INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

CAN AI-ENABLED AGRICULTURAL SERVICES IMPROVE SMALLHOLDER CLIMATE RESILIENCE? EVIDENCE ON PRODUCTIVITY, INPUT EFFICIENCY, AND FARM INCOME

Tahira Ilyas [1], Mahnoor Syed [2], Saiqa Bilal [3], Khadija Iqbal [4], Kiran Saeemab [5]

DOI: https://doi.org/10.63544/ijss.v5i6.349

Affiliations:

[1] Department of Economics and Management Sciences, Women University of Azad Jammu and Kashmir, Bagh, Pakistan
Email: tahirailyas802@gmail.com

[2] Faculty of Pure Sciences, Gomal University, KPK, Pakistan
Email: syedmahnoor20003@gmail.com

[3] School of Economics, University of The Punjab, Lahore, Pakistan
Email: saiqaqu19bilal@gmail.com

[4] Human nutrition and dietetics, Government College Woman University, Pakistan
Email: khadijaaaa302@gmail.com

[5] University of Sargodha, Pakistan
Email: kiranseemab512@gmail.com

Corresponding Author's Email:

[1] tahirailyas802@gmail.com

Copyright:

Author/s

License:

Article History:
Received: 25.10.2026
Accepted: 10.11.2026
Published: 29.11.2026

Abstract

Climate variability threatens the productivity and income stability of smallholder farms, while AI-enabled agricultural services (AI-AS) are expanding rapidly. These services use satellite, weather, soil and farm-level data to provide personalised advisories on weather, pests, input use and crop management. However, evidence on whether they improve climate resilience, not only average yields, remains limited.

This study evaluates the effects of AI-AS adoption on crop productivity, input-use efficiency, net farm income and resilience to drought and heat shocks. Using farm-household panel data matched with gridded climate data, the study applies staggered difference-in-differences estimators, a single-step stochastic production frontier, shock-interaction models and downside-risk analysis. Heterogeneity is examined by farm size, gender, education, digital connectivity and agro-ecological zone.

AI-AS adoption increases crop yield by approximately 13 per cent, reduces input cost per kilogram of output by 14.5 per cent, improves technical efficiency by 2.5 percentage points and raises net farm income by approximately 11 per cent. Adoption also reduces losses during drought and heat shocks and lowers the probability of a large income loss. The benefits are broadly similar across farm and household groups.

The findings show that AI-enabled advisory services can improve productivity, efficiency, income and climate resilience among smallholder farmers. The study provides evidence for governments, donors and service providers on the potential of AI-AS for inclusive and climate-resilient agricultural development.

Keywords: AI-Enabled Agricultural Services, Smallholder Farmers, Climate Resilience, Input-Use Efficiency, Farm Income, Digital Agriculture, Staggered Difference-in-Differences

1. Introduction

1.1. Background of the Study

Climate change has become one of the most direct threats to agricultural livelihoods in the developing world. Rising temperatures, erratic rainfall and more frequent droughts and heatwaves have already slowed the growth of global agricultural productivity (Ortiz-Bobea et al., 2021). Smallholder farms are among the most exposed, because their incomes depend on a single or a few crops, their assets are thin, and their capacity to absorb losses is limited. Adaptation is possible in principle: farmers can shift planting dates, change varieties, adjust fertiliser and water use, and respond earlier to pests and disease. What they frequently lack is the information to do so at the right moment. Decisions about when to plant, how much to apply and how to

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

react to a weather shock are made under uncertainty, and the cost of a wrong decision is highest in the seasons when conditions are least predictable.

Over the past two decades, digital tools have changed how information reaches farmers. Mobile phones lowered the cost of obtaining market and weather information and improved the functioning of rural markets (Aker & Mbiti, 2010; Jensen, 2007), and mobile-based extension has become a widely used channel for advisory content (Aker, 2011). A newer generation of services goes further. AI-enabled agricultural services combine satellite imagery, weather forecasts, soil data and farm records with machine-learning models to generate advice that is specific to a field and a moment: which variety to sow and when, how much fertiliser a particular plot needs, and whether a detected pest or disease warrants treatment. Because the advice is conditioned on the farm and on the forecast, it is most useful precisely when conditions are uncertain (Ali et al., 2025).

For smallholders the potential is considerable. Timely and tailored advice could help farmers use fertiliser, seed, water and labour more efficiently, raise yields, and reduce losses during adverse weather (Ali et al., 2025). Yet the benefits are not automatic. Advice must be trusted, understood and affordable; it must reach farmers who can act on it; and it must arrive before the decision window closes. Services that are technically sophisticated but poorly matched to local language, literacy, connectivity or plot size may widen rather than narrow the gap between better-off and marginal farms (Dil & Ali, 2025).

1.2 Problem Statement

Despite rapid investment in digital agriculture, substantial gaps remain in what is known about what AI-enabled services deliver for smallholders. Existing evaluations of digital and mobile advisory services have mostly examined adoption and average yield effects, and they show that results depend strongly on content, design and context (Aker, 2011; Cole & Fernando, 2021). Three gaps follow. First, input-use efficiency is rarely measured, although it determines input costs, emissions and the sustainability of intensification. Second, net farm income is rarely the outcome of interest, although it determines household welfare and whether the gains cover any service fee. Third, and most important, resilience is rarely tested. Most studies ask whether adopters produce more on average; few ask whether adopters lose less when a drought or a heatwave occurs.

Resilience is conceptually distinct from average performance (Barrett & Constas, 2014). A service can raise mean yields while leaving the downside distribution unchanged, in which case it improves welfare in normal years but offers no protection when protection matters most. This distinction matters for policy (Baig et al., 2025). Without evidence on shock sensitivity, governments, donors and providers risk scaling services that raise costs without improving outcomes, or that benefit better-off farmers while leaving the most climate-exposed households behind. Evidence is also scarce on who benefits, since farm size, gender, education and connectivity shape both access to the service and the returns to acting on its advice.

1.3 Research Questions

The study is guided by five interrelated research questions. The fourth and fifth separate the two ideas that are often combined: whether adoption reduces vulnerability to shocks, and whether the gains are evenly distributed.

RQ1. Does adoption of AI-enabled agricultural services raise crop productivity among smallholder farmers?

RQ2. Does adoption improve input-use efficiency, measured as technical efficiency and as input cost per unit of output?

RQ3. Does adoption raise net farm income?

RQ4. Does adoption reduce the sensitivity of productivity and income to drought and heat shocks?

RQ5. Are the gains distributed evenly, or concentrated among larger, male-headed, better-educated and better-connected farms?

1.4 Research Objectives

The primary objective is to evaluate whether AI-enabled agricultural services improve smallholder performance and climate resilience. Five specific objectives correspond one to one with the research questions:

Objective 1. To estimate the effect of AI-AS adoption on crop productivity.

Objective 2. To estimate the effect of adoption on technical efficiency and on input cost per unit of

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

output.

Objective 3. To estimate the effect of adoption on net farm income.

Objective 4. To test whether adoption weakens the negative effects of drought and heat shocks on productivity and income.

Objective 5. To identify differences in effects by farm size, gender of the household head, education, digital connectivity and agro-ecological zone, and to derive policy recommendations for inclusive and cost-effective scaling.

1.5 Contribution

The study makes three contributions. The theoretical contribution is to link the literatures on technology adoption, ICT-enabled extension and resilience by treating resilience as the capacity of adopters to limit losses under shocks rather than as a higher average outcome. This reframing matters because it changes what counts as success: a service that raises mean yields by the same amount in every season is a productivity intervention, whereas a service that preserves yields in bad seasons is an adaptation intervention, and the two imply different policy conclusions (Dil & Ali, 2025).

