quality, farmers' behavior, and shrimp farming outcomes based on the high-frequency field data recorded on the digital devices. Although these should not be taken as causal effects, we observe that years of schooling is associated with better maintenance of water quality but technical knowledge nor experience are not. This finding complements Laajaj and Macours (2026), who show that multidimensional learning capacity, rather than knowledge of any single input, constrains technology adoption in agriculture and extends their insight to the digital farming context. Better water maintenance, indicated as less incidence of water stress and stable level of DO, is associated with greater harvest volume and larger shrimp size. While scientific relations between water quality and shrimp farming outcomes are public knowledge in scientific literature, these are normally explored using model ponds operated by scientists under controlled environment. The novelty of our finding is that it is based on farmers' day-to-day activities over different cycles among different types of farmers in various environment, thus more generalizable and practical for policy considerations.

The remainder of the paper is organized as follows. Section 2 describes the study context and experimental design. Section 3 details the estimation strategy. Section 4 presents the results and robustness checks, and Section 5 concludes with policy implications.


2 Study Context

2.1 Vietnam Shrimp Sector

The shrimp sector in Southern Vietnam proliferated following a 2000 government decree allowing the conversion of low-altitude rice paddies into aquaculture ponds in the Mekong Delta (Tran et al., 2013). High salinity in local waterways, which previously suppressed rice yields, proved advantageous for shrimp cultivation, leading to a steady increase in returns per hectare (Figure 1). Our study site, Ca Mau Province, is Vietnam's largest shrimp producer. Many processing companies exist, providing large employment opportunities for the local economy. The industry relies heavily on smallholders who supply local collectors, traders, and large-scale processing factories. While some processors maintain integrated production, from our interviews, even the largest company has capacity to source only 20% of their raw material requirements, underscoring the sector's systemic dependence on small-scale producers.

Shrimp production systems range from extensive to super-intensive methods, categorized by stocking density (e.g., Joffre et al., 2018; Nguyen et al., 2019). We focus on intensive (70–150 pcs/m²) and super-intensive (250+ pcs/m³) systems for our study as maintenance of water quality is much more crucial than in other farming methods. In these high-density environments, water quality maintenance is a binding constraint on productivity, requiring precise management of dissolved oxygen, nutrient loads, and soil contamination via plastic pond liners. However, most shrimp farmers were relying on simple test kits to measure water quality at the time of survey while digital water quality meters were available in the Vietnam market.

Figure 1: Change in the gross output of product per hectare land area



2.2 Sampling and Intervention

We obtained a list of shrimp farmers from local government office in two selected districts in Ca Mau Province, which are characterized by high concentrations of super-intensive producers in 2022. From these two districts, three communes were chosen based on their higher density of shrimp farmers. Our intention was to interview all shrimp farmers in these three communes on the list as we are interested in social network among farmers. However, due to sector exits of many shrimp farmers, when the survey began, we found that the list was not up to date. Thus, our enumerators visited village leaders in these three communes first to obtain the updated list of farmers in each village and interviewed all the intensive and super-intensive farmers in all three communes. Our initial 2023 baseline surveyed 304 farmers in total. Upon follow-up in 2024, 210 farmers remained in the sample. The 30.9% attrition rate, while higher than anticipated, was primarily driven by sector exits; many farmers ceased operations due to rising input costs and stagnant global prices, migrating to industrial employment in neighboring provinces. Our survey instrument captured socio-economic characteristics, production and sales data, technical knowledge, and social networks.

While we initially randomly selected 130 farmers from the original list for our intervention, only 43 of them were found in the field. Further, during our intervention of 10 months, we had many leavers over the course of time, which prompted us to assign additional farmers for our treatment based on advice by our local counterparts. This amounted to 42, totaling 85 farmers being invited for treatment. As this latter group of farmers may introduce bias, we use the 43 farmers who were in the original list as the invited farmers. Among them, we consider those who joined our intervention from the beginning as treated farmers since using those who stayed until the end of 10 months may introduce other bias. This left us with 20 farmers who are considered as having participated in the treatment. We conduct robustness checks on these small numbers and use ITT estimation for our main analyses.

In April 2023, we launched the intervention via a workshop held in the center of Ca Mau province, inviting farmers, CEO and technical experts from JALA Tech, the VSSA, and our research team. Following the workshop, treated farmers received access to the JALA “Baruno” water quality measurement device, which is linked to the JALA application and automatically stores various water quality parameters once it is measured. Financial and stock management data could also be recorded digitally at farmers’ discretion.

To ensure data quality, we collaborated with the VSSA to deploy 13 field assistants who visited treated farmers’ ponds twice daily (morning and afternoon) to record water parameters if farmers agreed. Note that these assistants were not aquacultural experts, so no technical advice was given when they visited farmers. The assistants were purely technical operators of the water quality meters. While assistants ensured consistent biophysical data collection, the frequency of management record-keeping within the app remained at the farmers’ discretion. Field assistants had meetings with VSSA coordinators to clarify any questions they had. Our intervention was conducted between April 2023 and January 2024.

2.3 Summary Statistics

Table 1 and Table A1 show the descriptive statistics for our farmers based on the treatment assignment status and take-up status. We find that most of our shrimp farmers practice super intensive farming methods. Most of our farmers are male, which is typically the case for many high-value crops, and the average age is about 48 years old. Over 70% of farmers use radio or television to obtain information on shrimp farming while around 50% of them use internet to obtain the same information. It is rare for farmers to have any contracts with buyers (around 10-14%). About 80 to 90% of the farmers have received some kinds of training on good shrimp production methods. On the knowledge about prohibited elements, while about 60% of the farmers answered that they know which elements are prohibited to use (i.e., such as antibiotics), only 6-9% of farmers were able to mention names of specific products which contain these prohibited elements.

Some variables are statistically significantly different between the treated and control groups. Particularly, years of education completed is slightly longer for those who are assigned treatment and for those who participated in the treatment than the control group. The same tendency is observed for the cooperative membership. Some variables that indicate knowledge level are also higher for the treatment group while the frequency of meeting extension officers is higher for control group than the treated group. We consider these differences in the following analyses.

Table 1: Descriptive Statistics: Invited vs. Control at Baseline (Estimation Sample)

	Control			Invited			Diff.	<i>p</i> -value
	Mean	Median	SD	Mean	Median	SD		
<i>N</i>	219			43				
Super-intensive (=1)	0.991	1.000	0.096	1.000	1.000	0.000	0.009	0.529
Experience (yrs)	13.207	13.000	5.338	12.581	12.000	4.707	-0.626	0.475
Married (=1)	0.977	1.000	0.150	1.000	1.000	0.000	0.023	0.317
gender (=1: male)	0.889	1.000	0.314	1.000	1.000	0.000	0.111	0.022
Age	48.774	48.000	10.434	49.884	49.000	9.447	1.110	0.518
Education (yrs)	7.977	8.000	2.364	8.628	9.000	1.940	0.651	0.091
Residence in commune (yrs)	44.161	46.000	14.350	48.860	49.000	10.535	4.699	0.042
Household size	3.862	4.000	1.713	3.698	4.000	1.301	-0.164	0.553
No. of ponds	2.051	2.000	1.288	1.907	2.000	1.151	-0.144	0.497
Cooperative member	0.198	0.000	0.400	0.465	0.000	0.505	0.267	0.000
Farmer group member	0.041	0.000	0.200	0.000	0.000	0.000	-0.041	0.175
Owns smartphone	0.912	1.000	0.283	0.930	1.000	0.258	0.018	0.703
Info: TV/radio	0.714	1.000	0.453	0.814	1.000	0.394	0.100	0.180
Info: internet	0.447	0.000	0.498	0.558	1.000	0.502	0.111	0.183
Sales contract	0.129	0.000	0.336	0.163	0.000	0.374	0.034	0.555
Received training	0.811	1.000	0.392	0.907	1.000	0.294	0.096	0.130
Aware of prohibited inputs	0.032	0.000	0.177	0.140	0.000	0.351	0.107	0.003
Knows prohibited elements	0.553	1.000	0.498	0.791	1.000	0.412	0.238	0.004
Knows prohibited products (specific)	0.041	0.000	0.200	0.163	0.000	0.374	0.121	0.002
Extension officer (freq.)	5.783	6.000	1.172	5.070	5.000	1.387	-0.714	0.000
Postlarvae stocked (no.)	4.4e+05	3.0e+05	4.2e+05	5.2e+05	4.0e+05	5.2e+05	8.2e+04	0.263
Total pond area (ha)	0.550	0.240	0.958	0.929	0.200	2.049	0.379	0.061
ln(Revenue)	6.303	6.356	1.135	6.286	6.232	0.782	-0.018	0.925
ln(Cost)	5.843	5.874	0.682	5.847	5.844	0.777	0.004	0.972
ln(Profit)	5.494	5.535	1.344	5.299	5.340	1.262	-0.194	0.428
ln(Revenue/ha)	7.585	7.874	1.415	7.477	7.770	1.111	-0.109	0.647
ln(Cost/ha)	7.074	7.298	0.893	6.910	7.067	0.928	-0.165	0.273
ln(Profit/ha)	6.883	7.081	1.443	6.721	7.030	1.322	-0.162	0.537
Disease rate (avg)	0.206	0.000	0.395	0.368	0.000	0.464	0.163	0.017
Disease (total)	0.309	0.000	0.639	0.442	0.000	0.548	0.133	0.203
Disease (any, =1)	0.221	0.000	0.416	0.419	0.000	0.499	0.197	0.006

We assess whether our rather high attrition rate is systematic in Table A8. Attrition does not differ significantly between the invited and control groups (Panel A; 27.9% against 33.8%, $p = 0.46$), which is the comparison that matters for internal validity. Being invited or participated does not predict attrition (Panel C). Baseline characteristics do not jointly predict attrition either (Panel D; $F = 1.44$, $p = 0.10$), and dropping the production-system indicator (i.e., whether super-intensive or intensive) leaves them jointly insignificant ($F = 1.22$, $p = 0.24$). A handful of individual characteristics differ between stayers and attriters (Panel B), but not jointly. The composition of attriters likewise does not differ significantly between arms (interaction $F = 1.02$, $p = 0.44$). Attrition therefore appears close to ignorable; even so, we report Lee bounds in Table A9 as a conservative check that does not rely on that assumption.

