Learning to Leap: Platform Sequencing and Productivity in Emerging Market Agriculture

Soniya Gupta-Rawal[1]
PhD Candidate, Judge Business School, University of Cambridge, United Kingdom
ORCID: 0000-0002-1344-3105
+44 (0) 7435 775162
sg2001@jbs.cam.ac.uk

Paul Kattuman
Professor of Economics and Director of Studies in Management
Judge Business School, University of Cambridge, United Kingdom
+44 (0)1223 764136
p.kattuman@jbs.cam.ac.uk

Luisa Corrado
Professor of Economics
Department of Economics, University of Rome Tor Vergata, Italy
Visiting Research Fellow, Judge Business School, University of Cambridge, Cambridge, UK
luisa.corrado@uniroma2.it

Jaideep Prabhu
Professor of Marketing
Judge Business School, University of Cambridge, United Kingdom
+44 (0)1223 765468
j.prabhu@jbs.cam.ac.uk

Submitted for NBER SI26 Summer Institute 2026
Under the Entrepreneurship Workshop (Productivity, Innovation, and Entrepreneurship) and Development Economics

---

This document is an output from the research initiative `Private Enterprise Development in Low-Income Countries' (PEDL), a programme jointly funded by the Centre for Economic Policy Research (CEPR) and the Foreign, Commonwealth and Development Office (FCDO), contract reference PEDL_LOA_101024_Gupta-Rawal. The views expressed are not necessarily those of CEPR or FCDO.

[1] Correspondence regarding this article should be addressed to Soniya Gupta-Rawal, Judge Business School, University of Cambridge, Email: sg2001@jbs.cam.ac.uk

Learning to Leap: Platform Sequencing and Productivity in Emerging Market Agriculture

Abstract

New digital platforms in emerging markets often fail not for lack of demand but because adoption barriers prevent users from accessing high-value services. We study how staged adoption resolves this problem, using a setting in which Indian farmers transition from WhatsApp-based crop advisories to a satellite-driven AI platform without requiring intermediate technologies. Using a daily farm-level panel of 904 farms and satellite-based vegetation indices as objective measures of productivity, we estimate event-time and regression kink designs around two discrete adoption events. Registration with the messaging service produces modest, immediate productivity gains. Enrollment in the AI platform produces larger, persistent gains, with NDVI slope effects that compound over the post-enrollment period. Returns scale with farm size and nearly double under drought conditions, identifying an information-value mechanism: precise, timely advice matters most when input misallocation is most costly. Small farms face short-run adjustment costs before recovering; medium and large farms realize stable improvements from enrollment onward. The findings provide causal evidence that digitizing agronomic information generates measurable farm-level productivity gains, that staged platform design shapes the size and timing of those gains, and that returns are heterogeneous in ways that have direct implications for how digital ventures price, sequence, and target services in low-income agricultural markets.

Keywords: Digitization; Entrepreneurship; Agricultural productivity; AI adoption; Emerging markets; Innovation diffusion; Satellite data.

JEL Codes: O33 (Technological Change), O13 (Agriculture), Q16 (Agricultural R&D), L86 (Information and Internet Services), O12 (Microeconomic Analyses of Economic Development), D83 (Search, Learning, Information).

1. Introduction

Digital platforms in emerging markets face a structural problem: the services most likely to raise productivity are also the ones hardest to adopt. Users lack the habits, trust, and complementary skills to extract value from complex tools, and platform designers face a cold-start problem where the highest-value tier is unreachable without first building a base. One solution is staged adoption, where a low-friction entry channel acquires and habituates users before a higher-capability platform is introduced. Whether this sequencing produces productivity gains, and for whom, is an open empirical question.

We study this question in Indian agriculture, where the agri-tech startup Farmonaut operates a two-stage digital advisory service. Farmers first register for WhatsApp-based crop advisories and later enroll in JEEVN AI, a satellite-driven platform that delivers field-level analytics, irrigation scheduling, and risk alerts. The transition from messaging to an AI portal is the empirical footprint of digital leapfrogging: users move from near-zero formal decision support to a full-stack AI service without passing through intermediate technologies (Antonelli, 1991; Mody and Dahlman, 1992). Indian agriculture provides a credible setting for this analysis. The sector employs roughly 43 percent of the workforce and contributes about 18 percent of GDP, yet production is constrained by erratic monsoons, fragmented landholdings, and limited access to timely, farm-specific advice (Mulla, 2013; Majumdar and Kumar, 2020). These conditions create both high demand for information services and strong barriers to adopting them, making the returns to platform sequencing large in principle and observable in practice.

The empirical challenge is isolating the productivity effect of each adoption stage. Prior work on technology adoption in developing economies documents correlations between digital tool use and outcomes but rarely achieves causal identification, particularly when adoption unfolds in steps and self-selection into each stage is plausible (Zhu and Kraemer, 2005; Heeks, 2022). A further limitation is reliance on self-reported outcomes, which introduces measurement error and recall bias. We address both problems by using daily satellite-based vegetation indices as objective, high-frequency productivity outcomes and by exploiting discrete adoption events, registration, and AI enrollment as identification cutoffs in event-time and regression kink designs. The setup delivers two separate treatment effects: a level effect at registration and a

slope effect at AI enrollment, each identified from local variation around the adoption date within a symmetric 101-day window.

Three empirical questions organize the analysis. First, does staged adoption raise farm-level productivity, and do the two stages produce distinct impact profiles? Second, how do returns vary across farm size, and what does this heterogeneity reveal about the constraints that limit value creation? Third, are gains amplified under weather stress, and if so, what does this imply about the mechanism through which digital information generates productivity effects?

Using a daily panel of 904 farms across five subdistricts, we find that messaging registration produces modest immediate gains in the Normalized Difference Vegetation Index (NDVI), most clearly for medium farms. AI platform enrollment produces larger, persistent slope effects that compound over the post-enrollment period. Returns scale with farm size and nearly double under drought conditions, consistent with an information-value mechanism: precise, timely advice matters most when input misallocation is most costly, and decision windows are narrowest. Small farms exhibit short-run adjustment costs before recovery, suggesting onboarding design as a margin that affects the distributional consequences of platform rollout.

The paper makes three contributions to research on digitization, productivity, and platform economics. First, it provides causal evidence that digitizing agronomic information generates measurable farm-level productivity gains, using satellite outcomes that sidestep the self-report problems common in this literature. Second, it shows that the level-versus-slope distinction between adoption stages is empirically sharp, consistent with learning-by-using at the AI tier rather than a pure activation effect. Third, it identifies state-dependent returns, in which gains are largest under resource stress, thereby adding an information-value mechanism to theories of technology diffusion grounded in the option value of timely advice under uncertainty (Dixit and Pindyck, 1994).

Taken together, the findings speak to how digital ventures create and distribute productivity gains in low-income markets, and what platform architecture and sequencing decisions determine whether that value reaches users with the least capacity to capture it.