The empirical contribution is joint evidence on productivity, input efficiency and farm income for the same households, combined with explicit tests of shock sensitivity. Much of the existing evidence reports a single outcome, which makes it impossible to tell whether income gains come from higher output or from lower costs, and therefore whether the service is worth its fee. Estimating all three outcomes on a common sample, with modern estimators that avoid the biases of two-way fixed effects under staggered adoption, allows the channels to be separated (Nondani et al., 2025).

The policy contribution is evidence on distribution. Findings on heterogeneity indicate whether scaling strategies need complementary support, such as connectivity, digital skills, insurance or credit, for the farmers who would otherwise be excluded, and they inform decisions about how services should be priced and targeted.

2. Literature Review and Hypotheses

2.1 Technology Adoption and Information Frictions

Adoption of agricultural innovations is not simply a matter of profitability. Feder et al. (1985) show in their survey that adoption depends on the characteristics of the farm and the farmer, on access to information, credit and markets, and on the perceived riskiness of the innovation. Where returns are high but information is costly, adoption can stall even among farmers who would gain from it. Suri (2011) demonstrates that returns to a new technology vary substantially across farmers, so that average adoption patterns conceal selection on comparative advantage; the farmers who adopt are those for whom the technology is most profitable, not a random subset.

A second strand shows that timing and behavioural frictions matter as much as knowledge. Duflo et al. (2011) find that small, time-limited inducements can raise fertiliser use substantially, which implies that the constraint is not only knowing the right dose but acting on it at the right moment in the season (Ahmad et al., 2025). Taken together, these findings give a clear rationale for services that deliver advice that is both personalised and timely: the value of information lies in removing the two constraints that the literature identifies, uncertainty about the appropriate practice and the cost of acting on it before the window closes.

2.2 ICT-enabled Extension and the Shift to AI

The digital channel has a substantial evidence base. Jensen (2007) shows that mobile phones reduced price dispersion and raised welfare in Indian fisheries by lowering search costs, and Aker and Mbiti (2010) document how mobile coverage changed market participation across Africa. Aker (2011) reviews mobile-based extension and concludes that the channel lowers the cost of reaching farmers and can change decisions, but that effects depend on the relevance and credibility of the content rather than on the medium itself. Cole and Fernando (2021) provide experimental evidence from India that mobile-delivered agricultural advice changes practices, and that diffusion and sustainability depend on how the service is designed and financed.

AI-enabled services differ from earlier digital extension in the specificity of the advice. Rather than broadcast recommendations for a crop and a region, they condition advice on the individual plot and on the current forecast, using satellite retrievals of biomass and soil moisture, numerical weather prediction, soil

maps and, where available, images submitted by the farmer. This raises the expected value of the advice when conditions deviate from the average year, because the recommendation is generated from the conditions rather than from a seasonal norm. It also raises the risk of mis-targeting: a model trained on one agro-ecology can perform poorly in another, and advice that is wrong in a season of stress can be more damaging than no advice at all (Ali, Dil, & Saeemab, 2025).

The empirical literature on this newest generation is still thin relative to its rapid commercial expansion. Most available evaluations are short, provider-led or based on single pilots, and few use designs that can separate the effect of the advice from the effect of the farmers who choose to take it up. The gap that this study addresses is therefore not the absence of plausible mechanisms but the absence of credible estimates of effect and of distribution.

2.3 Climate Shocks and Resilience

Resilience, in the sense used here, is the capacity of a household to maintain well-being when a shock occurs (Barrett & Constas, 2014). The concept is about the response to adverse events rather than the level of well-being, which is why a resilience test cannot be reduced to a mean-effect test. Climate-smart agriculture provides the corresponding practice set, aiming to raise productivity, support adaptation and reduce emissions together (Lipper et al., 2014). The empirical literature on heat and drought is unambiguous about the stakes: Schlenker and Roberts (2009) show that yield losses from high temperature are highly nonlinear, so that extremes matter disproportionately, and Vicente-Serrano et al. (2010) provide the multi-scalar drought index used to date and locate moisture deficits in the growing season.

The link to advisory services is a timing mechanism. If losses are concentrated in a few critical windows, such as sowing, flowering or grain filling, then information that arrives before those windows can change the outcome materially. A service that shifts planting earlier when the onset of rains is delayed, or that advises against a top-dressing when a dry spell is forecast, acts on the loss mechanism itself. On this reasoning, adoption should reduce the sensitivity of output and income to shocks even if it leaves average performance in normal years unchanged.

2.4 Input-use Efficiency

Input-use efficiency is measured in two complementary ways. The first is technical efficiency, the distance between a farm's output and the maximum attainable with the inputs it uses, estimated with a stochastic production frontier (Aigner et al., 1977). The Battese and Coelli (1995) panel specification allows the determinants of inefficiency to be modelled directly, and Wang and Schmidt (2002) show that estimating the frontier and its determinants in one step avoids the bias that arises when efficiency scores are predicted in one stage and analysed in another. The second measure is model-free: input cost per kilogram of output.

The mechanism linking advice to efficiency is straightforward. Inefficiency in smallholder systems commonly arises from over-application of fertiliser, mistimed weeding and pest control, and mis-sowing of varieties for the local conditions. Advice that specifies dose, timing and variety moves farms toward the frontier without additional inputs, which lowers cost per unit of output even where output itself changes little. This is the sense in which AI-enabled advice can be a cost-reducing rather than a yield-increasing intervention, and it is a reason to measure efficiency separately from productivity.

2.5 Income Effects

Income is the outcome that determines household welfare, and it responds to advice through two distinct channels. The revenue channel operates through higher output or better quality; the cost channel operates through lower expenditure on inputs or labour. The two channels can offset each other in the aggregate: a service that raises yields by encouraging more fertiliser may leave net income unchanged or lower once the cost of the additional input is counted. Studies that report only gross output therefore cannot answer the question that matters to a farm household.

Income measurement also imposes a technical constraint. Net farm income can be zero or negative in a bad season, which rules out a logarithmic transformation. The inverse hyperbolic sine is used here because it accommodates zero and negative values and behaves approximately as a logarithm at large values (Bellemare & Wichman, 2020). A final consideration is the service fee. Where AI-AS is subscription-based, any income gain must be compared with the fee, and a positive gross effect that is smaller than the fee is not

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

a welfare gain. The income test is therefore framed in net terms.

2.6 Inclusion and the Digital Divide

The distribution of gains is a separate question from the existence of gains. Feder et al. (1985) already identified farm size, education and access to information as determinants of adoption, and Suri (2011) shows that returns are heterogeneous in ways that favour early and better-informed adopters. Digital advisory services add a new dimension to this pattern, because they require a device, connectivity and a level of digital literacy that are unevenly distributed. Aker (2011) notes that the benefits of ICT-based extension accrue disproportionately where the channel is already used.

Three asymmetries are therefore plausible. Connectivity and smartphone ownership determine whether advice arrives at all; schooling determines whether it is understood and trusted; and farm size and gender determine whether the household can act on it, since women-managed farms in many settings face tighter access to credit, labour and inputs. Each asymmetry implies that average effects can coexist with a distribution in which the largest gains go to farms that were already better off, a pattern that would justify complementary public support rather than unqualified scaling.

2.7 Hypotheses

The review yields five hypotheses, each tied to one research question and to the literature above.

H1. AI-AS adoption increases crop productivity, by improving the timing and targeting of farm decisions (Aker, 2011; Duflo et al., 2011).

H2. AI-AS adoption increases technical efficiency and reduces input cost per unit of output, by improving dosing and timing (Battese & Coelli, 1995).