3 Estimation Strategy

We use analysis of covariance (ANCOVA) as our primary estimator:

$$Y_{i1} = \alpha + \beta Treat_i + \lambda Y_{i0} + X_i' \theta + \mu_c + \varepsilon_i \quad (1)$$

where Y_{i1} is farmer i 's financial performance or the probability of disease outbreak at the endline, $Treat_i$ denotes participation to the JALA intervention, Y_{i0} is farmer i 's baseline outcome value, X_i is a vector of baseline farmer characteristics, μ_c is the commune fixed effects, and ε_{it} is an error term. Another natural specification for a two-period panel is differencing approach, and we use this for robustness checks (see equation 2)

$$Y_{it} = \alpha + \beta_1 Post_t + \beta_2 Treat_i \times Post_t + \delta_i + \gamma_t + \varepsilon_{it} \quad (2)$$

where Y_{it} is farmer i 's outcome at time t , $Post_t$ is a dummy for the endline period, δ_i is the farmer fixed effects, and γ_t is the time fixed effects. This is the special case $\lambda = 1$ of the ANCOVA model.

As McKenzie (2012) shows, the relative precision of the two estimators depends on the autocorrelation ρ of the outcome, with $\text{Var}(\hat{\beta}_{ANCOVA})/\text{Var}(\hat{\beta}_{DID}) = (1 + \rho)/2$. In our data ρ is close to zero for the disease outcomes and between 0.04 and 0.31 for the financial ones, so differencing roughly doubles the sampling variance.

Precision is not the only consideration. Because the invited group happens to have a higher baseline disease rate than the control group (Table 1), and because $\rho \approx 0$ means that baseline disease carries almost no information about endline disease, differencing mechanically imports that imbalance as an apparent treatment effect. Conditioning on Y_{i0} rather than differencing it out avoids this. Under random assignment, classical measurement error in Y_{i0} attenuates λ but leaves β unbiased, so the lagged dependent variable is included for identification rather than for precision alone.

The estimation sample is restricted to the 43 farmers drawn from the original randomization list plus the 219 never-invited farmers. The 42 farmers added mid-intervention are excluded to avoid bias. For the LATE we define take-up as participation from the start of the intervention (20 of the 43), rather than as staying until the end of our intervention (49 farmers), because the latter conditions on ten months of survival in the program and builds selection into the treatment variable itself.

To allow for spillovers from treated neighbours, all specifications also control for the number of *other* invited farmers in the same village, together with the number of farmers surveyed there so that the coefficient compares villages of similar size. Village fixed effects cannot be used alongside this measure, since they absorb the village total and leave only own assignment. With 27 villages, conventional cluster-robust inference over-rejects. We therefore report wild cluster bootstrap p-values and confidence intervals throughout (Rademacher weights, 9,999 replications), with conventional cluster-robust standard errors in parentheses for reference.

A primary identification challenge is non-random self-selection into the treatment take-up. Farmers who opted into the intensive ten-month monitoring regime may differ systematically from those who declined. To mitigate this endogeneity, we first employ Local Average Treatment Effects (LATE) using the treatment assignment as an instrument for treatment participation. We use treatment assignment (*Invited*) as an instrumental variable for actual JALA device take-up. For this instrument to be valid, two conditions must hold. The first condition, relevance, is satisfied: the Kleibergen–Paap first-stage statistic (Table 3) confirms a strong relationship between assignment and take-up. The second condition, the exclusion restriction, requires that treatment assignment affects farming outcomes only through actual use of the JALA device, and not through any other channel.

We argue the exclusion restriction is plausible for three reasons. First, while the 43 invited farmers consist of only a part of our original random assignment, making our invitation not perfectly random, the deviations from randomization were driven by administrative factors, i.e., an outdated farmer registry, rather than by farmer characteristics or potential outcomes. Consistent with this, our attrition analysis (Table A8) shows that the invited and control groups left the sample at statistically indistinguishable rates ($p = 0.46$), suggesting that the assignment process is not systematically related to farmer productivity. Second, the exclusion restriction would be violated if treatment assignment affected outcomes through channels other than actual device use; for example, if merely being invited conveyed useful information or altered farmer behavior. However, the invitation itself carried no agronomic content; it was simply an offer to participate in the monitoring program. Farmers who were assigned to treatment but declined to use the device (*never – takers*) received no water quality data and no access to the JALA app’s management tools. In other words, for never-takers, the experience of being assigned was observationally equivalent to being in the control group, leaving no plausible channel through which assignment alone could have affected their farming outcomes. Third, the introductory workshop held at the start of the intervention was open only to assigned farmers; however, the workshop focused on device operation rather than

agronomic advice, limiting the scope for information spillovers to affect outcomes independently of device use. Nonetheless, we acknowledge that the exclusion restriction is an assumption that cannot be directly tested, and interpret our LATE estimates with this caveat in mind.

To test robustness, we also conduct Propensity Score Matching with Difference-in-Differences (PSM-DID). By restricting the analysis to a balanced sample of treated and control units with common support, we reduce bias stemming from observable pre-treatment differences. We further assess sensitivity to unobserved confounding using the Oster (2019) bounding approach and the robustness values of Cinelli and Hazlett (2020), and check that no single village drives the results by re-estimating the model dropping one village at a time. Although our attrition analysis confirms the attrition is not significant, we also report Lee (2009) trimming bounds for prudence.

We further conduct heterogeneity analyses based on the education level, technical knowledge possessed, and experience of shrimp farming to examine whether the impacts of JALA interventions depend on human capital of individual farmers.

4 Results

4.1 Average effects

Table 2 reports the intention-to-treat estimates. Panel A shows that access to the device reduced disease substantially. The average probability of an outbreak across a farmer's ponds falls by 0.200 ($p = 0.006$) against a control-group endline mean of 0.247, the total number of outbreaks across a farmer's ponds falls by 0.404 ($p = 0.000$) against 0.441, and the probability of any outbreak falls by 0.282 ($p = 0.000$) against 0.324. These are large effects: relative to the control group they represent reductions of roughly 80%, 90% and 86% respectively.

Panel B shows the financial outcomes. Log revenue rises by 0.536 ($p = 0.054$) and log profit by 0.627 ($p = 0.089$); log costs also rise, by 0.422 ($p = 0.030$). The revenue and profit gains are consistent with the disease results, i.e., fewer outbreaks leave a larger share of the stocked shrimp to reach harvest, while the parallel rise in costs indicates that sustaining the larger surviving stock through the cycle was not costless, and revenue nonetheless rises by enough to leave profit higher. The corresponding per-hectare outcomes are smaller and not significant, which is expected given that the intervention should raise output per pond rather than expand pond area, and that dividing by area adds measurement error.

Table 2: Effect of Access to the Water Quality Device: Intention to Treat (ANCOVA)*Panel A: Disease Outcomes*

	Disease (avg)	Disease (total)	Disease (any)
Assigned to treatment (ITT)	-0.200*** (0.065)	-0.404*** (0.124)	-0.282*** (0.080)
Other invited farmers in village	0.002 (0.013)	-0.013 (0.028)	-0.003 (0.017)
<i>N</i>	176	176	176
Bootstrap <i>p</i> (treatment)	0.006	0.000	0.000
Bootstrap CI lower	-0.343	-0.716	-0.462
Bootstrap CI upper	-0.060	-0.151	-0.120
Bootstrap <i>p</i> (spillover)	0.913	0.746	0.891
Control mean (endline)	0.247	0.441	0.324

Notes: ANCOVA estimates of the intention-to-treat effect. The dependent variable is the 2024 outcome; the 2023 baseline outcome is included as a regressor. Sample restricted to farmers drawn from the original randomization list ($n=43$) and never-invited farmers ($n=219$); farmers added later to replace attriters are excluded because their exposure was shorter. All specifications include commune fixed effects and baseline farmer characteristics (age, years of shrimp farming experience, years of schooling, years of residence in the commune, postlarvae stocked, and total pond area) and institutional-connection controls (cooperative membership, training, input knowledge). To allow for spillovers from treated neighbours, the number of *other* invited farmers in the same village is included, together with the number of farmers surveyed there, so that the coefficient compares villages of similar size. The count excludes the farmer's own assignment, and is computed over all invited farmers, including those added mid-intervention who are excluded from the estimation sample but could still transmit information. Village fixed effects cannot be used alongside this measure: they absorb the village total, leaving only own assignment and making the two collinear. Parentheses report conventional cluster-robust standard errors; *p*-values and confidence intervals are from a wild cluster bootstrap (village clusters, 9,999 replications), which corrects the over-rejection of conventional inference with 27 clusters. Significance stars follow the bootstrap *p*-values (treatment and spillover rows).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Financial Outcomes

	ln(Rev)	ln(Cost)	ln(Profit)	ln(Rev/ha)	ln(Cost/ha)	ln(Profit/ha)
Assigned to treatment (ITT)	0.536* (0.226)	0.422** (0.151)	0.627* (0.299)	0.384 (0.237)	0.249 (0.165)	0.443 (0.289)
Other invited farmers in village	0.061 (0.024)	0.061** (0.015)	0.023 (0.042)	0.065* (0.028)	0.058** (0.021)	0.017 (0.033)
<i>N</i>	169	176	141	169	176	141
Bootstrap <i>p</i> (treatment)	0.054	0.030	0.089	0.140	0.187	0.184
Bootstrap CI lower	-0.017	0.047	-0.145	-0.168	-0.124	-0.311
Bootstrap CI upper	1.018	0.735	1.285	0.924	0.632	1.102
Bootstrap <i>p</i> (spillover)	0.193	0.041	0.715	0.078	0.034	0.648
Control mean (endline)	6.170	5.688	5.316	7.684	7.190	6.825