2. Theoretical Framing

2.1 Digital leapfrogging as staged platform adoption

Technology diffusion in emerging markets rarely follows the linear path assumed in standard models. Infrastructure gaps, institutional frictions, and uneven complementary assets produce non-linear adoption trajectories in which late-adopting users skip intermediate technologies and move directly to higher-capability services (Antonelli, 1991; Mody and Dahlman, 1992; Steinmueller, 2001). In practice, this leapfrogging often requires a staging mechanism. A low-friction entry point acquires and habituates users, building the baseline knowledge and trust that make a higher-capability platform productive once introduced. The empirical footprint of this process is a two-stage adoption sequence with distinct impact profiles at each stage: a level effect at entry and a slope effect as learning accumulates at the advanced tier.

This framing connects to a broader literature on platform economics and product ladders. New ventures in low-income markets face a version of the cold-start problem: the services with the highest productivity potential are also the hardest to adopt without prior exposure (Prahalad, 2004; Steinmueller, 2001). A lightweight entry channel resolves this by generating engagement, trust, and basic complementary skills before the high-value tier is introduced. The channel sequence is therefore not merely a marketing choice but a mechanism for capability building that shapes the size and timing of productivity gains. Research on IT business value confirms that complementary assets, training, and process alignment are necessary conditions for impact, not adoption alone (Melville, Kraemer, and Gurbaxani, 2004; Bharadwaj, 2000).

2.2 Adoption costs, learning, and heterogeneous returns

Channel design mediates adoption costs and thus shapes technology trajectories. Messaging platforms reduce cognitive load, work on low-end devices with intermittent connectivity, and deliver time-bound prompts that lower the activation cost of behavior change. AI portals demand higher activation energy but enable deeper engagement, continuous monitoring, and decision support for complex, multi-variable problems. Classic diffusion mechanisms, including perceived usefulness, ease of use, and compatibility with existing practices, interact with these

channel frictions to determine the speed of uptake and learning dynamics (Rogers, 2003; Zhu and Kraemer, 2005).

Absorptive capacity, the ability to recognize, assimilate, and exploit new information, generates heterogeneity in realized value from information services (Cohen and Levinthal, 1990). Users with greater prior knowledge and operational capacity can act faster on advisory content, translating information into changes in input timing and resource allocation more efficiently. Learning-by-using predicts short-run adjustment costs as users recalibrate routines to a new information environment, followed by performance gains as new practices are internalized (Rosenberg, 1982). These mechanisms together predict a size and capacity gradient in returns: larger, better-resourced operations should realize gains faster and capture more value from the AI tier, while smaller operations face a longer adjustment before productivity improves.

Evidence from adjacent settings supports this prediction. Mobile money adoption in Kenya generated large welfare gains for poor households through simple, familiar interfaces, but returns were heterogeneous and depended on network density and access to liquidity (Suri and Jack, 2016). Work on digital inequality cautions that access alone does not guarantee benefits; users who bear adoption costs without the complementary assets to act on information can experience neutral or negative net returns (Heeks, 2022). Frugal innovation research shows that low-cost, user-centric designs can widen access and serve as stepping stones to advanced capabilities, but the step-up requires deliberate design (Zeschky, Widenmayer, and Gassmann, 2011; Prabhu, 2017).

2.3 Information value under uncertainty

The productivity value of a digital advisory system is not constant. It rises with the cost of the decisions it informs and with the penalty for getting those decisions wrong. In agricultural settings, heat stress and water deficits tighten the decision window for irrigation and input application, magnifying the shadow price of mistimed actions (Bastiaanssen et al., 2000; Fensholt et al., 2013). The option value of timely, precise information is therefore highest exactly when conditions are most demanding (Dixit and Pindyck, 1994).

This generates a state-dependent payoff structure in which advisory systems function as information insurance: they reduce output variance when risk is high by improving the quality of resource allocation under stress. The prediction is sharp that post-adoption productivity gains should be largest under hot, dry conditions, and it is testable with high-frequency satellite data capturing vegetation responses to both advisory inputs and weather conditions. This mechanism also has implications for how digital ventures in agriculture should price and position their services. A platform that delivers its largest gains under resource stress aligns naturally with drought management, water-use efficiency, and resilience-linked public programs.

2.4 Satellite data as a productivity measurement tool

Agriculture is well-suited for evaluating digital advisory systems because satellite indices provide objective, high-frequency outcomes that can be linked directly to managerial decisions. Prior work establishes the agronomic validity of this approach. Landsat and MODIS imagery have been used to estimate wheat yields in India with high accuracy (Becker-Reshef et al., 2003). Vegetation indices, including NDVI, are effective for monitoring crop biomass and stress responses at scale in developing-country settings (Lobell and Burke, 2010). Thermal remote sensing approaches estimate evapotranspiration and improve yield forecasting by capturing crop water use (Bastiaanssen, 2000). A recurring constraint for optical data in tropical regions is cloud contamination; synthetic aperture radar is robust to cloud cover and improves classification accuracy in these settings (Massey et al., 2018). These trade-offs motivate explicit cloud controls and multi-index robustness checks in empirical designs.

The use of satellite outcomes directly addresses a core limitation of the technology adoption literature: reliance on self-reported productivity measures, which introduce recall bias and measurement error. Daily vegetation indices observed at the farm level provide an outcome independent of respondent recall, capture agronomic change at the frequency relevant to advisory interventions, and are observable for the full pre- and post-adoption period. This makes satellite data particularly well-suited for event-time and regression kink designs where identification depends on detecting sharp changes in outcome trajectories around discrete adoption events.

The literature yields four predictions that organize the empirical analysis. First, registration with a low-friction messaging channel should produce an immediate level-type improvement in vegetation outcomes by lowering attention and action costs and triggering basic behavioral change (Davis, 1989; Bhattacherjee, 2001). Second, enrollment in the AI portal should produce slope-type gains as task-technology fit improves for more complex decisions and as learning-by-using accumulates over the post-enrollment period (Goodhue and Thompson, 1995; Rosenberg, 1982). Third, productivity effects should be larger under heat and drought conditions, when the option value of precise, timely information is highest (Dixit and Pindyck, 1994; Pavlou and El Sawy, 2006). Fourth, returns should increase with farm size and operational capacity, consistent with absorptive capacity theory and the resource-based view, such that smaller farms face larger short-run adjustment costs before gains materialize (Cohen and Levinthal, 1990; Barney, 1991).

3. Empirical Setting and Data

India's smallholder agriculture operates under binding resource constraints and rising climate variability. Monsoon dependence, fragmented landholdings, uneven access to irrigation, and thin agronomic support systems make timely, field-specific information especially valuable. At the same time, rapid mobile penetration and ubiquitous use of low-data messaging applications have created a practical pathway for digital leapfrogging: farmers can access decision support without first adopting costly intermediate technologies. This combination of high information need and low-friction access provides a natural context for studying how channel design shapes the adoption and use of AI-enabled, satellite-based advisory services, with implications for sustainability and climate resilience.