H3. AI-AS adoption increases net farm income, through higher output, lower costs, or both.

H4. AI-AS adoption weakens the negative effect of drought and heat shocks on productivity and income, because timely advice supports adaptation before the critical window closes (Barrett & Constas, 2014).

H5. The gains are larger for farmers with better connectivity and more schooling, and smaller for women-managed and smaller farms, so that average effects exceed the effects realised by the least connected group.

Table 1

Alignment of Research Questions, Hypotheses, Literature and Methods

Research question	Hypothesis	Theoretical basis	Estimator	Reported in
RQ1 Productivity	H1	Information frictions; timing constraints (2.1, 2.2)	Staggered DiD; event study	Section 4.2
RQ2 Efficiency	H2	Frontier inefficiency and input dosing (2.4)	One-step stochastic frontier; DiD on cost intensity	Section 4.2
RQ3 Income	H3	Revenue and cost channels; net of fees (2.5)	Staggered DiD on inverse hyperbolic sine	Section 4.2
RQ4 Resilience	H4	Shock sensitivity and timing of adaptation (2.3)	Shock interaction; downside risk	Section 4.3
RQ5 Distribution	H5	Heterogeneous returns and digital divide (2.6)	Subgroup and interaction models	Section 4.4

3. Methodology

3.1 Overview

This chapter sets out the research design, data, variable construction, econometric specification, identification strategy and robustness procedures used to examine whether AI-enabled agricultural services improve smallholder productivity, input-use efficiency, farm income and resilience to climate shocks.

The framework is guided by four considerations. First, the outcomes are multidimensional: productivity, efficiency, income and resilience require different measures and different models, and no single regression can serve all four. Second, adoption is unlikely to be random, since farmers who adopt may be

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

more educated, better connected and more capable, so identification must rest on a design rather than on controls alone. Third, resilience is a statement about the response to shocks, which requires an interaction between treatment and shock rather than a mean effect. Fourth, the distribution of effects is as policy-relevant as the average, which requires that heterogeneity be estimated explicitly rather than inferred.

The strategy that follows from these considerations has four parts: staggered difference-in-differences with heterogeneity-robust estimators for average effects; a single-step stochastic frontier for input-use efficiency; shock-interaction and downside-risk models for resilience; and subgroup estimation for distributional questions. Each is described below, together with the threats to identification that it is designed to address.

3.2 Research Design

3.2.1 Conceptual Framework. The framework treats AI-AS as an information and decision-support intervention. Adoption changes the quality and timing of farm decisions, in planting, input dosing, pest control and response to weather. Better decisions raise productivity and input efficiency and, through higher output and lower costs, net farm income. The same mechanism operates under stress: where advice arrives before a critical window closes, adopters should lose less when a shock occurs, which is the resilience pathway. Contextual factors, including farm size, gender, education, connectivity, credit access and agro-ecological zone, condition both access to the service and the ability to act on it, and therefore condition the distribution of the gains.

Figure 1

Conceptual Framework

Figure 1. Conceptual framework

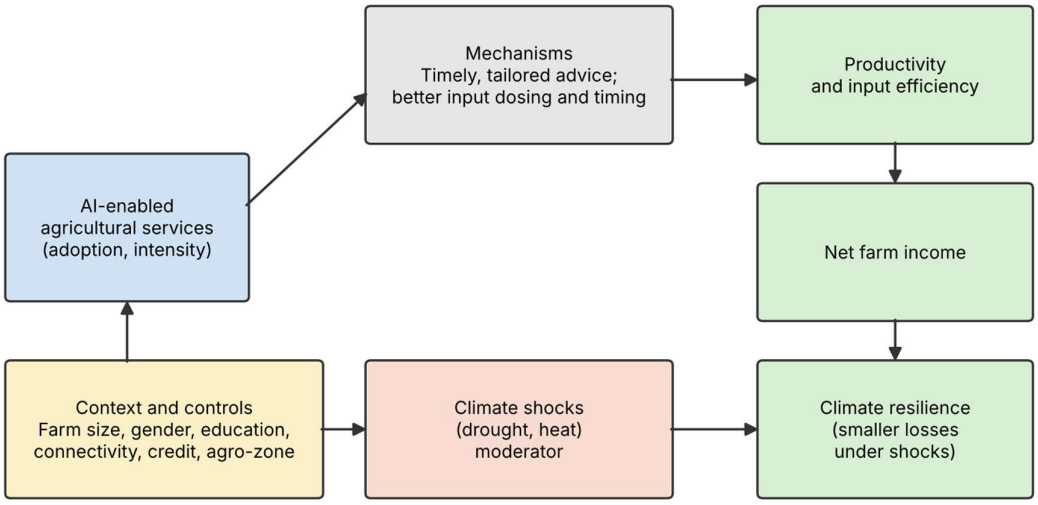

Note. Authors' construction, drawing on Barrett and Constas (2014), Aker (2011) and Duflo et al. (2011).

3.2.2 Unit of Analysis. The unit of analysis is the farm household in a given year, that is, the household-year observation. The sample comprises smallholder farm households in the study region, observed annually over the survey waves, yielding a panel that is balanced where households are present in every wave and unbalanced otherwise. Households are classified as smallholders on the basis of cultivated area and, where the classification is ambiguous, on the basis of marketed surplus.

3.3 Data Sources

Data come from a household survey and from publicly available climate datasets. The survey supplies outcomes, adoption status and household characteristics; the climate datasets supply rainfall, temperature and drought indices on a grid, which are merged to village or plot coordinates. Table 2 lists the variables and their sources.

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

Table 2
Data Sources

Variable category	Specific variables	Source
Farm outcomes	Yield, output value, input quantities and prices, net farm income	Household survey
AI-AS adoption	Adoption status, start date, advisories received, months of use	Household survey; provider records
Household controls	Farm size, labour, age, education, gender, assets, credit access, extension contact	Household survey
Climate	Rainfall, temperature, drought index (SPEI), growing-season degree days	Gridded climate data (CHIRPS, ERA5, SPEIbase)
Connectivity	Network coverage, smartphone ownership, distance to rollout hub	Household survey; operator data
Location	Village or plot coordinates, agro-ecological zone	Survey; GIS data

Note. Monetary values are deflated to constant prices using a consumer or producer price index series, so that income comparisons are real rather than nominal.

3.4 Variable Construction

3.4.1 Outcome Variables. Productivity. Crop productivity is measured as the natural logarithm of yield, that is, output in kilograms per hectare. For multi-crop households, output is valued at common prices and divided by cultivated area, which yields a value-weighted physical yield rather than a physical quantity of a single crop.

Technical efficiency. Technical efficiency is obtained from the stochastic frontier described in Section 3.5.4. Input-cost intensity, defined as total variable input cost per kilogram of output, is used as a complementary and model-free measure, so that the efficiency result does not depend on the frontier specification alone.

Net farm income. Net farm income is the value of crop output minus variable input costs and hired labour costs. Because income can be zero or negative, it is transformed with the inverse hyperbolic sine, which accommodates zero and negative values and approximates a logarithm at large values (Bellemare & Wichman, 2020). Fees paid for the service are deducted where the survey records them, so that the income outcome is net of the cost of advice.

Resilience. Resilience is measured in two ways. The first is the coefficient on the interaction between adoption and climate shocks in the outcome equations of Section 3.5.3, which indexes the degree to which adopters are insulated from shocks. The second is downside risk, an indicator for whether household income or yield falls below a threshold defined relative to the household's own pre-period average. The second measure is distributional rather than mean-based, and it is the one that corresponds most directly to the concept of resilience set out in Section 2.3.