Notes: See notes to the disease panel. ANCOVA intention-to-treat estimates with commune fixed effects, baseline outcome, farmer characteristics and institutional-connection controls. Wild cluster bootstrap inference (village clusters, 9,999 replications).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effect of Using the Water Quality Device: Local Average Treatment Effect (ANCOVA)

Panel A: Disease Outcomes

	Disease (avg)	Disease (total)	Disease (any)
Participated from start (LATE)	-0.422** (0.180)	-0.862** (0.370)	-0.597** (0.238)
Other invited farmers in village	0.002 (0.013)	-0.012 (0.027)	-0.003 (0.017)
<i>N</i>	176	176	176
Weak-ID F (KP rk Wald)	15.44	15.31	15.52
Bootstrap <i>p</i> (treatment)	0.032	0.028	0.019
Bootstrap <i>p</i> (spillover)	0.916	0.725	0.897
Control mean (endline)	0.247	0.441	0.324

Notes: ANCOVA local average treatment effects. The endogenous variable is participation from the start of the intervention, instrumented by assignment from the original randomization list. Take-up among the assigned is 20/43. Participation from the start is used rather than participation through the end, because the latter conditions on surviving the ten-month intervention. Weak-identification F is the Kleibergen-Paap rk Wald statistic; the Stock-Yogo 10% critical value is 16.38. Wild cluster bootstrap inference (village clusters, 9,999 replications); significance stars follow the bootstrap *p*-values. Because the LATE rescales the ITT by a take-up rate below one-half, its magnitude should be read alongside the ITT.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Financial Outcomes

	ln(Rev)	ln(Cost)	ln(Profit)	ln(Rev/ha)	ln(Cost/ha)	ln(Profit/ha)
Participated from start (LATE)	1.067* (0.517)	0.895* (0.388)	1.268* (0.592)	0.766 (0.509)	0.531 (0.380)	0.898 (0.566)
Other invited farmers in village	0.055 (0.027)	0.062** (0.016)	0.013 (0.043)	0.061* (0.027)	0.059** (0.019)	0.010 (0.033)
<i>N</i>	169	176	141	169	176	141
Weak-ID F (KP rk Wald)	16.53	15.50	14.74	16.53	15.16	14.64
Bootstrap <i>p</i> (treatment)	0.078	0.052	0.079	0.189	0.229	0.180
Bootstrap <i>p</i> (spillover)	0.282	0.040	0.831	0.097	0.026	0.821
Control mean (endline)	6.170	5.688	5.316	7.684	7.190	6.825

Notes: ANCOVA local average treatment effects on financial outcomes. The endogenous variable is participation from the start of the intervention, instrumented by assignment from the original randomization list; take-up is 20 of 43. Weak-identification F is the Kleibergen-Paap rk Wald statistic. Wild cluster bootstrap *p*-values (village clusters, 9,999 replications); significance stars follow the bootstrap *p*-values. With only twenty participants these estimates are imprecise and are reported for completeness.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3 instruments participation with assignment. Take-up was 20 of the 43 assigned farmers, and the Kleibergen-Paap statistic is around 15 in all specifications, above the Stock-Yogo 10% threshold. The disease estimates roughly double, as expected under one-sided non-compliance: -0.422 for the average rate ($p = 0.032$), -0.862 for the count ($p = 0.028$) and -0.597 for incidence ($p = 0.019$). Log revenue rises by 1.067 ($p = 0.078$), log profit by 1.268 ($p = 0.079$) and log costs by 0.895 ($p = 0.052$); under the bootstrap the financial estimates are significant at the 10% level.

These magnitudes should not be read literally. As our non-compliance is one-sided, the Wald estimator scales the intention-to-treat effect by the take-up rate of 0.465 (Angrist and Pischke, 2009), and the resulting disease estimates exceed the counterfactual disease rate of any group observed in the data. Compliers had a baseline disease rate of 0.348, so a reduction of 0.422 would imply a negative post-treatment rate while this is a probability that is bounded within 0 and 1. This indicates either that the exclusion restriction does not hold exactly or simply that twenty participants is too small to detect the effect. We report the local average treatment effect for completeness, but rely on the intention-to-treat estimates throughout.


4.2 Spillovers

All specifications control for the number of *other* invited farmers in the same village, which lets us ask whether treated neighbours affected untreated ones. For the disease outcomes the coefficient is essentially zero ($p > 0.7$), so the disease results are not contaminated by spillovers onto the control group. For the financial outcomes the coefficient is positive and significant for costs (0.061, $p = 0.041$; and 0.058 per hectare, $p = 0.034$), but not for revenue or profit. We are reluctant to read this as a genuine spillover. As village fixed effects cannot be included alongside a village-level count without collinearity, the coefficient may simply pick up unobserved village characteristics correlated with both invitation and input use. We report it for transparency and treat it as exploratory.

4.3 Heterogeneity by human capital

Our motivating hypothesis was that the returns to complex biophysical information would be skill-biased. Table 4 tests this directly, interacting treatment assignment with three moderators: more than nine years of schooling, technical knowledge of dissolved oxygen, and shrimp farming experience above the median. In Vietnam primary education lasts five years, lower secondary education ends at nine years, and higher secondary education ends at twelve. The distribution of schooling in our sample shows a sharp decline at nine years and no farmer has more than twelve (Figure 2), which motivates the threshold. For technical knowledge, we asked farmers five questions on dissolved oxygen—the appropriate level, the level at which shrimp tend to die, when it tends to increase, how it can be raised, and the signs of a low reading—and formed an index, splitting at the median.

Figure 2: Distribution of Education Level



None of the eighteen interaction terms is statistically significant, and the point estimates are small relative to the average effects. For disease (any), the education interaction is -0.018 ($p = 0.894$) against an average effect of -0.282 (Table 2). Table A4 confirms that this is not an artefact of the estimator: the difference-in-differences and ANCOVA interactions agree in showing no moderation by these human capital indicators. We also verified that the result does not depend on the threshold, using a median split, a continuous measure and a four-category specification.

This rejects our hypothesis raised in the beginning, and it is worth being explicit that it overturns what we expected. If the mechanism were domain-specific knowledge, technical knowledge of water quality or accumulated experience should have amplified the treatment effect. If it were general cognitive capacity, as the skill-biased technological change literature (Autor et al., 2003; Acemoglu and Autor, 2011) and selective attention theory (Schwartzstein, 2014) suggest, formal education should have mattered. Neither pattern appears. Access to high-frequency water quality data benefited farmers across the education distribution in our case. We return in Section 4.5 to why the exploratory sensor data nonetheless points to a role for education.

Table 4: Heterogeneity by Schooling, Technical Knowledge and Experience (ANCOVA)

	(1) Education > 9 years	(2) Tech. knowledge DO score ≥ 3	(3) Experience ≥ 14 years	<i>N</i>
<i>Panel A: Disease outcomes</i>				
Disease (avg)	-0.008 [0.943]	-0.033 [0.819]	0.054 [0.509]	176
Disease (total)	-0.075 [0.693]	-0.111 [0.587]	0.027 [0.810]	176
Disease (any)	-0.018 [0.894]	-0.025 [0.869]	0.048 [0.571]	176
<i>Panel B: Financial outcomes</i>				
ln(Revenue)	-0.205 [0.797]	0.117 [0.743]	-0.154 [0.723]	169
ln(Cost)	0.146 [0.722]	0.083 [0.659]	0.107 [0.754]	176
ln(Profit)	0.445 [0.581]	-0.175 [0.758]	-0.056 [0.922]	141

Notes: Each cell reports the coefficient on the interaction between treatment assignment and the moderator named in the column heading, from an ANCOVA regression that also includes the baseline outcome, the main effect of the moderator, commune fixed effects and the full set of baseline controls. Each row of Panel A and Panel B therefore comes from three separate regressions, one per moderator. Moderators are split at the sample median. Wild cluster bootstrap *p*-values in brackets (village clusters, 9,999 replications). The sample is farmers from the original randomization list (*n*=43) and never-invited farmers (*n*=219). None of the three moderators produces a systematic interaction, so the average effects reported earlier are not driven by a particular subgroup.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.4 Robustness

The results are robust to estimating equation (2) by difference-in-differences (Table A2, Table A3) and to DID propensity score matching (Table A6), although as discussed in Section 3 the differenced estimates are less precise and, for the disease outcomes, absorb the baseline imbalance between arms. For the matching estimates we present the propensity score estimation in Table A5, the score densities in Figure A1 and the love plot in Figure A2; covariate balance improves markedly after matching, though one-to-one nearest neighbour matching reduces the sample.

Because assignment was not fully random, we next ask how strong unobserved selection would have to be to explain the results away. The Oster (2019) statistic δ reports the degree of selection on unobservables, relative to selection on the observed covariates, that would drive the coefficient to zero. Setting $R_{max} = 1.3R^2$, we obtain $\delta = 5.44$ for the average disease rate, 7.55 for the count, 4.12 for incidence, 4.64 for log revenue and 6.75 for log profit. Since $|\delta| > 1$ in every case, unobserved confounders would have to be four to eight times as important as the full set of observed characteristics—which includes age, experience, schooling, pond area, stocking, cooperative membership and training—to overturn the estimates.