The advisory provider in this study is Farmonaut, which delivers satellite-derived agronomic insights through two complementary channels. The first is a lightweight, push-based WhatsApp stream designed for broad reach, low cognitive load, and low bandwidth. Messages provide concise recommendations on irrigation timing, fertilizer use, and crop health vigilance, drawing on up-to-date satellite data and local weather. The second is JEEVN AI, a feature-rich portal integrating multi-spectral satellite analytics with automation to surface field-level status, trends,

and risk flags. Whereas WhatsApp lowers activation barriers and encourages routine engagement, the portal supports deeper exploration (for example, maps, history, and drill-down views) and sustained planning. Together, the channels operationalize a staged adoption pathway: familiar messaging for initial onboarding and habit formation, followed by an AI-enhanced environment for users who seek more granular, continuous decision support. Web Appendix—WA1 shows the WhatsApp advisory, and WA2 displays the AI advisory platform.

The study area comprises multiple subdistricts within a single Indian state, spanning irrigated and rainfed conditions. Farms are heterogeneous in operational scale, ranging from very small plots to larger holdings typical of peri-urban and commercially oriented pockets. Cropping systems follow regional calendars, with monsoon-aligned Kharif cultivation and a subsequent Rabi season, during which water management remains pivotal. This agro-ecological diversity and shared exposure to monsoonal dynamics provide variation in baseline risk and potential returns to timely advisory information.

Seasonality is dominated by the southwest monsoon, which shapes sowing windows, vegetative growth, and irrigation needs. Intra-season weather volatility (heat spikes, dry spells, and intense rainfall) creates narrow decision windows in which satellite-informed advice can shift outcomes by improving the timing of water and input use. The setting is therefore well suited to examining resilience mechanisms. If advisories are valuable, their benefits should be most visible under heat and moisture stress when the costs of mis-timed actions are highest.

Farmer engagement begins with registration to receive WhatsApp advisories, typically supported by short onboarding interactions and vernacular messaging. Enrollment in JEEVN AI is available to a subset of users, reflecting a transition from periodic prompts to a continuous, self-serve decision environment. In practice, advisory content covers irrigation scheduling, early stress detection, and watchlists for potential risks. Operationally, the provider's backend fuses satellite observations with local weather feeds to generate high-frequency farm-specific guidance. This staged user journey, registration via messaging, and subsequent portal enrollment, anchors the empirical analysis while remaining grounded in real operational workflows. Figure 1 shows the framework for the contextual setting of staged digital adoption in agriculture.

Insert Figure 1 here

The setting reflects common features of emerging-market agriculture: uneven access to formal extension, variable digital literacy, and heterogeneous capacity to act on information. Messaging-based advisories exploit existing social and communication habits. At the same time, the AI portal caters to farmers and intermediaries who can devote attention to planning, record-keeping, and comparative analysis across plots and over time. Both channels operate in a context of cost and water pressures that foreground sustainability: better-timed irrigation and input use can reduce waste, protect yields, and strengthen climate resilience without requiring large capital expenditures.

The coexistence of a low-friction messaging stream and a higher-capability AI portal provides clear behavioral stages that map to plausible mechanisms: activation and habit formation versus planning and optimization. The monsoon-driven production cycle generates predictable seasonality alongside frequent weather shocks, creating meaningful variation in the value of information. Heterogeneity in farm size and water access further sharpens design-relevant contrasts in constraints and capabilities. These contextual features make the setting appropriate for testing how digital channel sequencing enables AI technology adoption in agriculture and assessing whether advisory value scales with conditions' most demanding.

Previous research establishes the technical feasibility and agronomic relevance of satellite monitoring for yield, drought, and precision management, and documents promising AI integrations. Yet gaps remain on three fronts. *First*, there is limited causal evidence on the impact of delivery channels, particularly the sequencing from low-friction messaging to higher-capability AI portals; more broadly, we are not aware of management research that directly theorizes and tests digital leapfrogging as a staged adoption process in emerging markets, tracing user behavior as it moves from simple prompts to complex analytics within the same season. *Second*, the role of environmental stress as a moderator of advisory returns is underexplored, despite sound theoretical grounds for resilience-driven heterogeneity and the likelihood that attention, learning, and action differ under weather extremes. *Third*, few studies map adoption

effects across farm-size or wealth segments, where constraints and complementarities differ systematically and shape behavioral responses.

Addressing these gaps, this paper leverages daily satellite outcomes and platform metadata to test staged adoption effects at two discrete events, registration (WhatsApp) and enrolment (AI portal), using event-time estimators and regression kink designs with explicit cloud and weather controls. By linking channel strategy to objective agronomic outcomes and documenting compelling effects under heat and dry conditions, the study contributes to theories of digital leapfrogging, staged technology adoption, and user behavior in emerging markets, while speaking directly to sustainability and climate-resilience positioning in agricultural innovation.

4. Methodology

4.1 Data and Measures

The dataset used in this study is derived from Farmonaut's AI-powered satellite advisory system. It includes multiple data sources relevant to evaluating the impact of satellite-based advisories on agricultural productivity. The data is structured into two main components: farm-specific static information and dynamic time-series data, enabling a comprehensive assessment of productivity outcomes. Table 1 shares the description of the variables in the data.

Insert Table 1 here

The study uses a proprietary daily panel from a digital advisory provider operating in India. The panel spans March 2023 to November 2023 and covers 904 anonymized farms across five subdistricts. The unit of analysis is the farm-day. Static attributes include an anonymous farm identifier, subdistrict, crop type, and a categorical land-size band. Farm size strata use area-band midpoints: Small ≤ 3,000 m², Medium 3,000–10,000 m², Large >10,000 m² (659/207/38 farms). Enrollment occurred in three closely spaced cohorts; registration preceded enrollment within the same season, so the stages are modeled separately.

Two adoption events anchor the analysis. *Registration* marks entry to a low-complexity messaging advisory delivered via WhatsApp. *Enrollment* activates an AI-enabled, satellite-driven advisory. Because enrollment typically follows registration in the same season, the two stages are modeled separately to avoid collinearity. For the event-time analysis, enrollment occurs on three calendar dates that define the cohorts: July 2 (cohort 23193), July 5 (cohort 23196), and July 7 (cohort 23198). The primary analysis uses *only cloud-free observations* and restricts the event window to [−99, +101] days around each farm's enrollment date. For registration-stage models, symmetric windows are constructed around the registration date. This ensures balanced pre and post periods in all event analyses.

Vegetation outcomes are derived from multi-spectral satellite imagery. NDVI is the primary outcome because it summarizes canopy vigor and biomass. Additional indices are constructed for robustness: NDWI for moisture content, SAVI for soil-adjusted vegetation, NDRE for chlorophyll and stress, RVI for growth rate, and RSM for vegetation cover.

Weather and season controls, like daily minimum and maximum temperature, relative humidity, and precipitation (probability and amount), are aligned to each farm day from gridded products. Season indicators for Kharif and Rabi absorb systematic seasonal differences in vegetation dynamics. Subdistrict indicators control for location-specific factors.

To study resilience, extreme weather conditions are coded as hot (max ≥ 35°C), cold (min ≤ 20°C), wet (rain ≥ 30 mm), and dry (rain ≤ 4 mm), with indicators for daily extremes and multi-day streaks. The extreme-day and streak indicators allow explicit tests of whether advanced advisories deliver larger gains under agronomic stress. This links adoption to a resilience mechanism that is directly relevant to a sustainability-oriented marketing strategy.