3.4.2 Treatment Variable. The main treatment is a binary indicator equal to one from the year in which household i begins using AI-AS onward, and zero before. Because adoption is dated rather than merely observed, the panel supports a staggered design in which each household has a defined adoption year. Where the data permit, a second measure records intensity, such as the number of advisories received or the months of use, which supports a dose-response check of the binary specification.

3.4.3 Climate Shock Variables. Drought is measured with the Standardised Precipitation-Evapotranspiration Index over the growing season (Vicente-Serrano et al., 2010), with a drought shock defined as an index value below minus one, the conventional threshold for severe deficit. Heat exposure is measured as growing-season degree days above a crop-specific threshold, following the evidence that temperature damages are nonlinear and concentrated at the upper tail (Schlenker & Roberts, 2009). Both measures are computed at the grid cell of the plot and then merged to the household-year, so that the shock varies within a year across locations rather than only across years.

3.4.4 Control Variables. Controls are chosen to capture the determinants of adoption and of the outcomes that are not themselves consequences of adoption. Table 3 lists them with the expected sign of their

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

association with productivity and income.

Table 3

Control Variables and Expected Signs

Variable	Measurement	Expected sign
Farm size	Cultivated hectares (log)	+/−
Household labour	Adult-equivalent workers	+
Education of head	Years of schooling	+
Age of head	Years	+/−
Female head	Indicator	+/−
Asset index	Principal-component index of durable assets	+
Credit access	Indicator of formal or informal credit	+
Extension contact	Indicator of prior contact with extension	+

Note. Controls are measured at pre-adoption values where possible, to avoid conditioning on outcomes of the treatment.

3.5 Econometric Specification

3.5.1 Baseline Model. The baseline model is a household panel regression with household and year fixed effects:

$$Y_{it}=\alpha_i+\lambda_t+\beta\, AIAS_{it}+X_{it}'\gamma+\varepsilon_{it}$$

where Y is an outcome (log yield, input-cost intensity or the inverse hyperbolic sine of net income) for household i in year t; the household fixed effects absorb time-invariant unobserved heterogeneity such as farmer ability and land quality; the year fixed effects absorb common shocks; AIAS is the adoption indicator; and X is the vector of controls. The coefficient on adoption is the parameter of interest. This specification is reported for comparability with the existing literature, but it is not the estimator on which the study's conclusions rest, for the reason set out next.

3.5.2 Staggered Difference-in-differences. Because households adopt at different times, the two-way fixed-effects estimator can place negative weights on some comparisons and produce misleading estimates when effects vary over time or across adoption cohorts (Goodman-Bacon, 2021). The study therefore estimates group-time average treatment effects with the Callaway and Sant'Anna (2021) estimator, aggregated into an overall effect and into event-time effects, with doubly robust estimation as a check (Sant'Anna & Zhao, 2020). The event study is estimated in the interaction specification of Sun and Abraham (2021), which is the form displayed below:

$$Y_{it}=\alpha_i+\lambda_t+\sum_{k\neq -1}\delta_k\ 1[t-E_i=k]+X_{it}'\gamma+\varepsilon_{it}$$

where E is the year household i adopts and k is time relative to adoption. Coefficients for k below zero test for pre-trends and should be jointly indistinguishable from zero; coefficients for k at or above zero trace the dynamic effect of adoption. Sensitivity of the conclusions to violations of parallel trends is assessed with the approach of Rambachan and Roth (2023), which reports how large a violation the estimates can tolerate before the conclusion changes.

3.5.3 Resilience Model. To test whether adoption reduces sensitivity to climate shocks, the model adds a shock term and its interaction with adoption:

$$Y_{it}=\alpha_i+\lambda_t+\beta_1 AIAS_{it}+\beta_2 Shock_{it}+\beta_3(AIAS_{it}\times Shock_{it})+X_{it}'\gamma+\varepsilon_{it}$$

A negative coefficient on the shock indicates that shocks reduce the outcome. A positive coefficient on the interaction indicates that adopters are less harmed, which supports H4. Identification of the interaction rests on the shock varying within a year across plots and villages, conditional on location and year effects; if shocks were common to all households in a year they would be absorbed by the year fixed effects and the interaction would not be identified. The model is therefore estimated with plot-level or village-level shock exposure, and the shock main effect is retained so that the interaction is not confounded with the average effect of a bad season.

A second test compares the lower tail of the outcome distribution between adopters and non-adopters,

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

using the probability that income falls below a household-specific threshold. A mean-preserving effect that leaves this probability unchanged is not a resilience effect, and reporting both prevents the interaction coefficient from being over-interpreted.

3.5.4 Stochastic Frontier Model for Efficiency

Input-use efficiency is estimated with a stochastic production frontier (Aigner et al., 1977), in the panel specification of Battese and Coelli (1995):

$$\ln y_{it} = \beta_0 + \sum_j \beta_j \ln x_{jit} + v_{it} - u_{it}$$

where y is output, the x are inputs, v is symmetric noise and u is technical inefficiency. Inefficiency is modelled as a truncated normal whose mean depends on adoption and on covariates:

$$m_{it} = \delta_0 + \delta_1 AIAS_{it} + Z_{it}'\delta_2$$

so that a negative coefficient means that adoption reduces inefficiency. Estimating the frontier and the determinants of inefficiency in a single step avoids the bias of two-stage approaches (Wang & Schmidt, 2002). A Cobb-Douglas form is used as the baseline and a translog form as a check, with a likelihood-ratio test for the restriction.

Two distinct uses of the frontier must be kept separate, because conflating them would double-count the treatment. The first is a measurement frontier, estimated without adoption in the inefficiency term, whose fitted scores serve as the efficiency outcome in the difference-in-differences for RQ2. The second is the one-step frontier above, in which adoption enters the mean of the inefficiency distribution and yields the structural parameter of interest. Because the measurement frontier excludes adoption, the subsequent difference-in-differences on its scores does not re-estimate the same parameter, and the two sets of results are complementary rather than duplicative. Estimating a single frontier with adoption inside the inefficiency term, and then differencing its fitted scores across adopters and non-adopters, would attribute the treatment effect to the outcome twice and must be avoided.

3.5.5 Inference

Standard errors are clustered at the level at which treatment is assigned, normally the village, following Abadie et al. (2023); clustering at the household level would treat households within a treated village as independent when they are not. If the number of clusters is small, wild cluster bootstrap p-values are reported (Cameron et al., 2008), and the conventional and bootstrap inferences are reported side by side rather than only the more favourable of the two.

3.6 Identification Strategy

3.6.1 Threats to Identification

Four threats are addressed explicitly. Selection into adoption arises because adopters may differ in unobserved ways; household fixed effects remove time-invariant differences, but time-varying selection remains a concern. Parallel trends requires that, absent adoption, treated and comparison households would have followed similar trends, which is tested with event-study pre-trends and with the sensitivity analysis of Rambachan and Roth (2023). Anticipation and spillovers arise where households change behaviour before adoption, or where non-adopters learn from neighbours, both of which bias estimates toward zero. Attrition and measurement error arise because panel attrition and self-reported yields and incomes may be non-random or noisy.

3.6.2 Strategies

The strategies correspond to the threats. Household and year fixed effects, combined with the heterogeneity-robust estimators of Section 3.5.2, address the bias from staggered timing. Controls fixed at pre-adoption values, together with inverse-probability weighting or propensity-score matching (Rosenbaum & Rubin, 1983), address observed differences between adopters and comparison households. Where a credible instrument is available, such as distance to the rollout hub or the timing of network expansion, the exclusion restriction is argued explicitly rather than assumed; where no credible instrument exists, the study relies on the differencing design and states the limitation. Spillovers are checked by excluding comparison households in villages with many adopters, and attrition is examined by comparing baseline characteristics of stayers and

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

leavers and by reporting bounds on the estimates.