The Cinelli and Hazlett (2020) framework gives a complementary result. The Robustness Value computed for the disease outcomes ($RV_{q=1}$) is 16.85% for the average rate, 18.93% for the count and 19.56% for incidence. This means that an unobserved confounder would have to explain more than roughly a fifth of the residual variance in *both* treatment assignment and the outcome to drive our result estimate to zero. The corresponding values for retaining statistical significance of our result at the 5% level ($RV_{q=1, \alpha=0.05}$) are lower, at 2.68%, 5.12% and 5.87%, so statistical significance is more fragile than the point estimate. However, benchmarking against cooperative membership, which is the strongest observed predictor of assignment in our model, even a confounder of comparable strength is far from being sufficient to overturn the results.

Finally, we re-estimate the model dropping one village at a time to check that no single cluster drives the findings. The full-sample coefficients of disease outcomes -0.200 , -0.404 and -0.282 (Table 2) move within $[-0.226, -0.161]$, $[-0.454, -0.318]$ and $[-0.317, -0.243]$ respectively. The sign and approximate magnitude are preserved throughout, and no village accounts for more than a fifth of the estimated effect.

Our attrition analysis shows that the attrition does not differ significantly by treatment assignment (27.9% among invitees against 33.8% among controls, $p = 0.455$), and the composition of attriters does differ across arms (Table A8). Nonetheless, for prudence we report Lee bounds in Table A9 (Lee, 2009). Because attrition is higher in the control group, the bounds are obtained by trimming between 8 and 10% of the retained control observations, which equalizes response rates across arms. All five outcomes have bounds that exclude zero, and so do the confidence intervals of Imbens and Manski (2004): $[-0.42, -0.14]$ for disease incidence and $[0.15, 1.16]$ for log revenue. The covariate-adjusted ANCOVA estimates lie inside these bounds. Worst-case selection into response can therefore attenuate the estimates, but it cannot overturn their sign or significance.

Table A10 reports what our design was able to detect. We deliberately do not report power computed at the observed effect size as such a quantity is a monotone transformation of the p -value and therefore adds nothing to the estimates and confidence intervals already reported (Hoenig and Heisey, 2001). Instead we report the minimum detectable effect (MDE) at 80% power, which depends only on the achieved standard error and not on the estimate, and express it relative to the control group so that its economic size is readable.

The design was only able to detect large effects. For the average disease rate, the MDE is 0.182, or 74% of the control mean; for disease incidence it is 0.224, or 69%. For the financial outcomes the MDE corresponds to an 88% increase in revenue and a 131% increase in profit. Two things follow. First, our disease estimates are significant not because the design is well powered in any absolute sense, but because the effects are themselves very large, i.e., the reductions of roughly 80 to 90% of the control mean. Second, and more importantly, the design could not have detected moderate effects. For example, an intervention that reduced disease by a third of the control mean or raised revenue by 30% would not have been detected. Our null results should therefore be read as uninformative about moderate effects rather than as evidence of their absence. This applies with particular force to the heterogeneity analysis in Section 4.3: the absence of significant interactions rules out large differences between education groups, but is entirely consistent with moderate ones. Because take-up is only 20 of 43 assigned farmers, the MDEs for the local average treatment effect are roughly twice as large again, and the same caution applies more strongly there.

4.5 What the sensor records suggest about mechanism

The treatment effects above tell us that the device worked, but not how. The devices themselves recorded water quality at high frequency for the duration of the intervention, and the app recorded feeding. We use these records to look at management behaviour directly. Note that these analyses are exploratory. They are available only for device users, so there is no control group and no causal claim. We report them because the data are unusual and the patterns are informative about where to look next.

Table 5 relates the stability of dissolved oxygen over a crop cycle to farmer characteristics. Farmers with more than nine years of schooling maintained markedly steadier dissolved oxygen: the within-cycle standard deviation is lower by 0.372 ($p = 0.038$) against a mean of 0.720 among less educated farmers, and the coefficient of variation is lower by 0.120 ($p = 0.005$) against a mean of 0.188. The result holds with schooling entered continuously, where each additional year reduces the coefficient of variation by 0.017 ($p = 0.004$). Neither technical knowledge of dissolved oxygen nor farming experience shows a comparable association. Education therefore appears to matter for the execution of water management even though it does not moderate the average return to the device.

Table 6 asks whether farmers adjusted feeding after observing a water quality shock. Once the underlying growth trend in feed volume is removed, the response is essentially zero and does not vary with education: pooling the daily feeding records of all device users, feed rises by 26.4% in the three days after a stress event against 27.6% in comparable non-stress windows, and the within-cycle difference that nets out the growth trend is -0.84 percentage points (Table 6 reports the corresponding breakdown by schooling). Whatever farmers did in response to the information, adjusting feed volume was not the main channel.

Table 7 connects management to harvest outcomes. Instability of dissolved oxygen is associated with smaller shrimp as we find that a one-unit increase in the coefficient of variation raises the number of shrimp per kilogram by 768.9 ($p = 0.046$). A higher share of days under water stress is associated with both lower harvested weight (-1.133 , $p = 0.063$) and smaller shrimp (86.0 shrimp per kilogram, $p = 0.087$). We measure recording intensity as counts per culture day rather than as raw counts, because failed harvests truncate cycles and so mechanically depress raw counts; with this correction, recording intensity itself is unrelated to outcomes. Table 8 finds no consistent heterogeneity in these links.

Taken together, these associations suggest a narrower role for education than our original hypothesis proposed. Education does not determine whether a farmer benefits from being given the device, but it is associated with steadier day-to-day water management, and steadier water management is in turn associated with larger shrimp. This is closer to Laajaj and Macours (2026), who find that the capacity to integrate multidimensional information constrains technology adoption, than to a story in which education gates access to the benefits of information altogether. We emphasise that the chain is correlational at every link and that the sample is 87 to 105 crop cycles among device users.

Table 5: Water Quality Management and Schooling among Device Users

	Education > 9 years		Years of schooling	
	(1) DO SD	(2) DO CV	(3) DO SD	(4) DO CV
Education > 9 years	-0.372** (0.144)	-0.120** (0.037)		
Years of schooling (centered)			-0.048* (0.021)	-0.017*** (0.006)
No. of recording days	0.010 (0.004)	0.002 (0.001)	0.009 (0.004)	0.002 (0.001)
Observations	87	87	87	87
Bootstrap p	0.038	0.005	0.051	0.004
Mean (low education) Village FE	0.720	0.188	0.667	0.169
Farmer controls	Yes	Yes	Yes	Yes

Notes: Sample: device users only, cycle-level observations restricted to cycles with at least 10 water-quality recording days, because the within-cycle standard deviation is mechanically attenuated when few readings are available. DO SD is the within-cycle standard deviation of dissolved oxygen; DO CV is the coefficient of variation. Columns (1)–(2) use a binary indicator for more than nine years of schooling; columns (3)–(4) use years of schooling, mean-centered, to show that the association does not depend on where the threshold is placed. Farmer controls are baseline age, years of shrimp farming experience, technical knowledge of dissolved oxygen, years of residence in the commune, postlarvae stocked, total pond area, cooperative membership and training. Years of schooling is excluded because the reported coefficient is itself a function of schooling. Number of recording days is controlled throughout. Parentheses report conventional standard errors clustered at the farmer level. Stars and the p -value row are both from a wild cluster bootstrap (farmer clusters, 9,999 replications), not from conventional asymptotic thresholds. **These are descriptive associations among device users, not causal effects of schooling:** education was not randomly assigned and no non-user comparison group has sensor records.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Feeding Response to Water Quality Stress and Schooling among Device Users

	All stress events		Averaged by cycle	
	(1)	(2)	(3)	(4)
	Feed (kg)	Feed (%)	Feed (kg)	Feed (%)
Education > 9 years	-0.178 (1.936)	14.343 (7.766)	5.780 (5.396)	9.491 (10.120)
pH stress (0/1)	1.269 (4.132)	-11.095 (4.737)	-17.633 (9.066)	-86.135 (42.271)
DO stress (0/1)	-2.571 (3.975)	-8.318 (3.618)	-8.280 (13.761)	-24.857 (31.924)
Salinity stress (0/1)	2.827 (3.378)	5.762 (8.258)	8.083 (17.627)	34.540 (41.601)
Shrimp age when stress occurred	-0.047 (0.034)	-0.827 (0.099)	-0.054 (0.061)	-0.508 (0.210)
Observations	1,638	1,638	100	100
Bootstrap p	0.936	0.166	0.418	0.408
Bootstrap CI lower	-5.509	-8.339	-4.155	-15.609
Bootstrap CI upper	5.654	37.162	21.669	33.099
Mean (low education)	8.539	26.392	8.151	28.417
Village FE	Yes	Yes	Yes	Yes
Farmer controls	Yes	Yes	Yes	Yes
Cultivation stage control	Yes	Yes	Yes	Yes

Notes: The dependent variable is the difference in average daily feed quantity between the three days after and the three days before a water quality stress event, where a stress event is a day on which pH, dissolved oxygen, salinity or temperature falls outside the recommended range. Columns (1)–(2) use every stress event. Feeding is recorded daily, so each window contains actual observations, but windows around successive events overlap. Any resulting attenuation applies equally to both education groups. Columns (3)–(4) average all events within a production cycle so that each cycle contributes one observation. A daily distributed-lag specification is not feasible: water quality is recorded every two days at the median, so requiring two or more consecutive daily readings reduces the sample sharply. Farmer controls are baseline age, years of shrimp farming experience, technical knowledge of dissolved oxygen, years of residence in the commune, postlarvae stocked, total pond area, cooperative membership and training. Years of schooling is excluded because the reported coefficient is itself a function of schooling. Cultivation stage is days of culture at the stress event. Parentheses report conventional standard errors clustered at the farmer level; p -values are from a wild cluster bootstrap (9,999 replications). **These are descriptive associations among device users, not causal effects of schooling.** Significance should be read from the bootstrap p -value row rather than from conventional asymptotic thresholds, so no stars are shown.