4.2 Empirical Estimations

Much of the prior literature documents correlations between satellite indices and productivity. Fewer studies use quasi-experimental designs to causally identify the impact of advisory systems on outcomes. Regression kink design (RKD) and event-time approaches allow estimation of local treatment effects around well-defined adoption events. In settings with staged

adoption, a double-event design is especially informative: registering for a messaging channel can plausibly produce an immediate level shift (activation and basic guidance), while enrolling in an AI portal may induce a slope change (learning, planning, and more precise timing). Difference-in-discontinuities offers a robustness strategy when adoption dates cluster or secular shocks need to be differenced. Bringing these designs to high-frequency satellite outcomes enables fine-grained tests of channel-specific mechanisms and resilience moderation.

4.2.1 Empirical Strategy

The empirical strategy combines three complementary designs to separate immediate level changes from gradual slope changes and to recover heterogeneous effects by farm size and weather stress. The designs are: an event-regression kink design, a natural cubic spline model of post-adoption dynamics, and a saturated event-time panel with interactions for extreme conditions and size strata. Windows are symmetric around each farm's adoption date to ensure balanced pre and post-periods. Registration and enrollment are estimated separately to avoid collinearity in event timing.

Model 1. Event regression kink design. The event RKD traces the path of Y_{it} around the adoption date using a saturated set of event time dummies and, where relevant, a visibility weight for image clarity. The specification is:

$$Y_{it} = \sum_{k=-K}^{K} \beta_k \, . event_{ik} + \sum_{k=-K}^{K} \gamma_k \, . (event_{ik} \times intensity_{it}) + X'_{it}\delta + \alpha_i + \varepsilon_{it}$$

In this equation, Y_{it} represents the vegetation index (NDVI, NDWI, SAVI, NDRE, RVI), where $event_{ik}$ = 1 if day t is k days from the adoption date of farm i, with one bin omitted as the reference period. The set $\{\beta_k\}$ traces the dynamic response. A persistent step up in β_k for $k >$ 0 relative to pre-event bins and a steeper profile for post-event bins are interpreted as level and slope changes, respectively. The term $intensity_{it}$ equals 1 on cloud-free days and 0 otherwise, and is used as an interaction to downweight low-quality scenes in robustness checks. The

interaction term ($event_{ik} \times intensity_{it}$) accounts for the influence of satellite image clarity on measurement precision. The vector X'_{it} includes time-varying controls such as weather and soil controls, as well as season fixed effects, while α_i denotes time-invariant heterogeneity such as slope, soil type, and long-run management. The error term ε_{it} captures unobserved shocks.

Model 2. Natural cubic spline model of post-adoption dynamics. The analysis fits a natural cubic spline in days since adoption to allow flexible curvature without saturating the timeline with event dummies. The specification is

$$Y_{it} = \beta_0 + \beta_1 \mathcal{D}_{it} + \sum_{j=1}^{4} \theta_j S_j(days\ since\ adoption) + X'_{it}\delta + \alpha_i + \varepsilon_{it}$$

where $\mathcal{D}_{it}$ marks the post adoption window and $S_j(\cdot)$ represents the natural j-th spline basis function of days since adoption, with j =1,…,4. For registration-stage models, the spline is a function of the number of days since registration. For enrollment-stage models, the spline is a function of the number of days since enrollment. Spline coefficients are not interpreted in isolation. They jointly describe the curvature of the recovery path and are used to plot fitted post-adoption trajectories by size class.

Model 3. Event time panel with weather and size interactions. The event time model is extended to study resilience and heterogeneity by capacity, with interactions for extreme conditions and farm size classes. Let e_{itk} be the event time dummies as above. Let $Extreme_{it}$ be a daily indicator for hot, cold, wet, or dry days, or a streak indicator that marks days that lie within runs of extremes. Let $Size_i$ index the three size strata. The specification is

$$Y_{it} = \alpha_i + \sum_k \beta_k . e_{itk} + \sum_k \theta_k (e_{itk} \times Extreme_{it}) + \sum_k \psi_k (e_{itk} \times Extreme_{it} \times Size_i) + X'_{it}\lambda + \varepsilon_{it}$$

The coefficients $\{\theta_k\}$ recover how post-enrollment dynamics differ on extreme days. The coefficients $\{\psi_k\}$ recover how those differences vary by farm size. This design produces the pre versus post contrasts shown for dry day conditions by cohort and size class in the results.

4.2.2 Identification Strategy

Our identification relies on the assumption that, within local windows around each adoption date, unobserved determinants of vegetation evolve smoothly through the cutoffs and that no other discrete policy or agronomic shocks coincide with those dates. To support this assumption, we conduct a set of validation exercises.

First, we examine pre-trend diagnostics. In event-time plots, coefficients for the pre-adoption bins are jointly close to zero. We report joint F-tests showing that these pre-event coefficients do not differ statistically, reducing concern about anticipatory behavior or differential trends before adoption. *Second*, we assess bandwidth and kernel sensitivity. Estimates are stable when we narrow or widen the local window and switch across alternative kernel weightings. For the spline specification, varying knot locations in grid searches yield similar fitted trajectories, indicating that results are not an artifact of a particular smoothness choice. *Third*, we implement donut and placebo tests. Excluding observations immediately adjacent to the adoption cutoff (a donut) leaves results intact, suggesting they are not driven by transient transients around the event. Placebo cutoffs placed in the pre-period do not reproduce the post-adoption slope break, supporting a causal interpretation of the adoption date. *Fourth*, we include common-shock controls. Calendar-time fixed effects at a coarse frequency, or cohort-by-date indicators when appropriate, absorb shocks shared across farms on the same day. Results remain unchanged, indicating that common shocks are not driving the findings. *Fifth*, we account for measurement quality. Weighting by scene visibility, or restricting to higher-quality scenes, yields estimates substantively similar to the baseline, reassuring that conclusions are not sensitive to variation in image clarity. *Sixth*, we check outcome robustness. Repeating the analyses on alternative vegetation and water-status indices preserves the qualitative patterns, with magnitude differences consistent with each index's sensitivity to canopy vigor and moisture conditions.

All models are estimated using ordinary least squares with farm fixed effects and time-varying controls (e.g., daily weather and season indicators). We report heteroskedasticity-robust standard errors clustered at the farm level; as a robustness check, we also use two-way clustering by farm and calendar date. For regression–kink specifications, we implement robust bias correction where applicable. Dynamic effects are summarized using event-time graphs and compact post-adoption contrasts (e.g., averages over selected post-event bins). Heterogeneity is presented as differences in these contrasts across extreme weather states and farm-size strata. Where families of related hypotheses are tested, we apply multiple-testing adjustments.