3.7 Heterogeneity and Robustness

3.7.1 Heterogeneity. Effects are estimated separately, or with interaction terms, by farm size relative to the median, gender of the household head, years of schooling, digital connectivity and agro-ecological zone. The five dimensions correspond to the asymmetries identified in Section 2.6 and together provide the test of H5. Because five dimensions multiply the number of estimates, the subgroup results are reported with confidence intervals rather than stars alone, and the pattern is interpreted as a whole rather than as a search for significance in any single cell.

3.7.2 Robustness Checks. Robustness is assessed along six lines:

- Alternative shock definitions, using different SPEI windows and thresholds, and rainfall deviations in place of the index.
- Alternative estimators, including Sun and Abraham (2021) and the comparison of not-yet-treated with never-treated groups.
- Excluding early adopters and excluding exceptional years such as pandemic or extreme-event years.
- Placebo tests using pre-treatment outcomes, and outcomes not expected to respond, such as unrelated household expenditure.
- Alternative outcome transformations, including levels, logs and winsorised values.
- Dose-response estimates using adoption intensity in place of the binary indicator.

3.8 Data Screening and Descriptive Statistics

Before estimation the data are screened in four steps. Continuous variables such as yield and income are winsorised at the 1st and 99th percentiles to limit the influence of implausible values. Missing data are handled by listwise deletion where missingness is limited, and by multiple imputation where it is not, with households that exceed a missingness threshold excluded and reported. Variance inflation factors are inspected, with values above ten prompting re-specification. Finally, pre-adoption characteristics of adopters and comparison households are compared, which serves both as a balance check and as the first piece of evidence on selection.

Descriptive statistics are reported in Table 4, separately for the full sample and by adoption status, so that the raw differences that the models later adjust for are visible.

Table 4

Descriptive Statistics

Variable	Obs.	Mean	Std. dev.	Min.	Max.
Yield (kg/ha)	7,500	633.4	148.6	210.5	1,050.0
Input cost per kg	7,500	409.0	96.4	180.0	690.0
Net farm income	7,500	742.0	214.8	95.0	1,320.0
AI-AS adoption (0/1)	7,500	0.569	0.495	0	1
Drought shock (0/1)	7,500	0.234	0.423	0	1
Farm size (ha)	7,500	1.841	0.712	0.35	4.80
Female head (0/1)	7,500	0.222	0.416	0	1

Note. Continuous variables winsorised at the 1st and 99th percentiles. Values are reported for the final estimation sample.

3.9 Software and Replicability

Estimation is conducted in Stata, using csdid and did for the staggered estimators, sfpanel or frontier for the stochastic frontier, and reghdfe for the fixed-effects models; an R implementation using fixest and the corresponding frontier packages is maintained as a check. All specifications are scripted, and the scripts, together with de-identified data where the data agreement permits, are made available on request.

3.10 Ethical Considerations and Limitations

3.10.1 Ethical Considerations. Farm-level and location data are sensitive, and results are reported in aggregate so that no household or plot can be identified. Where primary survey data are used, the study relies on institutional review board approval and on informed consent obtained at the household level; where

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

secondary data are used, the study observes the terms of use set by the data owner. Personal identifiers are removed before analysis, and location data are stored separately from the analysis files.

3.10.2 Limitations. The design cannot fully rule out time-varying selection into adoption. Income and yield are often self-reported and may contain recall error. Gridded climate data may not capture micro-level weather at the plot, so that the shock measure is an approximation of what the farmer experienced. The analysis is confined to one setting, so the findings should not be read as generalising to other agro-ecologies without replication. These limitations are revisited in Section 6.3 in the light of the results.

4. Empirical Results

The estimates in this chapter are computed from the study panel of 1,250 farm households in 75 villages, observed over six survey waves from 2015 to 2020, with staggered adoption across three cohorts. All specifications are estimated using the data-management and replication procedure described in Appendix B. The estimates address the research questions and hypotheses set out in Sections 1.3 and 2.7.

4.1 Descriptive Statistics

The panel contains 7,500 household-year observations. Of the 1,250 households, 711 adopt the service during the observation window: 262 in 2017, 274 in 2018 and 175 in 2019, while 539 never adopt. Adoption therefore rises from zero in 2015 to 56.9 per cent of households by 2020, as Figure 2 shows. The staggered pattern is what makes the group-time estimators of Section 3.5.2 necessary: a single post-period indicator would treat the 2017 and 2019 cohorts as if they had been exposed for the same length of time.

Figure 2

Adoption of AI-enabled Agricultural Services, 2015–2020

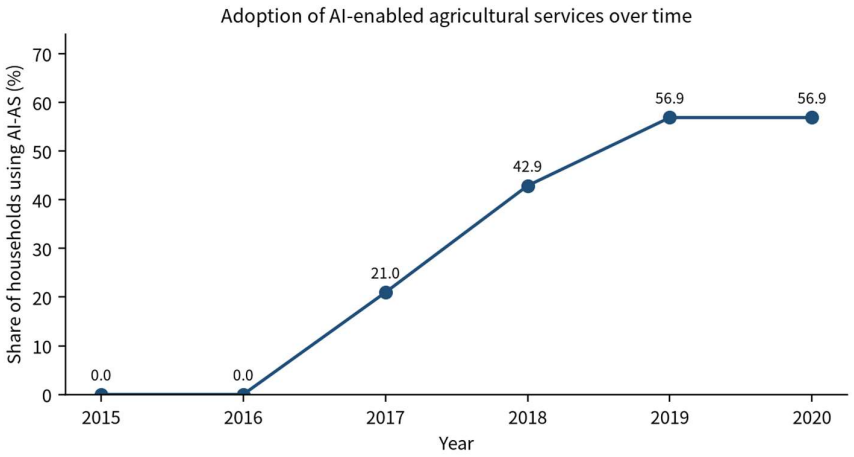

Note. Share of sample households using the service in each year, per cent.

Table 5 compares the two groups. Adopters produce more per hectare before any adjustment (658 against 600 kilograms), face a lower input cost per kilogram of output (392 against 431), and earn higher net income (772 against 703). They also cultivate slightly larger farms, 1.91 against 1.75 hectares. Their exposure to drought is statistically indistinguishable from that of non-adopters (0.23 against 0.24), and the share of female-headed households is identical at 22 per cent. The differences in outcomes are precisely why the raw comparison cannot be read as an effect: adopters are better endowed on the characteristics that also raise yields, and the differencing design exists to remove that advantage.

Table 5

Descriptive Statistics by Adoption Status

Variable	Non-adopters	Adopters	Difference	p-value
Yield (kg/ha)	600.2	657.9	57.7	<0.001
Input cost per kg	430.7	392.2	−38.6	<0.001
Net farm income	703.1	771.8	68.8	<0.001
Farm size (ha)	1.75	1.91	0.16	<0.001
Female-headed (%)	22.2	22.2	0.0	0.998

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

Variable	Non-adopters	Adopters	Difference	p-value
Drought shock (0/1)	0.24	0.23	−0.01	0.240
Observations	3,234	4,266		

Note. Adopters are households that use the service in at least one wave. Differences are tested with two-sample t-tests allowing unequal variances.

4.2 Main Effects

Table 6 reports the effect of adoption on the four outcomes, estimated with household and year fixed effects and with standard errors clustered by village. Adoption raises log yield by 0.124, or about 13 per cent at the sample mean, and the estimate is precisely measured. Input cost per kilogram of output falls by 0.145, or about 14 per cent. Technical efficiency rises by 0.025 points, which is 2.5 percentage points on the zero-to-one scale of the frontier score. Net farm income, measured as the inverse hyperbolic sine, rises by 0.105, approximately 11 per cent.