Table 7: Water Quality Management and Harvest Outcomes

	(1) ln(weight)	(2) Shrimp per kg	(3) Days of culture	(4) Failed harvest
DO coefficient of variation	-6.005 [0.139]	768.856** [0.046]	-13.771 [0.697]	0.740 [0.372]
Share of days with stress	-1.133* [0.063]	85.995* [0.087]	0.303 [0.976]	0.227 [0.202]
Water readings per culture day	-3.464 [0.320]	-177.056 [0.686]	-55.208 [0.172]	-0.568 [0.339]
Feeding records per culture day	-1.251 [0.404]	56.714 [0.958]	-28.991 [0.243]	-0.162 [0.646]
Observations	90	92	92	92
Cycle length controlled	Yes	Yes	—	Yes
Village FE	Yes	Yes	Yes	Yes
Farmer controls	Yes	Yes	Yes	Yes

Notes: Each cell is a separate regression of the harvest outcome in the column on the water management measure in the row. The sample is production cycles for which both sensor records and harvest records exist. Shrimp per kg is measured at the final harvest of the cycle, so a lower value means larger shrimp. Failed harvest is an indicator that the cycle ended with no marketable output. Recording intensity is expressed per day of culture rather than as a raw count, and days of culture is controlled in columns (1), (2) and (4). Raw counts would mechanically track cycle length, and because a failed cycle is terminated early they would make frequent recording appear to prevent failure. Farmer controls are baseline age, years of schooling, years of shrimp farming experience, technical knowledge of dissolved oxygen, years of residence in the commune, postlarvae stocked, total pond area, cooperative membership and training. Wild cluster bootstrap *p*-values in brackets (farmer clusters, 9,999 replications). **These associations are exploratory: water management is a farmer choice, not an experimental manipulation, so the estimates cannot be read as causal effects of management on yield.**

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Water Quality Management and Harvest Outcomes: Heterogeneity

		(1) ln(weight)	(2) Shrimp per kg	(3) Days of culture	(4) Failed harvest
Schooling > 9 yrs	DO CV	-9.510	973.218**	18.939	1.489
	× characteristic	9.331	-476.241	-70.455	-1.849
Experience ≥ 14 yrs	DO CV	-1.039	399.238*	7.635	1.056
	× characteristic	-13.193	893.665	-90.135	-0.852
DO knowledge ≥ 3	DO CV	0.610	591.210	-10.150	0.873
	× characteristic	-9.201	257.845	-4.734	-0.165

Notes: Each block reports one regression per column in which the coefficient of variation of dissolved oxygen is interacted with the farmer characteristic named in the row. The first row of each block is the association for farmers without the characteristic; the second is the difference for those with it. Village fixed effects and baseline farmer controls are included throughout. The control set omits the continuous variable from which the row's characteristic is constructed, since conditioning on it would reduce the reported coefficient to a discontinuity at the threshold rather than the difference between the groups; the other two characteristics remain controlled, so the three rows are mutually adjusted. Wild cluster bootstrap *p*-values determine the stars (farmer clusters, 9,999 replications). **Exploratory associations, not causal effects.**

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Conclusion

This paper asked whether smallholders can benefit from complex, high-frequency bio-physical information, and whether those benefits depend on human capital. Working with an Indonesian aquatech startup, we gave shrimp farmers in the Mekong Delta region of Vietnam free access to a water quality meter and a farm management app for ten months.

We find that they can. Access to the device reduced the probability of a disease outbreak by 0.28 against a control mean of 0.32, and raised revenue and profit. The disease results are precisely estimated and survive worst-case bounds on selective attrition; the financial results point the same way but sit at the margin of what this sample can detect, and we treat them as the weaker of the two findings.

We expected these benefits to be skill-biased, and they were not. Neither schooling, nor technical knowledge of water chemistry, nor farming experience moderated the treatment effect. This is the opposite of what the skill-biased technological change and selective attention frameworks would predict, and it is encouraging: it suggests that when complex information is delivered in a usable form, farmers across the education distribution can act on it. We note that our farmers practise super-intensive aquaculture and are therefore not the most marginalized producers in the region; whether the same holds for less sophisticated operations is an open question.

The sensor records suggest a more subtle role for education. Farmers with more schooling maintained more stable dissolved oxygen over the crop cycle, while technical knowledge and experience showed no such association, and instability of dissolved oxygen was in turn associated with smaller shrimp. Education may therefore matter for the day-to-day execution of water management even where it does not moderate the average return to the device. These associations are not causal and we present them as a direction for future work rather than as a finding.

Two limitations bear on interpretation. Take-up was low, at twenty of forty-three assigned farmers, which is typical of long intervention periods in which farmers weigh the time and effort of participating, and attrition reached 30.9%. Neither appears to threaten the sign or significance of our results, but both limit precision, and better designs to encourage participation should be considered for future studies. For policy, our results are more optimistic than the narrative in which digital tools widen rural inequality: hardware alone delivered measurable gains without requiring farmers to be pre-selected on education. As “Aquaculture 4.0” expands, the open design question is therefore not whether smallholders can use complex information, but how to raise take-up, and future research should explore what bundling of hardware, subscription and support achieves that.

References


- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 4, pp. 1043–1171). Elsevier.
- Aker, J. C. (2011). Dial “A” for agriculture: A review of information and communication technologies for agricultural extension in developing countries. *Agricultural Economics*, 42(6), 631–647.
- Aker, J. C., & Fafchamps, M. (2015). Mobile phone coverage and producer markets: Evidence from West Africa. *The World Bank Economic Review*, 29(2), 262–292.
- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist’s companion*. Princeton University Press.
- Arouna, A., Michler, J. D., Yergo, W. G., & Saito, K. (2021). One size fits all? Experimental evidence on the digital delivery of personalized extension advice in Nigeria. *American Journal of Agricultural Economics*, 103(2), 596–619.
- Asche, F., Eggert, H., Oglend, A., Roheim, C. A., & Smith, M. D. (2022). Aquaculture: Externalities and policy options. *Review of Environmental Economics and Policy*, 16(2), 282–305.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Cinelli, C., & Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1), 39–67. <https://doi.org/10.1111/rssb.12348>
- Cole, S., & Sharma, G. (2018). The promise and challenges of implementing ICT in Indian agriculture. *India Policy Forum*, 14(1), 173–240.
- Cole, S. A., & Fernando, A. N. (2021). ‘Mobile’izing agricultural advice: Technology adoption, diffusion and sustainability. *The Economic Journal*, 131(633), 192–219.
- Fabregas, R., Kremer, M., Lowes, M., On, R., & Zane, G. (2025). Digital information provision and behavior change: Lessons from six experiments in East Africa. *American Economic Journal: Applied Economics*, 17(1), 527–566.
- Fu, X., & Akter, S. (2016). The impact of mobile phone technology on agricultural extension services delivery: Evidence from India. *The Journal of Development Studies*, 52(11), 1561–1576.

- Hoenig, J. M., & Heisey, D. M. (2001). The abuse of power: The pervasive fallacy of power calculations for data analysis. *The American Statistician*, 55(1), 19–24.
- Hörner, D., Bouguen, A., Frölich, M., & Wollni, M. (2022). Knowledge and adoption of complex agricultural technologies: Evidence from an extension experiment. *The World Bank Economic Review*, 36(1), 68–90.
- Imbens, G. W., & Manski, C. F. (2004). Confidence intervals for partially identified parameters. *Econometrica*, 72(6), 1845–1857.
- Jensen, R. (2007). The digital divide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The Quarterly Journal of Economics*, 122(3), 879–924.
- Joffre, O. M., Poortvliet, P. M., & Klerkx, L. (2018). Are shrimp farmers actual gamblers? An analysis of risk perception and risk management behaviors among shrimp farmers in the Mekong Delta. *Aquaculture*, 495, 528–537.
- Kruk, S. R. L., Bush, S. R., & Phillips, M. (2024). Federating ‘Aquaculture 4.0’ for data-driven social and environmental sustainability. *Marine Policy*, 169, Article 106355.
- Laajaj, R., & Macours, K. (2026). The complexity of multidimensional learning in agriculture. *Econometrica*, 94(2), 465–503. <https://doi.org/10.3982/ECTA22974>
- Lafferty, K. D., Harvell, C. D., Conrad, J. M., Friedman, C. S., Kent, M. L., Kuris, A. M., Powell, E. N., Rondeau, D., & Saksida, S. M. (2015). Infectious diseases affect marine fisheries and aquaculture economics. *Annual Review of Marine Science*, 7, 471–496.
- Larochelle, C., Alwang, J., Travis, E., Barrera, V. H., & Dominguez Andrade, J. M. (2019). Did you really get the message? Using text reminders to stimulate adoption of agricultural technologies. *The Journal of Development Studies*, 55(4), 548–564.
- Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies*, 76(3), 1071–1102.
- McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics*, 99(2), 210–221.
- Nacher, D., & Schündeln, M. (2022). The demand for advice: Theory and empirical evidence from farmers in Sub-Saharan Africa. *The World Bank Economic Review*, 36(1), 91–113.
- Nakasone, E., Torero, M., & Minten, B. (2014). The power of information: The ICT revolution in agricultural development. *Annual Review of Resource Economics*, 6(1), 533–550.