5. Results and Findings

This section reports the effects of two adoption stages and layers of heterogeneity based on farm size and extreme weather. Table 2 summarizes the spline models centered on the registration date. Medium farms exhibit the most precise and immediate improvement in NDVI. The event indicator for medium farms is positive and significant (+7.782, $p < .01$), and the first spline basis term is significantly positive (+74.293, $p < .001$), indicating a sharp post-registration rise that tapers thereafter. Small farms also move in the right direction, with an early positive spline (+26.793, $p < .05$) consistent with incremental gains once simple messages arrive on a familiar channel. Large farms show muted and statistically insignificant movement around registration.

Together, these patterns indicate that the low-complexity entry point is most effective where there is enough organizational capacity to act quickly, yet not so much complexity that existing routines slow adoption. This aligns with staged adoption and assimilation perspectives, in which medium-scale operators often rapidly convert information into action (Rogers 2003; Zhu and Kraemer 2005).

Insert Table 2 here

The early post-enrollment period often dips below the pre-trend, and the trajectory then rises substantially in later spline segments. Small farms illustrate this V-shape, with an initial decline (−1.382, $p < .05$) and a mid-period drop (−12.383, $p < .001$), followed by a pronounced recovery (+9.722, $p < .001$). Medium farms show a mid-period dip (−14.799, $p < .001$) and significant late gains (+11.268, $p < .001$). Large farms display the most clearly positive late response (+12.594, +8.016, +13.307, all significant). The pattern is consistent with learning costs and workflow realignment as users shift from simple messages to precision, map-based recommendations, with larger payoffs once practices are adapted. In digital leapfrogging terms, the WhatsApp stage builds familiarity and trust, while the AI stage creates the step change in capability once initial frictions are overcome (Antonelli 2017; Heeks 2022; Steinmueller 2001).

The early negative spline segments after enrollment may appear counterintuitive, particularly for small farms. Three mechanisms can reconcile the evidence. *First*, adoption learning costs can temporarily disrupt routines while farmers recalibrate irrigation cycles and nutrient schedules to align with satellite-driven thresholds. *Second*, measurement can become stricter after enrollment because cloud-free days are more frequent in the analysis window, and the advisory prompts activity exactly when crops are stressed, which can lower NDVI before the corrective actions take effect. *Third*, the enrollment date often falls near mid-season, so growth-stage transitions can coincide with the cutoff, resulting in temporary dips even as the trajectory improves. The subsequent positive slope and late-period gains in all size strata indicate that net improvements follow these short-run adjustments.

The enrollment results show slower starts followed by larger payoffs, especially for medium and large farms. Figure 2 demonstrates that payoffs are most salient when the weather is adverse, a setting where decision support is most valuable.

Insert Figure 2 here

Figure 2 plots average predicted NDVI by cohort and farm size, comparing pre- and post-enrollment periods on dry days, defined as precipitation at most 4 mm. Dry conditions are informative because they tighten resource constraints and raise the value of precise irrigation advice. Across cohorts, NDVI is higher after enrollment for all size classes, yet the magnitude and stability of gains differ. Medium and large farms exhibit the largest and most stable post-enrollment improvements, while small farms improve modestly with wider confidence intervals. Web Appendix—WA3 to WA6 shows the predicted NDVI by cohort and farm size for regular and extreme weather days.

The size gradient is consistent with a capacity-based mechanism, in which pumps, labor scheduling, and access to inputs are less binding, and farmers can translate AI guidance into timely irrigation and input decisions during moisture stress. The finding also links the service to resilience outcomes central to global market sustainability positioning. The evidence supports a sequenced adoption pathway in which a low-friction channel initiates behavior change, and an advanced platform delivers cumulative gains once users adapt. The pattern is consistent with digital leapfrogging theory that emphasizes windows of opportunity and staged capability building in late-adopting markets (Antonelli 1991; Mody and Dahlman 1992; Rogers 2003; Zhu and Kraemer 2005; Heeks 2022; Steinmueller 2001).

The identification strategy assumes that unobserved determinants of vegetation evolve smoothly across the adoption cutoffs, whereas registration or enrollment induces a discontinuity in level or slope. Event-time plots with saturated pre-period bins show flat or declining NDVI before enrollment across cohorts, and registration-centered plots show similar stability, thereby limiting concerns about anticipatory behavior. Results persist when replacing NDVI with NDWI, SAVI, NDRE, RVI, or RSM; when adding weather-by-season interactions; when dropping soil controls; and when restricting to a single season. Trimming outliers by interquartile rules does not change the signs or significance.

Heterogeneity findings are not artifacts of binning, as models using the continuous area midpoint reproduce the exact monotone ordering by scale. Although the large-farm stratum is small, late-period gains after enrollment remain statistically different from zero. The flat pre-

trends, stable re-estimations, null placebos, and consistency across indices support a causal interpretation for the observed gains from staged digital adoption.

Collectively, three implications follow. *First*, market entry in resource-constrained contexts benefits from intentional sequencing. Simple, ubiquitous channels such as messaging build reach and credibility, then AI platforms can be introduced with onboarding that targets the expected early dip. *Second*, segmentation by capacity is essential. Small farms may require bundles that include short training modules, vernacular support, and low-cost irrigation checklists to compress the learning curve. Large farms can be approached with integration features and multi-plot dashboards that leverage existing assets. *Third*, positioning the service as climate resilience, not only productivity, is credible because post-enrollment gains in vegetation indices are largest under dry conditions.

6. Discussion

The evidence from 904 farms across five subdistricts points to three clean findings. Registration with the messaging channel produces immediate, modest productivity gains, most clearly for medium farms. Enrollment in the AI portal produces larger, persistent slope effects after a short adjustment period, with the strongest late gains for medium and large farms. Both sets of gains are amplified under dry conditions, when the cost of mistimed input decisions is highest. Flat pre-trends, stable re-estimations across alternative indices, and null placebo cutoffs support a causal reading of these patterns.

6.1 Staged adoption as a productivity mechanism

The level-versus-slope distinction between registration and enrollment effects is the central empirical finding, and it has a natural interpretation in terms of learning-by-using (Rosenberg, 1982). WhatsApp registration reduces the cost of accessing agronomic information and prompts straightforward behavioral changes, resulting in an immediate improvement in vegetation outcomes. AI portal enrollment changes the nature of the decision environment: users shift from receiving periodic prompts to engaging with continuous, field-level analytics that support planning and optimization. This transition takes time. The temporary dip in NDVI after

enrollment, most pronounced for small farms, is consistent with short-run workflow disruption as farmers recalibrate irrigation cycles and input schedules to satellite-driven thresholds (Jovanovic, 1982; Arrow, 1962). Once practices align with the new information environment, the slope turns positive and late-period gains emerge across all size classes.

This pattern connects to a broader literature on IT business value. Brynjolfsson and Hitt (2000) show that the productivity returns to information technology are larger and more persistent when firms invest in complementary organizational changes alongside the technology itself. The adjustment period observed here is the agricultural equivalent: the platform's value is not unlocked by access alone but by the reconfiguration of routines around the information it provides. Messaging registration builds the baseline habits and trust that make the AI portal productive once introduced, rather than substituting for it. The front-end is a complement to the full-stack layer, not a replacement (Milgrom and Roberts, 1990).