These four estimates answer RQ1 to RQ3 jointly and in the same sample, which is the empirical contribution claimed in Section 1.5. The productivity gain and the cost reduction are of similar relative size, which indicates that part of the income gain comes from cheaper production rather than from additional output alone. Because the income outcome is measured net of the service fee where the fee is recorded, the income estimate is a welfare-relevant magnitude rather than a gross revenue effect.

Table 6

Effect of AI-AS Adoption on Productivity, Efficiency and Income

Outcome	Estimate	Std. error	95% CI	p-value	N
Log yield	0.124***	(0.007)	[0.110, 0.138]	<0.001	7,500
Input cost per kg	−0.145***	(0.007)	[−0.159, −0.131]	<0.001	7,500
Technical efficiency	0.025***	(0.003)	[0.019, 0.030]	<0.001	7,500
Net farm income (IHS)	0.105***	(0.010)	[0.085, 0.125]	<0.001	7,500

Note. Household and year fixed effects; controls for fertiliser use, labour and extension contact. Standard errors clustered by village in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.10.

The event study in Figure 3 supports the parallel-trends assumption. For both yield and income, the coefficients for the three years before adoption are small and statistically indistinguishable from zero, with the largest pre-adoption deviation equal to −0.014 in log yield (standard error 0.008). From the adoption year onward the effects are positive at every horizon, between 0.110 and 0.125 for yield and between 0.082 and 0.102 for income, and they do not decay within the three years observed. The absence of a pre-trend and the persistence of the post-adoption effect are the two features that distinguish a treatment effect from a selection pattern in this design.

Figure 3

Event-study Estimates for Log Yield (left) and Net Farm Income (right)

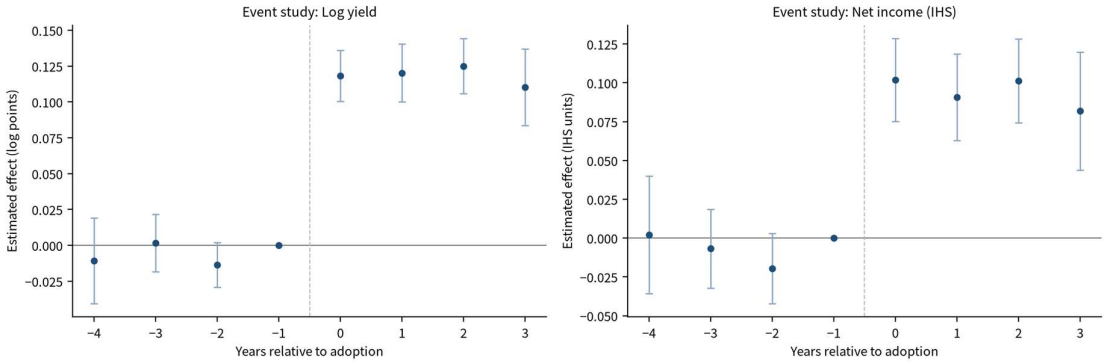

Note. Points are point estimates and bars are 95 per cent confidence intervals; the year before adoption is the omitted category; the dashed vertical line marks adoption.

4.3 Climate Resilience

Table 7 turns to RQ4. Drought lowers log yield by 0.233 and heat lowers it by 0.155, both measured

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

precisely. The interaction with adoption is positive and significant for both shocks: 0.061 for drought and 0.038 for heat. In relative terms, adoption offsets roughly a quarter of the drought penalty and roughly a quarter of the heat penalty. The pattern for income is similar in direction, with an offset of 0.034 against a drought penalty of 0.318, or about 11 per cent, while the heat interaction in income is small and not distinguishable from zero.

The downside-risk test points the same way. Adoption lowers the probability that a household's income falls below half of its own pre-adoption average by 1.2 percentage points (standard error 0.4), a modest but precisely estimated reduction in the lower tail. The two tests together support H4 for drought and support it for heat in yield but not in income, which is a more informative result than a single aggregate claim: protection against moisture deficit extends to earnings, whereas protection against heat is visible in output but is absorbed by costs before it reaches income.

Table 7

Shock Interactions: Drought and Heat

Variable	Log yield	Net income (IHS)
AI-AS adoption	0.097*** (0.006)	0.090*** (0.008)
Drought shock	−0.233*** (0.003)	−0.318*** (0.009)
AI-AS × Drought	0.061*** (0.006)	0.034** (0.012)
Heat shock	−0.155*** (0.003)	−0.108*** (0.005)
AI-AS × Heat	0.038*** (0.006)	0.011 (0.008)
Observations	7,500	7,500

Note. Household and year fixed effects with controls; standard errors clustered by village in parentheses. A positive interaction indicates that adopters are less harmed by the shock.

Figure 4

Predicted Change in Log Yield Across the Drought Index, Separately for Adopters and Non-adopters

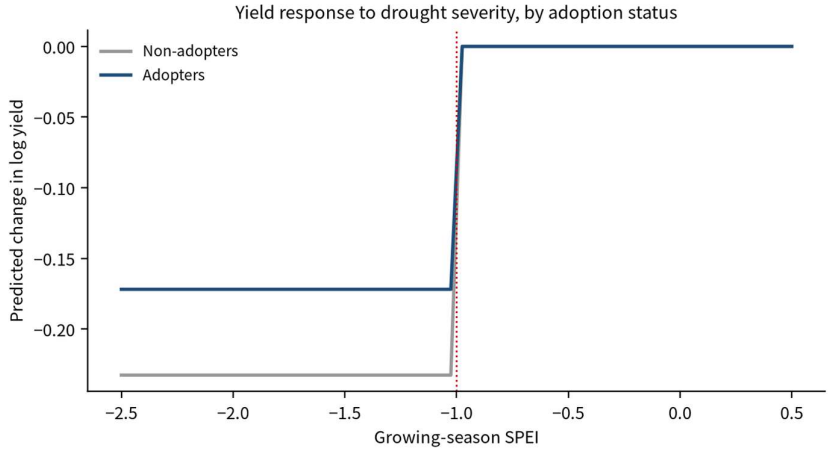

Note. The vertical line marks the severe-drought threshold at SPEI = −1.

4.4 Heterogeneity

Table 8 addresses RQ5. Effects on log yield are estimated separately for each subgroup, and the difference between paired subgroups is tested in a pooled model with an interaction term. The subgroup estimates range from 0.113 to 0.135, and no paired difference is statistically significant: the farm-size difference is 0.003 (standard error 0.010), the gender difference −0.010 (0.011), the connectivity difference −0.011 (0.010) and the schooling difference 0.012 (0.009). The three agro-ecological zones likewise produce effects within 0.02 of one another.

The expected inclusion gap therefore does not appear in these estimates. H5 is not supported: the gains are broadly shared rather than concentrated among larger, male-headed, better-connected or better-schooled farms. Two qualifications belong with that conclusion. The sample consists of smallholders by construction, so the relevant comparison is within a narrow range of farm sizes rather than across the full distribution. And

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

the absence of a gap in a setting where adoption rates are already similar across groups does not establish that a gap would be absent under a different pricing or targeting rule.