- Nguyen, T. A. T., Nguyen, K. A. T., & Jolly, C. (2019). Is super-intensification the solution to shrimp production and export sustainability? *Sustainability*, 11(19), Article 5277.
- Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2), 187–204.
- Oyinbo, O., Chamberlin, J., Abdoulaye, T., & Maertens, M. (2022). Digital extension, price risk, and farm performance: Experimental evidence from Nigeria. *American Journal of Agricultural Economics*, 104(2), 831–852.
- Schwartzstein, J. (2014). Selective attention and learning. *Journal of the European Economic Association*, 12(6), 1423–1452.
- Suzuki, A., Olivia, S., Nam, V. H., & Lee, G. (2025). Contaminated water spillovers or peer effects? Determinants of disease outbreaks in shrimp farming in Vietnam. *Agricultural Economics*, 56(2), 210–227. <https://doi.org/10.1111/agec.12872>
- Tran, N., Bailey, C., Wilson, N., & Phillips, M. (2013). Governance of global value chains in response to food safety and certification standards: The case of shrimp from Vietnam. *World Development*, 45, 325–336.
- Van Campenhout, B., Spielman, D. J., & Lecoutere, E. (2021). Information and communication technologies to provide agricultural advice to smallholder farmers: Experimental evidence from Uganda. *American Journal of Agricultural Economics*, 103(1), 317–337

Appendix

Figure A1: Propensity Score Distribution Before and After Matching



(a) Before Matching

(b) After Matching

Figure A2: Covariate Balance Before and After Matching




Table A1: Descriptive Statistics, Participants vs. Non-participants

	Non-participant			Participant			Diff.	<i>p</i> -value
	Mean	Median	SD	Mean	Median	SD		
<i>N</i>		242			20			
Super-intensive (=1)	0.992	1.000	0.091	1.000	1.000	0.000	0.008	0.683
Experience (yrs)	13.042	13.000	5.256	13.850	13.500	5.050	0.808	0.508
Married (=1)	0.979	1.000	0.143	1.000	1.000	0.000	0.021	0.516
gender (=1: male)	0.900	1.000	0.301	1.000	1.000	0.000	0.100	0.139
Age	48.908	48.500	10.518	49.550	49.500	6.716	0.642	0.789
Education (yrs)	8.012	8.000	2.324	8.950	9.000	1.959	0.938	0.081
Residence in commune (yrs)	44.554	46.000	14.262	49.550	49.500	6.716	4.996	0.122
Household size	3.850	4.000	1.675	3.650	4.000	1.348	-0.200	0.604
No. of ponds	2.013	2.000	1.259	2.200	2.000	1.361	0.188	0.525
Cooperative member	0.217	0.000	0.413	0.550	1.000	0.510	0.333	0.001
Farmer group member	0.037	0.000	0.190	0.000	0.000	0.000	-0.037	0.380
Owns smartphone	0.912	1.000	0.283	0.950	1.000	0.224	0.037	0.564
Info: TV/radio	0.725	1.000	0.447	0.800	1.000	0.410	0.075	0.469
Info: internet	0.446	0.000	0.498	0.700	1.000	0.470	0.254	0.029
Sales contract	0.129	0.000	0.336	0.200	0.000	0.410	0.071	0.374
Received training	0.812	1.000	0.391	1.000	1.000	0.000	0.188	0.033
Aware of prohibited inputs	0.046	0.000	0.210	0.100	0.000	0.308	0.054	0.287
Knows prohibited elements	0.567	1.000	0.497	0.900	1.000	0.308	0.333	0.003
Knows prohibited products (specific)	0.054	0.000	0.227	0.150	0.000	0.366	0.096	0.087
Extension officer (freq.)	5.758	6.000	1.186	4.550	4.000	1.317	-1.208	0.000
Postlarvae stocked (no.)	4.4e+05	3.0e+05	4.2e+05	6.0e+05	4.5e+05	6.1e+05	1.6e+05	0.109
Total pond area (ha)	0.552	0.240	0.946	1.342	0.320	2.852	0.790	0.005
ln(Revenue)	6.292	6.332	1.094	6.410	6.601	0.984	0.118	0.649
ln(Cost)	5.833	5.857	0.686	5.979	5.971	0.821	0.146	0.369
ln(Profit)	5.465	5.492	1.313	5.455	5.843	1.589	-0.010	0.977
ln(Revenue/ha)	7.579	7.873	1.387	7.433	7.713	1.161	-0.147	0.655
ln(Cost/ha)	7.065	7.293	0.887	6.826	7.058	1.024	-0.239	0.253
ln(Profit/ha)	6.869	7.074	1.426	6.714	7.043	1.416	-0.155	0.676
Disease rate (avg)	0.219	0.000	0.403	0.392	0.000	0.475	0.172	0.071
Disease (total)	0.321	0.000	0.635	0.450	0.000	0.510	0.129	0.376
Disease (any, =1)	0.237	0.000	0.426	0.450	0.000	0.510	0.212	0.036

Notes: Participants are the twenty farmers who joined at the start of the intervention. The sample is the estimation sample used in the main text.

Table A2: Difference-in-Differences Estimates (ITT)*Panel A: Disease Outcomes*

	Disease (avg)	Disease (total)	Disease (any)
Assigned to treatment (ITT)	-0.316** (0.127)	-0.490** (0.201)	-0.419** (0.155)
Other invited farmers in village	-0.005 (0.019)	-0.030 (0.030)	-0.008 (0.023)
N	176	176	176
Bootstrap <i>p</i> (treatment)	0.035	0.027	0.016
Bootstrap CI lower	-0.597	-0.956	-0.769
Bootstrap CI upper	-0.021	-0.049	-0.077
Bootstrap <i>p</i> (spillover)	0.861	0.396	0.777
Control mean (endline)	0.247	0.441	0.324

Notes: Difference-in-differences estimates on the same sample and with the same controls as the ANCOVA specification. The dependent variable is the change in the outcome between 2023 and 2024, which constrains the coefficient on the baseline outcome to one. Sample restricted to farmers drawn from the original randomization list and never-invited farmers. Commune fixed effects, baseline farmer characteristics (age, years of shrimp farming experience, years of schooling, years of residence in the commune, postlarvae stocked, and total pond area) and institutional-connection controls (cooperative membership, training, input knowledge) are included, as is the number of other invited farmers in the same village together with village size, to allow for spillovers from treated neighbours. Parentheses report conventional cluster-robust standard errors; *p*-values and confidence intervals are from a wild cluster bootstrap (village clusters, 9,999 replications).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Financial Outcomes

	ln(Rev)	ln(Cost)	ln(Profit)	Ln(Rev/ha)	Ln(Cost/ha)	Ln(Profit/ha)
Assigned to treatment (ITT)	0.609** (0.227)	0.509** (0.154)	0.670 (0.479)	0.431 (0.245)	0.290** (0.130)	0.483 (0.466)
Other invited farmers in village	0.011 (0.047)	0.044 (0.022)	0.029 (0.047)	0.007 (0.036)	0.034 (0.019)	0.023 (0.044)
N	167	176	141	169	176	141
Bootstrap <i>p</i> (treatment)	0.029	0.011	0.217	0.122	0.044	0.339
Bootstrap CI lower	0.061	0.124	-0.533	-0.134	0.007	-0.694
Bootstrap CI upper	1.072	0.820	1.717	0.941	0.560	1.529
Bootstrap <i>p</i> (spillover)	0.852	0.315	0.607	0.873	0.150	0.623
Control mean (endline)	6.170	5.688	5.316	7.684	7.190	6.825

Notes: Difference-in-differences estimates on the same sample and with the same controls as the ANCOVA specification. The dependent variable is the change in the outcome between 2023 and 2024, which constrains the coefficient on the baseline outcome to one. Sample restricted to farmers drawn from the original randomization list and never-invited farmers. Commune fixed effects, baseline farmer characteristics (age, years of shrimp farming experience, years of schooling, years of residence in the commune, postlarvae stocked, and total pond area) and institutional-connection controls (cooperative membership, training, input knowledge) are included, as is the number of other invited farmers in the same village together with village size, to allow for spillovers from treated neighbours. Parentheses report conventional cluster-robust standard errors; *p*-values and confidence intervals are from a wild cluster bootstrap (village clusters, 9,999 replications).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Estimated on the same sample and with the same controls as the ANCOVA specification in Table 2. The two estimators differ only in how the baseline outcome enters: differencing constrains its coefficient to one, whereas ANCOVA estimates it. Because serial correlation in these outcomes is close to zero, the estimated coefficient is near zero and the differencing constraint is strongly rejected.

Table A3: Difference-in-Differences Estimates (LATE)*Panel A: Disease Outcomes*

	Disease (avg)	Disease (total)	Disease (any)
JALA × Post (IV)	-0.296** (0.1414)	-0.416* (0.230)	-0.468** (0.167)
Post	0.043 (0.050)	0.124 (0.086)	0.110* (0.061)
N	512	512	512
Within R^2	.	.	.
1st stage F	59.50	59.50	59.50
Control mean	0.206	0.309	0.221

Notes: Standard errors clustered at village level in parentheses. Treatment variable (JALA × Post) is instrumented by treatment assignment (Invited × Post). Farmer and time fixed effects included. Control mean is the unweighted mean of the outcome for the control group at baseline.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Financial Outcomes

	ln(Rev)	ln(Cost)	ln(Profit)	Ln(Rev/ha)	Ln(Cost/ha)	Ln(Profit/ha)
JALA × Post (IV)	-0.110 (0.307)	0.096 (0.237)	0.389 (0.399)	-0.221 (0.294)	0.068 (0.215)	0.160 (0.416)
Post	-0.255*** (0.090)	-0.194*** (0.070)	-0.317** (0.082)	-0.185* (0.108)	-0.136* (0.079)	-0.305*** (0.104)
Baseline disease × Post	1.021*** (0.355)	0.453 (0.284)	1.570*** (0.317)	1.755*** (0.446)	1.138*** (0.230)	2.281*** (0.278)
N	501	512	453	501	512	453
Within R^2	0.105	0.041	0.148	0.209	0.216	0.252
1 st stage F	60.57	60.7	60.57	60.57	60.57	60.57
Control mean	6.303	5.843	5.494	7.585	7.074	6.883