6.2 Heterogeneous returns and absorptive capacity

The size gradient in returns is sharp and consistent across cohorts and weather conditions. Medium farms respond fastest to messaging registration and scale gains reliably after AI enrollment. Large farms show the largest late-period slope effects once the portal is in place. Small farms improve but face wider confidence intervals and a more pronounced early dip.

Absorptive capacity offers the most direct explanation (Cohen and Levinthal, 1990). Farms with greater operational capacity, better access to pumps, more flexible labor scheduling, and greater financial slack can translate advisory content into action at lower marginal cost. Where these complements are available, the response to precise, timely information is fast and stable. Where they are absent, the information is valuable in principle but costly to act on, producing a longer adjustment period and noisier gains. This is consistent with evidence on the productivity dispersion effects of information technology more broadly: IT raises average productivity but widens the gap between high- and low-capability adopters, at least in the short run (Acemoglu et al., 2007; Bloom et al., 2012).

The finding also speaks to platform economics. Platforms serving heterogeneous user populations face a design tension between accessibility and capability. A single-tier product

either leaves value on the table for high-capacity users or imposes prohibitive adoption costs on low-capacity ones (Parker, Van Alstyne and Choudary, 2016; Rochet and Tirole, 2003). The two-stage architecture here resolves this tension by matching product complexity to user readiness at each stage. The identification of short-run adjustment costs for small farms further pinpoints where onboarding design can shift the distributional consequences of platform rollout without altering the underlying technology.

6.3 State-dependent returns and the option value of information

The amplification of enrollment effects under dry conditions is the paper's sharpest mechanism result. When precipitation falls below 4mm per day, post-enrollment NDVI gains are larger and more stable across all size classes. This is consistent with the option value of information rising under uncertainty (Dixit and Pindyck, 1994). Drought conditions tighten the window for effective irrigation and raise the penalty for mistimed input application. In this environment, a platform that delivers precise, field-level guidance on irrigation scheduling reduces the expected cost of a bad decision more than it would under average conditions. The advisory system functions as information insurance, reducing outcome variance when the stakes are highest (Hirshleifer and Riley, 1979).

This state dependence has implications beyond this setting. A growing literature on technology adoption in developing agriculture documents average treatment effects but rarely asks when those effects are largest (Cole and Fernando, 2021; Suri, 2011). The finding here suggests that the returns to digital advisory services are concentrated in exactly the periods when farmers are most exposed to yield-reducing shocks. This creates a natural alignment between the private value of platform adoption and the social value of reducing climate-driven productivity losses, an alignment that policy design can exploit through stress-contingent pricing or subsidy triggers.

6.4 Platform design and market entry in low-income settings

Three design principles follow from the empirical patterns. First, sequencing matters for both reach and productivity. A low-friction entry channel is not a concession to low-end users; it is the mechanism by which high-value AI services become accessible at all. Platforms that skip this stage and deploy full-stack AI directly will face adoption barriers that the messaging layer is

specifically designed to dissolve. Second, onboarding design is a distributional lever. The early dip for small farms after AI enrollment is not a fixed feature of smallholder agriculture; it reflects a mismatch between the platform's decision interface and the user's current workflow. Short, task-specific onboarding modules aligned with the transition period can accelerate this adjustment. Third, timing upgrade offers to forecast stress windows, rather than to average engagement levels, is likely to raise both conversion rates and measured returns.

These principles connect to the literature on market design in emerging markets. Prahalad (2004) argues that serving low-income users requires rethinking product architecture, not just pricing. The frugal-to-full-stack pathway demonstrated here is a concrete instance of this: a two-tier architecture that matches product complexity to user readiness at each stage, generating platform-wide gains while preserving a path to full-value adoption for users with greater capacity. Bandiera and Rasul (2006) show that social learning and network structure shape technology adoption in smallholder settings; a staging architecture that builds familiarity before introducing complexity can substitute for the peer observation effects that accelerate diffusion in denser networks.

The results also speak to the economics of new venture scaling in low-income markets. Farmonaut's staged deployment converts a broad base of low-engagement messaging users into high-value AI platform subscribers, illustrating how a digital venture can build a conversion funnel without requiring users to invest in prior infrastructure. The productivity gains documented here are the value proposition that sustains this funnel. Because gains are largest under moisture stress, alignment with drought management schemes, water-saving initiatives, and resilience-linked credit products is natural and potentially self-financing through reduced default risk during stress periods.

7. Implications and Future Directions

7.1 Managerial Implications

The results suggest a clear playbook for sequencing, targeting, and timing AI upgrades in resource-constrained markets.

Upgrade offers should be prioritized where operational capacity and payoffs are highest. Medium and large farms exhibit the most stable late gains after enrollment, particularly during dry conditions, so an initial focus on these segments is warranted. Practical screening rules can include recent consistency in responding to WhatsApp advisories, access to pumps or shared irrigation, cooperative membership, and prior responsiveness to irrigation prompts. These proxies for absorptive capacity are cheap to observe from platform engagement data and are likely to outperform simple demographic targeting (Cohen and Levinthal, 1990; Bloom et al., 2012). Smallholders should not be excluded; the AI portal should be bundled with short onboarding and micro-timing aids to compress the early-learning dip observed after enrollment.

Upgrades should be timed to when the value of information is highest. Nudges just before forecast dry spells or heat streaks, or after two to three weeks of regular messaging engagement, are likely to convert at higher rates and produce measurable gains. This is consistent with evidence that demand for information services rises with environmental uncertainty and that the timing of conversion relative to stress exposure materially affects realized returns (Dixit and Pindyck, 1994; Cole and Fernando, 2021).

Product design should preserve a frugal-to-full-stack pathway: use WhatsApp to build habits with short, vernacular prompts and field photos, then transition engaged users to the portal with a two-week guided sprint centered on irrigation timing and stress thresholds. For smallholders, add voice notes, checklists, and default reminders tied to forecast stress days. For medium and large farms, emphasize multi-plot dashboards, historical trend views, and exportable task lists that integrate with pumps and scheduling apps. Pricing can reinforce resilience by offering fee rebates or credits when verified portal use coincides with officially declared dry streaks, and by providing referral credits for cooperatives that onboard a threshold share of members. Managers should treat the upgrade as a conversion funnel, A/B test message format and cadence, track time from invite to first portal task, and instrument the onboarding sequence to identify which steps most effectively reduce the early dip (Brynjolfsson and Hitt, 2000).

Future managerial research should identify the messaging cadence that maximizes portal conversion without fatigue across seasons, isolate the specific onboarding tasks that most reduce the early dip for smallholders, and design pricing that varies with predicted volatility so that

farmers self-select efficiently into resilience-focused plans. It will also be valuable to develop reliable proxies for high return to the portal, such as cooperative links or pump ownership, and to evaluate whether these improve targeting beyond simple engagement metrics.

7.2 Academic Implications

The paper contributes to research on digitization, technology adoption, and platform economics by linking a staged sequence of channels to distinct empirical signatures: level-type activation from messaging registration and slope-type intensification after AI enrollment. It clarifies the role of environmental volatility by showing that the payoff to information is state-dependent, increasing sharply under dry conditions. It also delineates boundary conditions for task-technology fit by documenting returns that scale with operational capacity (Goodhue and Thompson, 1995; Brynjolfsson and Hitt, 2000).