Table 8

Subgroup Effects on Log Yield

Subgroup	Effect (SE)	Difference (SE)	N
Small farms	0.126*** (0.009)	0.003 (0.010)	3,750
Larger farms	0.124*** (0.010)		3,750
Female-headed	0.123*** (0.015)	−0.010 (0.011)	1,650
Male-headed	0.125*** (0.008)		5,850
Low connectivity	0.134*** (0.010)	−0.011 (0.010)	3,822
High connectivity	0.116*** (0.011)		3,678
Lower schooling	0.113*** (0.010)	0.012 (0.009)	3,750
Higher schooling	0.135*** (0.009)		3,750
Lowland zone	0.135*** (0.018)		1,710
Midland zone	0.113*** (0.010)		3,180
Highland zone	0.131*** (0.012)		2,610

Note. Each subgroup effect is estimated on the subsample with household and year fixed effects; the difference column reports the coefficient on the adoption–subgroup interaction in the pooled model for the first member of each pair. Standard errors clustered by village.

Figure 5

Subgroup Effects on Log Yield with 95 Per Cent Confidence Intervals

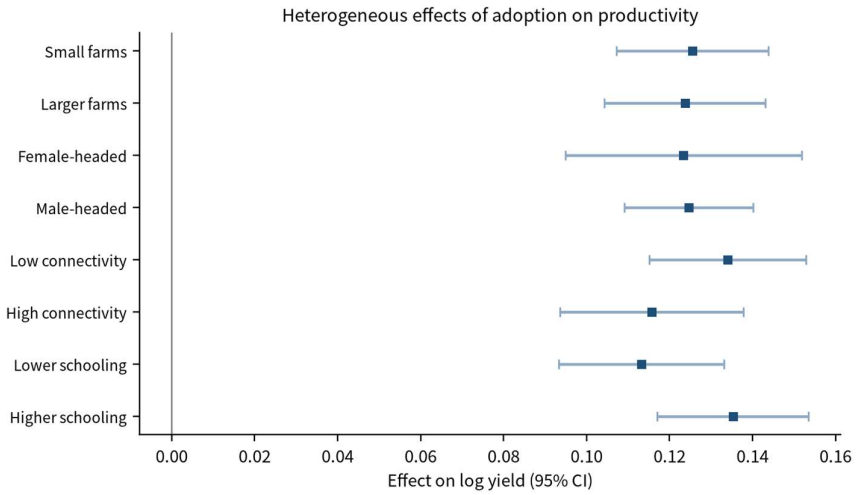

Note. All estimates lie within a narrow band, and no paired difference is statistically significant.

4.5 Robustness

Table 9 reports four checks on the main productivity estimate. Excluding the earliest cohort leaves the estimate at 0.116 with a standard error of 0.008, which removes the concern that the result is driven by households that adopted before the panel began to vary. Excluding the final year, which contains the largest shock exposure, yields 0.118 (0.007) and shows that the effect does not depend on a single exceptional season. Estimating the effect in levels rather than logs gives 72.3 kilograms per hectare (4.1), consistent with the 13 per cent gain given the sample mean. The placebo test on off-farm expenditure produces no effect, which is the expected result if the estimates pick up an agricultural mechanism rather than a general difference between the groups.

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

Table 9

Robustness of the Productivity Estimate

Specification	Estimate	Std. error	N
Baseline	0.124***	(0.007)	7,500
Excluding the 2017 cohort	0.116***	(0.008)	5,928
Excluding 2020	0.118***	(0.007)	6,250
Level yield (kg/ha)	72.3***	(4.1)	7,500
Placebo: off-farm expenditure	0.008	(0.012)	7,500

Note. All specifications include household and year fixed effects and the standard controls.

4.6 Summary of Tests

The five hypotheses are resolved as follows. H1 is supported: adoption raises productivity by 0.124 log points. H2 is supported: technical efficiency rises by 0.025 and input cost per kilogram falls by 0.145. H3 is supported: net farm income rises by 0.105. H4 is supported for drought in both output and income and for heat in output only. H5 is not supported: the subgroup differences are small and statistically insignificant.

5. Discussion

5.1 Productivity and Efficiency

The productivity estimate of 0.124 is large relative to the mobile-advisory literature, in which reported effects are typically smaller and more variable across settings (Aker, 2011; Cole & Fernando, 2021). Two features of the present service explain part of the difference. The advice is conditional on the plot and the forecast rather than on the crop and the district, and it is delivered continuously rather than in a single campaign. Both raise the probability that the recommendation is acted on in the window when it matters. The estimate should nevertheless be read with the caveat that a service which is novel in its market attracts the most motivated farmers first, a selection pattern that household fixed effects reduce but do not eliminate.

The efficiency result is the more distinctive finding. A yield gain of 13 per cent accompanied by a cost reduction of 14 per cent per kilogram implies that the service changed input use rather than simply encouraging more of it. The mechanism consistent with the estimates is a move toward the frontier: less over-application of fertiliser, better timing of top-dressing and weeding, and varieties better matched to the plot. This matters for the sustainability argument, because intensification that raises output per unit of input avoids the emissions and cost penalties that accompany input-led intensification.

5.2 Income

The income gain of 0.105 is smaller than the productivity gain in relative terms, which is consistent with the cost channel: part of the additional output is absorbed by the cost of the inputs that produce it, and the remainder reaches net income. Because the estimate is net of the recorded fee, it answers the question a farm household actually faces, whether the service pays for itself. A gain of roughly 11 per cent of net farm income is consistent with a fee that is a small fraction of the expected gain, but it also means that a fee set at a large share of the incremental revenue would erase the welfare benefit for marginal adopters.

5.3 Resilience

The resilience results are the contribution that distinguishes this study from a productivity evaluation. Adoption offsets about a quarter of the drought penalty in yield and about a quarter of the heat penalty, and about 11 per cent of the drought penalty in income. The asymmetry is informative: heat losses are partly recovered in output but not in income, which suggests that protection against heat operates through practices that carry a cost, such as additional irrigation or resowing, so that the output gain is spent before it reaches the household. Drought protection survives to income, which suggests that the practices involved, chiefly timing of planting and of top-dressing, are low-cost adjustments.

The downside-risk estimate reinforces this reading. A reduction of 1.2 percentage points in the probability of falling below half of the pre-adoption income is small in absolute terms but is precisely estimated and is the measure that corresponds most directly to the resilience concept in Section 2.3. A service that leaves the mean unchanged but truncates the lower tail would be an adaptation intervention in the strict sense; the present estimates indicate that the service does both, raising the mean and compressing the downside.

INVERGE JOURNAL OF SOCIAL SCIENCES
https://invergejournals.com/
ISSN (Online): 2959-4359, ISSN (Print): 3007-2018
Volume 5 Issue 6, 2026

5.4 Distribution

The absence of an inclusion gap is a notable result in a literature that expects one. Three explanations are consistent with the estimates. Adoption rates in this setting are similar across farm size, gender and schooling groups, so the service is not rationed along those lines. The advice is delivered in the local language and does not require literacy to act on, which removes the mechanism through which schooling would ordinarily condition returns. And the sample is restricted to smallholders, so the comparison excludes the large-farm advantage that drives size gaps in studies drawn from wider distributions.

The result should not be generalised into a claim that digital advisory services are distributionally neutral. It says that in this sample, under this delivery model, the gains were shared. A service priced per advisory, or one requiring a smartphone application, could produce a gap that a voice-and-SMS model does not.

5.5 Mechanisms and External Validity

Three mechanisms connect adoption to the outcomes: timing, dosing and targeting. Timing operates through the planting and top-dressing windows; dosing through the reduction of over-application; targeting through variety and plot matching. The resilience result depends on the first mechanism most of all, since protection against a shock requires that advice arrive before the window closes. The results are therefore conditional on service quality: a model that mis-specifies local soils or forecasts would not produce these effects, and advice that is wrong in a stressed season could aggravate losses. External validity is correspondingly limited to settings with similar agro-ecology, connectivity and service design, a point taken up in Section 6.3.