Notes: Standard errors clustered at village level in parentheses. Treatment variable (JALA × Post) is instrumented by treatment assignment (Invited × Post). Farmer and time fixed effects included. All specifications control for the baseline (2023) average disease rate interacted with the post indicator. Control mean is the unweighted mean of the outcome for the control group at baseline.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A4: Moderator Comparison: Difference-in-Differences and ANCOVA

		(1) Education	(2) Tech Knowledge	(3) Experience
Disease (avg)	DID	0.034 [0.850]	-0.207 [0.304]	-0.001 [0.996]
Disease (avg)	ANCOVA	-0.008 [0.938]	-0.033 [0.820]	0.054 [0.498]
ln(Revenue/ha)	DID	-0.162 [0.803]	0.135 [0.697]	-0.184 [0.756]
ln(Revenue/ha)	ANCOVA	0.166 [0.834]	0.133 [0.641]	-0.083 [0.846]
ln(Cost/ha)	DID	0.389 [0.383]	0.173 [0.468]	0.091 [0.823]
ln(Cost/ha)	ANCOVA	0.328 [0.368]	0.087 [0.601]	-0.034 [0.910]
ln(Profit/ha)	DID	0.423 [0.519]	0.097 [0.848]	0.16 [0.825]
ln(Profit/ha)	ANCOVA	0.750 [0.302]	-0.156 [0.737]	-0.112 [0.812]

Notes: Each cell reports the interaction between treatment assignment and the moderator named in the column heading, estimated by difference-in-differences and by ANCOVA on the same sample with the same controls. Wild cluster bootstrap *p*-values in brackets (village clusters, 9,999 replications). Commune fixed effects and the full set of baseline controls are included throughout. No moderator produces a systematic interaction under either estimator. The ANCOVA rows correspond to the estimates reported in the main text; the difference-in-differences rows are shown alongside them to confirm that the absence of moderation is not an artefact of the choice of estimator.

Table A5: Propensity Score Estimation (Logit)

	JALA User (=1)
Experience (years)	-0.033 (0.034)
Age	0.020 (0.025)
Education (years)	0.233*** (0.088)
Years in commune	0.010 (0.016)
Household members	-0.047 (0.101)
Number of ponds	0.119 (0.118)
Constant	-4.696*** (1.596)
N	302
Pseudo R ²	0.036
χ^2	9.70

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A6: Treatment Effects, PSM-DID*Panel A: Disease Outcomes*

	Disease (avg)	Disease (total)	Disease (any)
JALA $\times$ Post	0.042 (0.133)	0.067 (0.247)	-0.067 (0.169)
Post	-0.117 (0.082)	-0.067 (0.160)	-0.033 (0.133)
<i>N</i>	156	156	156
Within R^2	0.035	0.003	0.019
Control mean	0.286	0.381	0.310

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Financial Outcomes

	ln(Rev)	ln(Cost)	ln(Profit)	ln(Rev/ha)	ln(Cost/ha)	ln(Profit/ha)
JALA $\times$ Post	-0.475 (0.364)	0.057 (0.311)	0.182 (0.312)	-0.696* (0.383)	-0.135 (0.235)	0.241 (0.316)
Post	0.160 (0.344)	-0.225 (0.280)	0.022 (0.226)	0.624 (0.438)	0.210 (0.250)	0.074 (0.183)
<i>N</i>	155	156	139	155	156	139
Within R^2	0.032	0.043	0.012	0.052	0.019	0.027
Control mean	6.352	5.870	5.560	7.547	7.065	6.946

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Standard errors clustered at village level in parentheses. Sample restricted to 1:1 nearest-neighbor matched pairs with common support. PSM covariates: experience, age, education, years in commune, household size, number of ponds. Farmer and time fixed effects included. Control mean is the mean of the outcome for matched control units at baseline.

Table A7: Moderator Comparison, PSM-DID*Panel A: Disease Outcomes*

	(1) Education (>9 years)	(2) Tech Knowledge ($DO \geq 3$)	(3) Experience (≥ 14 years)
Disease (avg)			
JALA $\times$ Post ($\hat{\beta}_2$)	0.143 (0.135)	0.121 (0.220)	-0.103 (0.216)
JALA $\times$ Post $\times$ Moderator ($\hat{\beta}_4$)	-0.259 (0.323)	-0.104 (0.244)	0.282 (0.217)
<i>N</i>	156	156	156
Control mean	0.286	0.286	0.286
Disease (total)			
JALA $\times$ Post ($\hat{\beta}_2$)	0.27 (0.260)	0.313 (0.402)	-0.033 (0.371)
JALA $\times$ Post $\times$ Moderator ($\hat{\beta}_4$)	-0.637 (0.402)	-0.325 (0.464)	0.233 (0.341)
<i>N</i>	156	156	156
Control mean	0.381	0.381	0.381
Disease (any)			
JALA $\times$ Post ($\hat{\beta}_2$)	0.033 (0.183)	0 (0.291)	-0.2 (0.264)
JALA $\times$ Post $\times$ Moderator ($\hat{\beta}_4$)	-0.2 (0.359)	-0.076 (0.302)	0.256 (0.243)
<i>N</i>	156	156	156
Control mean	0.310	0.310	0.310

Notes: Standard errors clustered at village level in parentheses. Sample restricted to 1:1 nearest-neighbor matched pairs with common support. PSM covariates: experience, age, education, years in commune, household size, number of ponds. Moderator thresholds: Education > 9 years; DO knowledge score ≥ 3 (baseline median); experience ≥ 14 years (baseline median). Farmer and time fixed effects included. Control mean is the mean of the outcome for matched control units at baseline.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Financial Outcomes

	(1) Education (>9 years)	(2) Tech Knowledge ($DO \geq 3$)	(3) Experience (≥ 14 years)
ln (Revenue)			
JALA $\times$ Post (β_2)	-0.666 (0.390)	-0.653 (0.645)	0.19 (0.291)
JALA $\times$ Post $\times$ Moderator (β_4)	0.903 (1.079)	0.304 (0.712)	-1.266* (0.667)
<i>N</i>	155	155	155
Control mean	6.352	6.32	6.352
ln (Cost)			
JALA $\times$ Post (β_2)	0.05 (0.317)	0.144 (0.401)	0.042 (0.199)
JALA $\times$ Post $\times$ Moderator (β_4)	-0.315 (0.903)	-0.103 (0.504)	-0.049 (0.441)
<i>N</i>	156	156	156
Control mean	5.870	5.870	5.870
ln (Profit)			
JALA $\times$ Post (β_2)	-0.072 (0.327)	0.228 (0.640)	0.436 (0.514)
JALA $\times$ Post $\times$ Moderator (β_4)	0.438 (0.822)	-0.217 (0.732)	0.65 (0.811)
<i>N</i>	139	139	139
Control mean	5.560	5.560	5.560
ln (Revenue/ha)			
JALA $\times$ Post (β_2)	-0.941** (0.404)	-1.046 (0.668)	0.029 (0.273)
JALA $\times$ Post $\times$ Moderator (β_4)	1.486 (1.036)	0.563 (0.794)	-1.338* (0.727)
<i>N</i>	155	155	155
Control mean	7.547	7.547	7.547
ln (cost/ha)			
JALA $\times$ Post (β_2)	-0.191 (0.258)	-0.154 (0.561)	-0.119 (0.217)
JALA $\times$ Post $\times$ Moderator (β_4)	0.233 (0.821)	0.063 (0.582)	-0.066 (0.365)
<i>N</i>	156	156	156
Control mean	7.065	7.065	7.065
ln (Profit/ha)			
JALA $\times$ Post (β_2)	-0.008 (0.366)	-0.153 (0.572)	0.275 (0.478)
JALA $\times$ Post $\times$ Moderator (β_4)	0.526 (0.916)	0.464 (0.718)	-0.145 (0.833)
<i>N</i>	139	139	139
Control mean	6.946	6.946	6.946

Notes: Standard errors clustered at village level in parentheses. Sample restricted to 1:1 nearest-neighbor matched pairs with common support. PSM covariates: experience, age, education, years in commune, household size, number of ponds. Moderator thresholds: Education > 9 years; DO knowledge score ≥ 3 (baseline median); experience ≥ 14 years (baseline median). Farmer and time fixed effects included. Control mean is the mean of the outcome for matched control units at baseline.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A8: Attrition Analysis*Panel A: Attrition rate by treatment status*

	Attrition rate	N (Baseline)	Diff.	p-value
Full sample	0.309	304		
Control	0.338	219		
Invited (original, ITT)	0.279	43	-0.059	0.455
Invited (total, ITT)	0.235	85	-0.103	0.083
Non-participant	0.335	242		
Participant (TOT)	0.250	20	-0.085	0.440

Notes: A farmer is counted as having attrited if they did not consent to the 2024 follow-up survey. Diff. reports the difference in attrition rates: for the two *Invited* rows against the control group, and for *Participant* against *Non-participant*. *p*-value from a two-sample *t*-test. *Invited (original, ITT)* is the sample used in the main text, which restricts the treated group to the 43 farmers from the original randomization list; there the difference in attrition against the control group is small and insignificant. *Invited (total, ITT)* adds the 42 farmers assigned mid-intervention, who by construction were recruited later and were more likely to still be in contact with the survey team. *Non-participant* and *Participant (TOT)* split the main-analysis sample (control plus original list, *n*=262) by take-up from the start of the intervention (*n*=20 participants); the participant row conditions on having engaged with the device and so mechanically selects farmers who stayed in contact, and should not be read as an effect of take-up on attrition.