Several research extensions follow. Mechanism studies should combine advisory logs with timestamped field actions, such as pump run-time or irrigation starts, to separate attention from compliance and to map advice-to-action latency. This would allow direct estimation of how quickly information translates into behavior change and where in the advisory chain value is lost (Suri, 2011; Cole and Fernando, 2021). Longer-horizon work can track multi-season persistence, input intensity, profits, and water withdrawals to quantify joint productivity and resource-use outcomes, connecting the short-run NDVI effects documented here to the broader literature on agricultural total factor productivity (Foster and Rosenzweig, 2010). Network studies can test for neighborhood and cooperative spillovers in adoption and effects, distinguishing peer learning from congestion, building on evidence that social networks shape technology diffusion in smallholder settings (Bandiera and Rasul, 2006; Conley and Udry, 2010). Structural and behavioral models can estimate learning curves and switching costs, and model attention constraints and habit formation under uncertainty (Jovanovic, 1982; Fudenberg and Tirole, 1991). External validity should be assessed across crops, states, and countries with different monsoon regimes and institutional settings, including contexts where radar-based outcomes are required due to persistent cloud cover. Product-architecture experiments that randomize stress-window alerts, checklists, or voice notes can identify which components deliver the largest marginal gains for each farm-size segment.

7.3 Policy Implications

Sequenced digital programs can scale productivity-improving practices at low marginal cost by pairing wide-reach messaging campaigns with targeted AI upgrades for users that can act on granular analytics. Because gains are largest under moisture stress, advisory systems complement drought management and water-saving initiatives and create opportunities for public-private partnerships to time intensive advisories to forecast stress windows (World Bank, 2016).

Policy design should prioritize access for smallholders through vernacular onboarding, shared pump schemes, and simple irrigation timers that lower the marginal cost of acting on advisory content. Resilience-linked vouchers discounting portal fees during declared dry or heat streaks would directly align incentive structures with the periods when advisory value is highest. Integration with crop insurance and working-capital credit is a natural extension: premium rebates or improved loan terms when verified advisory use and compliance are observed during stress periods could reduce insurer default risk while accelerating platform adoption (Jensen, 2010; Barnett et al., 2008). Data governance requires explicit consent frameworks, privacy and data-sharing standards, and transparent public dashboards that report resilience outcomes using satellite metrics while protecting farm-level identities.

Future policy research should compare subsidy instruments, including vouchers, rebates, and results-based contracts, for their ability to accelerate upgrades without crowding out private demand (Kremer and Muralidharan, 2008). It should test whether verified advisory compliance can serve as an underwriting input for credit and insurance, and whether it reduces defaults or claims during stress periods. A further question is which training modalities scale best for smallholders across group sessions, phone-based micro-lessons, and peer ambassadors, and how cooperative or district-level procurement of information-enabled resilience, measured by vegetation or water-status improvements during stress windows, can be aligned with privacy and equity goals.

A research agenda that links advice, action, and long-run outcomes at the farm level can establish whether staged digital adoption constitutes a durable path to productivity growth in

emerging-market agriculture, or whether the gains documented here are short-run adjustments that plateau without further complementary investments in infrastructure and institutional capacity.

8. Conclusion

This paper set out to answer a concrete question: Does staged digital adoption generate measurable farm-level productivity gains, and if so, when and for whom? Using a daily farm-level panel of 904 farms, event-time estimators, and regression kink designs around two discrete adoption events, the answer is yes, with important qualifications about timing, farm size, and weather conditions.

Registration with the WhatsApp messaging service produces modest, immediate improvements in vegetation productivity, most clearly for medium farms. Enrollment in the AI portal produces larger, persistent slope effects after a short adjustment period, with the strongest late gains for medium and large farms. Both sets of effects are amplified under dry conditions, when the cost of mistimed input decisions is highest, and the option value of precise, timely information is largest. Flat pre-trends, null placebo tests, and consistency across alternative vegetation indices support a causal reading of these patterns.

The level-versus-slope distinction between the two adoption stages is the paper's central empirical finding. It is consistent with a learning-by-using mechanism in which the messaging stage builds the habits and baseline knowledge that make the AI platform productive, rather than serving as a substitute for it. The platform's value is not unlocked by access alone but by the reconfiguration of farming routines around the information it provides. This makes the frugal entry channel a productive complement to the full-stack AI layer, rather than a fallback for users who cannot access the higher tier.

The size gradient in returns maps directly to heterogeneity in absorptive capacity. Where operational complements, pumps, flexible labor, and access to inputs are available, farmers can quickly translate advisory content into action at low marginal cost. Where those complements are absent, the information is valuable in principle, but adjustment takes longer. Short-run learning costs for small farms are not a fixed feature of smallholder agriculture; they are a design problem

that targeted onboarding can address. The finding that gains compound under drought conditions positions the advisory bundle as information insurance, reducing outcome variance precisely when the stakes are highest.

For platform designers and new ventures operating in low-income agricultural markets, the evidence supports three practical conclusions. First, sequencing from a low-friction channel to a high-capability AI platform is not just a marketing ladder but a mechanism for building the user capacity that makes the AI tier productive. Second, upgrade timing relative to forecast weather stress materially affects both conversion rates and realized returns. Third, the distributional consequences of platform rollout are shaped by onboarding design, not just by underlying farm heterogeneity, which makes them tractable.

For policy, wide-reach messaging can expand access at low marginal cost, while AI upgrades can be channeled through cooperatives and schemes that value resilience outcomes. Integration with crop insurance and working-capital credit is a natural extension if verified advisory use during stress periods is treated as a signal of reduced default risk.

Several questions remain open. Linking advisory content to timestamped field actions would enable direct estimation of advice-to-action latency and the separation of attention effects from compliance effects. Multi-season panels that track yields, water withdrawals, and input intensity would connect the short-run NDVI gains documented here to longer-run productivity and resource-use outcomes. Network studies within cooperatives and villages can test whether adoption effects spill over to non-adopters through peer learning. Replication across crops, states, and countries with different monsoon regimes and institutional settings will establish how far these findings generalize and where platform architecture needs to adapt to local constraints.

The broader implication is straightforward. Digitizing agronomic information creates measurable productivity value in low-income agricultural markets, but that value is not automatic. It depends on how platforms are sequenced, who is targeted at each stage, and when upgrades are introduced relative to the conditions under which information is most valuable. Getting those design choices right is both an economic question and an entrepreneurial one.

References

Acemoglu, D., Aghion, P., Zilibotti, F. 2006. Distance to Frontier, Selection, and Economic Growth. *Journal of the European Economic Association* 4(1): 37-74.

Alzouma, G. 2005. Myths of digital technology in Africa: Leapfrogging development?. *Global Media and Communication*, 1(3): 339–356.

Antonelli, C. 1991. *The Diffusion of Advanced Telecommunications in Developing Countries*. OECD, Paris.

Antonelli, C. 2017. Digital knowledge generation and the appropriability trade-off. *Telecommunications Policy*, 41(10): 991–1002.