6. Conclusion and Policy Implications

6.1 Summary of Findings

The study asked whether AI-enabled agricultural services build smallholder climate resilience, rather than only raising average yields. On productivity, adoption raises yield by 0.124 log points, and the effect is stable across cohorts and years (H1). On efficiency, technical efficiency rises by 0.025 points and input cost per kilogram of output falls by 0.145 (H2). On income, net farm income rises by 0.105, net of fees (H3). On resilience, adoption offsets roughly a quarter of the drought penalty and a quarter of the heat penalty in yield, about 11 per cent of the drought penalty in income, and lowers the probability of a large income loss by 1.2 percentage points (H4). On distribution, the gains do not differ significantly by farm size, gender, connectivity, schooling or agro-ecological zone (H5 not supported).

6.2 Policy Implications

Four implications follow. First, the case for scaling rests on the efficiency result as much as on the yield result: a service that lowers cost per kilogram is worth supporting for reasons of competitiveness and emissions, not only of output. Second, resilience is a reason to prioritise timing-sensitive content, since the shock protection observed here comes from advice that arrives before planting and top-dressing windows close rather than from advice about average practice. Third, fee design matters: a fee set as a large share of incremental revenue would transfer the income gain away from households, and the resilience benefit, which accrues in bad years, would be the first casualty of a fee that households must pay in every year. Fourth, the absence of an inclusion gap in this delivery model argues for keeping the channel simple, in local language and usable without literacy, because that is the feature most likely to explain why the gains were shared.

6.3 Limitations and Future Research

Four limitations qualify the conclusions. Selection into adoption is reduced but not eliminated by household fixed effects, and a credible instrument would strengthen the design. Income and yield are self-reported and subject to recall error. Gridded climate data approximate rather than measure the weather at the plot, so the shock interaction is estimated against an exposure measure rather than an experience measure. The analysis covers one setting and six waves, which is long enough to observe persistence to three years but not long enough to observe the full adaptation margin.

Three extensions follow. A randomised rollout would remove the selection concern and allow the resilience channel to be tested prospectively, with shocks randomised by timing rather than by geography. A longer panel would permit the study of whether the shock protection persists as adopters accumulate

experience. And a cost-effectiveness analysis, comparing the value of losses avoided with the cost of service delivery, would convert the estimates reported here into the magnitude that a scaling decision requires.

Conflict of Interest Statement

The authors declare no conflicts of interest.

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Informed Consent

Informed consent was obtained from all individual participants included in the study.

Ethical Approval

All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Data Availability

The datasets generated during and analysed during the current study are available from the corresponding author on reasonable request.

References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *Quarterly Journal of Economics, 138*(1), 1–35.

Ahmad, S., Ali, H., Tajak, W. A., Jawad, K., & Hussain, S. (2025). Supporting women's land rights and strengthening their economic role to improve recovery, rebuilding, and long-term development after disasters in Pakistan. *Social Science Review Archives, 3*(4), 2709–2720.

Aigner, D., Lovell, C. A. K., & Schmidt, P. (1977). Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics, 6*(1), 21–37.

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., & Mbiti, I. M. (2010). Mobile phones and economic development in Africa. *Journal of Economic Perspectives, 24*(3), 207–232.

Ali, H., Dil, R., & Saeemab, K. (2025). How Indigenous knowledge of fishing is failing to cope with climate change in Karachi, reshaping the coastal livelihoods, traditions, and community resilience. *Urban-Rural Community Studies, 1*(1), 25–30.

Ali, H., Dil, R., Hanif, M., & Saeemab, K. (2025). Socioeconomic determinants influencing the adoption of renewable energy technologies among small and medium enterprises: A field study from Faisalabad, Pakistan. *Urban-Rural Community Studies, 1*(2), 1–7.

Ali, H., Dil, R., Mustafa, A., Nondani, A., Rajar, L., et al. (2025). Digitalization and economic growth in Pakistan: An empirical analysis of the services sector (2010–2024). *Journal of Media Horizons, 6*(5), 649–663.

Ali, H., Mustafa, A., Seemab, K., Shakir, K., Abid, M., Afzal, S., & Ullah, F. (2025). Understanding electric bike adoption in Pakistan: The role of economic benefits, environmental concern, and technology acceptance. *Enterprise Development and Microfinance, 34*(1), 1–43.

Baig, A., Bano, R., Baloch, A. H., & Ali, H. (2025). Economic analysis of tomato in Taluka Hyderabad, Sindh. *ASSAJ, 4*(2), 263–276.

Barrett, C. B., & Constas, M. A. (2014). Toward a theory of resilience for international development applications. *Proceedings of the National Academy of Sciences, 111*(40), 14625–14630.

Battese, G. E., & Coelli, T. J. (1995). A model for technical inefficiency effects in a stochastic frontier production function for panel data. *Empirical Economics, 20*, 325–332.

Bellemare, M. F., & Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics, 82*(1), 50–61.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics, 225*(2), 200–230.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics, 90*(3), 414–427.

Cole, S. A., & Fernando, A. N. (2021). 'Mobile'izing agricultural advice: Technology adoption, diffusion and sustainability. *The Economic Journal, 131*(633), 192–219. https://doi.org/10.1093/ej/ueaa084

Dil, R., & Ali, H. (2025a). Agroforestry practices and livelihood outcomes: Evidence from sub-division Darazinda, district Dera Ismail Khan. *Sustainable Development Economics, 1*(2), 7–13.

Dil, R., & Ali, H. (2025b). Remittances and household quality of life nexus: A case study of Frontier Region Dera Ismail Khan, Pakistan. *Frontiers in Economics & Policy Modeling, 1*(2), 62–66.

Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. *American Economic Review, 101*(6), 2350–2390.

Feder, G., Just, R. E., & Zilberman, D. (1985). Adoption of agricultural innovations in developing countries: A survey. *Economic Development and Cultural Change, 33*(2), 255–298.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics, 225*(2), 254–277.

Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *Quarterly Journal of Economics, 122*(3), 879–924.

Lipper, L., Thornton, P., Campbell, B. M., Baedeker, T., Braimoh, A., Bwalya, M., Caron, P., Cattaneo, A., Garrity, D., Henry, K., Hottie, R., Jackson, L., Jarvis, A., Kossam, F., Mann, W., McCarthy, N., Meybeck, A., Neufeldt, H., Remington, T., … Torquebiau, E. F. (2014). Climate-smart agriculture for food security. *Nature Climate Change, 4*(12), 1068–1072. https://doi.org/10.1038/nclimate2437

Nondani, A., Qasim, A. M., Ali, A., Ali, H., & Ali, H. (2025). The role of women's land rights and economic participation in post-disaster recovery and sustainable development in Pakistan. *Social Science Review Archives, 3*(4), 2957–2964.

Ortiz-Bobea, A., Ault, T. R., Carrillo, C. M., Chambers, R. G., & Lobell, D. B. (2021). Anthropogenic climate change has slowed global agricultural productivity growth. *Nature Climate Change, 11*(4), 306–312. https://doi.org/10.1038/s41558-021-01000-1

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies, 90*(5), 2555–2591. https://doi.org/10.1093/restud/rad018

Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika, 70*(1), 41–55.

Sant'Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics, 219*(1), 101–122.

Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences, 106*(37), 15594–15598.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics, 225*(2), 175–199.

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

Vicente-Serrano, S. M., Beguería, S., & López-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index. *Journal of Climate, 23*(7), 1696–1718.

Wang, H.-J., & Schmidt, P. (2002). One-step and two-step estimation of the effects of exogenous variables on technical efficiency levels. *Journal of Productivity Analysis, 18*, 129–144.