Panel B: Baseline characteristics of stayers and attriters

	Stayers			Attriters			Diff.	p-value
	Mean	Median	SD	Mean	Median	SD		
<i>N</i>		210			94			
Super-intensive (=1)	1.000	1.000	0.000	0.978	1.000	0.147	-0.022	0.032
Experience (yrs)	13.076	13.000	5.202	12.554	12.000	5.136	-0.522	0.421
Married (=1)	0.981	1.000	0.137	0.957	1.000	0.205	-0.024	0.225
gender (=1: male)	0.919	1.000	0.273	0.880	1.000	0.326	-0.039	0.288
Age	49.395	50.000	10.075	47.076	45.000	10.780	-2.319	0.073
Education (yrs)	7.948	8.000	2.322	8.620	9.000	2.343	0.672	0.022
Residence in commune (yrs)	46.286	48.000	13.092	41.000	41.500	15.028	-5.286	0.002
Household size	3.905	4.000	1.646	3.630	3.000	1.642	-0.274	0.183
No. of ponds	2.062	2.000	1.331	2.000	2.000	1.048	-0.062	0.693
Cooperative member	0.290	0.000	0.455	0.293	0.000	0.458	0.003	0.958
Farmer group member	0.029	0.000	0.167	0.043	0.000	0.205	0.015	0.507
Owns smartphone	0.914	1.000	0.281	0.935	1.000	0.248	0.02	0.546
Info: TV/radio	0.729	1.000	0.446	0.717	1.000	0.453	-0.011	0.842
Info: internet	0.457	0.000	0.499	0.511	1.000	0.503	0.054	0.391
Sales contract	0.133	0.000	0.341	0.141	0.000	0.350	0.008	0.853
Received training	0.814	1.000	0.390	0.880	1.000	0.326	0.066	0.156
Aware of prohibited inputs	0.057	0.000	0.233	0.033	0.000	0.179	-0.025	0.368
Knows prohibited elements	0.567	1.000	0.497	0.598	1.000	0.493	0.031	0.615
Knows prohibited products (specific)	0.067	0.000	0.250	0.054	0.000	0.228	-0.012	0.686
Extension officer (freq.)	5.667	6.000	1.235	5.543	6.000	1.262	-0.123	0.429
ln(Revenue)	6.299	6.327	1.119	6.291	6.497	0.948	-0.008	0.952
ln(Cost)	5.820	5.819	0.745	5.856	5.920	0.648	0.036	0.687
ln(Profit)	5.445	5.493	1.338	5.492	5.620	1.237	0.047	0.791
ln(Revenue/ha)	7.552	7.827	1.385	7.559	7.835	1.132	0.007	0.966
ln(Cost/ha)	7.046	7.244	0.842	7.004	7.221	0.924	-0.042	0.697
ln(Profit/ha)	6.821	7.039	1.449	6.878	7.070	1.228	0.056	0.765
Disease rate (avg)	0.225	0.000	0.406	0.251	0.000	0.417	0.026	0.616
Disease (total)	0.338	0.000	0.645	0.326	0.000	0.576	-0.012	0.878
Disease (any =1)	0.248	0.000	0.433	0.283	0.000	0.453	0.035	0.524

Panel C: Linear probability model of attrition

	(1)	(2)	(3)
Invited (=1)	-0.059 (0.079)	-0.033 (0.084)	
Participant (=1)			-0.097 (0.116)
Super-intensive (=1)		-0.759** (0.342)	-0.760** (0.342)
Experience (yrs)		-0.004 (0.006)	-0.003 (0.006)
Married (=1)		-0.282 (0.224)	-0.28 (0.224)
Gender (=1: male)		0.095 (0.134)	0.09 (0.134)
Age		0.001 (0.005)	0.001 (0.005)
Education (yrs)		0.021 (0.019)	0.02 (0.019)
Residence in commune (yrs)		-0.009*** (0.003)	-0.009*** (0.003)
Household size		-0.011 (0.019)	-0.012 (0.019)
No. of ponds		0.006 (0.024)	0.007 (0.024)
Cooperative member		0.013 (0.075)	0.016 (0.075)
Farmer group member		-0.012 (0.173)	-0.018 (0.172)
Owns smartphone		0.017 (0.109)	0.02 (0.109)
Info: TV/radio		0.023 (0.078)	0.022 (0.078)
Info: internet		-0.101 (0.09)	-0.094 (0.09)
Sales contract		0.043 (0.093)	0.042 (0.093)
Received training		0.152* (0.09)	0.155* (0.09)
Aware of prohibited inputs		-0.178 (0.155)	-0.189 (0.153)
Knows prohibited elements		-0.012 (0.066)	-0.011 (0.066)
Knows prohibited products (specific)		-0.057 (0.149)	-0.065 (0.149)
Extension officer (freq.)		-0.036 (0.031)	-0.04 (0.031)
<i>N</i>	262	260	260
<i>R</i> ²	0.002	0.118	0.12

Notes: Linear probability model (OLS). Dependent variable equals 1 if the farmer did not consent to the 2024 follow-up survey. The estimation sample is the main-analysis sample (control plus the original randomization list, n=262); on this sample the invited group coincides with the original list, so *Invited* in columns (1)–(2) is assignment to that list, and column (3) uses participation from the start of the intervention. Standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel D: Joint tests of attrition on baseline characteristics

	<i>F</i> -statistic	<i>p</i> -value
Joint test of all baseline characteristics	1.44	0.102
excluding production system (Super-intensive)	1.22	0.240
Composition of attriters across arms (interaction)	1.02	0.443

Notes: *F*-tests from linear-probability regressions of the attrition indicator on baseline characteristics. Rows 1–2 use all baseline farmers ($n=304$): row 1 tests the twenty baseline characteristics jointly, and row 2 re-estimates without the production-system (super-intensive) indicator. Row 3 uses the main-analysis sample (control plus the original randomization list, $n=262$) and tests the joint significance of the interactions between assignment to the original list and the twenty characteristics, i.e. whether the composition of attriters differs across arms.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: A farmer is counted as having attrited if they did not consent to the 2024 follow-up survey; 210 of the 304 baseline farmers were retained, so the attrition rate is 30.9%. Attrition does not differ significantly by treatment assignment. Panel D reports joint *F*-tests of the baseline characteristics: they do not jointly predict attrition, with or without the production-system indicator, and the composition of attriters does not differ significantly across arms. A few individual characteristics differ between stayers and attriters (Panel B), but not jointly, so attrition appears close to ignorable; Table A9 nonetheless reports bounds that do not require that assumption.

Table A9: Lee Bounds on the Treatment Effect

	Disease			Financial	
	(avg)	(total)	(any)	ln(Rev)	ln(Profit)
Unadjusted difference	-0.182	-0.345	-0.260	0.625	0.753
Lee bounds: lower	-0.247	-0.441	-0.324	0.490	0.493
Lee bounds: upper	-0.175	-0.334	-0.253	0.794	1.010
Imbens–Manski 95% CI: lower	-0.321	-0.570	-0.424	0.153	0.063
Imbens–Manski 95% CI: upper	-0.073	-0.157	-0.139	1.158	1.437
Trimming proportion q	0.082	0.082	0.082	0.101	0.162

Notes: Lee (2009) trimming bounds on the intention-to-treat effect, on the sample of farmers from the original randomization list ($n=43$) and never-invited farmers ($n=219$). Attrition is 33.8% in the control group and 27.9% among invitees; the difference is not statistically significant ($p=0.46$). Because attrition is higher among controls, the bounds are obtained by trimming a proportion q of the retained control observations from the top (lower bound) and from the bottom (upper bound), which equalizes the response rate across arms. The bounds are non-parametric and use no covariates, so they are not directly comparable to the covariate-adjusted ANCOVA estimates; they show the range of treatment effects consistent with any pattern of selection into response, under the assumption that invitation does not reduce the probability of responding for any farmer. Confidence intervals follow Imbens and Manski (2004), with bound standard errors from a village-cluster bootstrap (500 replications).

Table A10: Minimum Detectable Effects

Outcome	SD	ICC	N	Minimum detectable effect		Estimate
				Level	Relative	
<i>Panel A: Intention-to-treat (ANCOVA)</i>						
Disease (avg)	0.391	0.010	176	0.182	7364%	-0.200
Disease (total)	0.706	0.000	176	0.348	7892%	-0.404
Disease (any)	0.470	0.045	176	0.224	6902%	-0.282
ln(Revenue)	0.882	0.079	169	0.634	8849%	0.536
ln(Profit)	1.076	0.073	141	0.837	13099%	0.627
<i>Panel B: Local average treatment effect (ANCOVA-IV)</i>						
Disease (avg)	0.391	0.010	176	0.182	20417%	-0.422
Disease (total)	0.706	0.000	176	0.348	23472%	-0.862
Disease (any)	0.470	0.045	176	0.224	20598%	-0.597
ln(Revenue)	0.882	0.079	169	0.634	32554%	1.067
ln(Profit)	1.076	0.073	141	0.837	42590%	1.268

Notes: SD is the standard deviation of the 2024 outcome in the control group and ICC is the corresponding village-level intra-cluster correlation. The minimum detectable effect (MDE) is computed at 80% power and a 5% significance level as $2.802 \times SE$, where SE is the cluster-robust standard error actually achieved by the corresponding specification in Table 2 (Panel A) and Table 3 (Panel B). Using the realized standard error rather than reconstructing a design effect from assumed parameters accounts for village clustering, the baseline and covariate adjustment, and the achieved sample size simultaneously. The MDE depends only on the standard error, not on the estimated coefficient. *Level* is the MDE in the units of the outcome. *Relative* expresses it against the control group: for the disease outcomes as a percentage of the control mean, and for the log outcomes as the implied percentage change, $100 \times (e^{MDE} - 1)$. The estimate is reproduced in the last column for comparison only; we do not report power computed at the observed effect size, since it is a monotone transformation of the p -value and adds nothing to the estimates and confidence intervals reported elsewhere (Hoenig and Heisey, 2001). Take-up is 20 of 43 assigned farmers, which is why the LATE panel has an MDE roughly twice as large as the ITT panel.