Arrow, K. J. 1962. The Economic Implications of Learning by Doing. *Review of Economic Studies* 29(3): 155-173.

Bandiera, O., Rasul, I. 2006. Social Networks and Technology Adoption in Northern Mozambique. *Economic Journal* 116(514): 869–902.

Barnett, B. J., Barrett, C. B., Skees, J. R. 2008. Poverty Traps and Index-Based Risk Transfer Products. *World Development* 36(10): 1766-1785.

Barney, J. B. 1991. Firm resources and sustained competitive advantage. *Journal of Management*, 17(1): 99–120.

Bastiaanssen, W. G. M., Molden, D. J., & Makin, I. W. 2000. Remote sensing for irrigated agriculture: Examples from research and possible applications. *Agricultural Water Management*, 46(2): 137–155.

Becker-Reshef, I., Justice, C., & Sullivan, S. 2003. Agricultural yield forecasting using satellite data. *Remote Sensing of Environment*, 86: 271–282

Bharadwaj, A. S. 2000. A resource-based perspective on information technology capability and firm performance: An empirical investigation. *MIS Quarterly*, 24(1): 169–196.

Bhattacherjee, A. 2001. Understanding information systems continuance: An expectation–confirmation model. *MIS Quarterly*, 25(3): 351–370.

Bloom, N., Sadun, R., Van Reenen, J. 2012. Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review* 102(1): 167–201.

Brynjolfsson, E., Hitt, L. M. 2000. Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives* 14(4): 23–48.

Chen, Y., Farinelli, U., & Johansson, T. B. 2004. Technological leapfrogging–a strategic pathway to modernisation of the Chinese iron and steel industry?. *Energy for Sustainable Development*, 8(2): 30–38.

Cohen, W. M., & Levinthal, D. A. 1990. Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1): 128152.

Cole, S., Fernando, A. N. 2021. Mobile'izing Agricultural Advice: Technology Adoption, Diffusion, and Sustainability. *Economic Journal* 131(633): 192–219.

Conley, T. G., Udry, C. R. 2010. Learning about a New Technology: Pineapple in Ghana. *American Economic Review* 100(1): 35–69.

Davis, F. D. 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3): 319–340.

Dayasindhu, N., & Chandrashekar, S. 2005. Indian remote sensing program: A national system of innovation? *Technological Forecasting and Social Change*, 72(3): 287–299.

Dixit, A. K., & Pindyck, R. S. 1994. *Investment under uncertainty.* Princeton, NJ: Princeton University Press.

Dutta, S., Ray, B., & Mukherjee, G. 2015. Drought monitoring using remote sensing data in India. *Journal of Hydrology*, 523: 524–538.

Fensholt, R., Rasmussen, K., & Kaspersen, P. S. 2013. Assessing drought severity in Ethiopia using satellite-based indices. *Remote Sensing of Environment*, 127: 132–147.

Foster, A. D., Rosenzweig, M. R. 2010. Microeconomics of Technology Adoption. *Annual Review of Economics* 2(1): 395-424.

Fudenberg, D., Tirole, J. 1991. *Game Theory*. MIT Press.

Goodhue, D. L., & Thompson, R. L. 1995. Task–technology fit and individual performance. *MIS Quarterly*, 19(2): 213–236.

Heeks, R. 2022. Digital inequality beyond the digital divide: Conceptualizing adverse digital incorporation in the Global South. *Information Technology for Development*, 28(4): 688–704.

Hirshleifer, J., Riley, J. G. 1979. The Analytics of Uncertainty and Information: An Expository Survey. *Journal of Economic Literature* 17(4): 1375–1421.

Jain, M., Srivastava, R., & Kaushik, R. 2016. Mapping paddy fields in India using deep learning and satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 122: 233–246.

Jensen, R. 2010. The (Perceived) Returns to Education and the Demand for Schooling. *Quarterly Journal of Economics* 125(2): 515–548.

Jovanovic, B. 1982. Selection and the Evolution of Industry. *Econometrica* 50(3): 649–670.

Kremer, M., Muralidharan, K. 2008. Public and Private Schools in Rural India. In P. Peterson and R. Chakrabarti (Eds.), *School Choice International*. MIT Press.

Lobell, D.B. and Burke, M.B. 2010. On the use of satellite data for yield gap analysis. *Field Crops Research,* 119: 240–245.

Majumdar, S., & Kumar, A. 2020. AI and remote sensing for precision agriculture. *Agricultural Systems*, 187: 102982.

Massey, R., Gowda, P. H., & Yost, M. A. 2018. Remote sensing for precision agriculture in tropical climates. *Remote Sensing*, 10(10): 1629.

Melville, N., Kraemer, K., & Gurbaxani, V. 2004. Review: Information technology and organizational performance: An integrative model of IT business value. *MIS Quarterly*, 28(2): 283–322.

Milgrom, P., Roberts, J. 1990. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *American Economic Review* 80(3): 511–528.

Miller, R. L. 2015. Rogers' innovation diffusion theory (1962, 1995). In M. N. Al-Suqri & A. S. Al-Aufi (Eds.), *Information seeking behavior and technology adoption: Theories and trends*: 261–274. Hershey, PA: Information Science Reference.

Mody, A., & Dahlman, C. 1992. Performance and potential of information technology: An international perspective. *World Development*, 20(12): 1703–1719.

Mulla, D. J. 2013. Mapping and managing spatial patterns in soil fertility and crop yield. *Geoderma*, 199: 1–9.

Parker, G., Van Alstyne, M., Choudary, S. P. 2016. *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*. W. W. Norton.

Pavlou, P. A., & El Sawy, O. A. 2006. From IT leveraging competence to competitive advantage in turbulent environments: The case of new product development. *Information Systems Research*, 17(3): 198–227.

Prabhu, J. 2017. Frugal innovation: doing more with less for more. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *375*(2095): 20160372.

Prahalad, C. K. 2004. *The fortune at the bottom of the pyramid: Eradicating poverty through profits.* Upper Saddle River, NJ: Wharton School Publishing.

Rochet, J. C., Tirole, J. 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1(4): 990–1029.

Rogers, E.M. 2003. Elements of diffusion. *Diffusion of Innovations*, 5(1.38): 453.

Rosenberg, N. 1982. *Inside the Black Box: Technology and Economics*. Cambridge University Press.

Steinmueller, W. E. 2001. ICTs and the possibilities for leapfrogging by developing countries. *International Labour Review*, 140(2): 193–210.

Suri, T., & Jack, W. 2016. The long-run poverty and gender impacts of mobile money. *Science*, 354(6317): 1288–1292.

World Bank. 2016. *World Development Report 2016: Digital Dividends*. World Bank.

Zeng, J., Kumar, P., & Arora, V. K. 2015. Satellite remote sensing of soil moisture: Applications for irrigation management. *Agricultural Water Management*, 153: 1–15.

Zeschky, M., Widenmayer, B., & Gassmann, O. 2011. Frugal innovation in emerging markets. *Research-Technology Management*, 54(4): 38–45.

Zhu, K., & Kraemer, K. L. 2005. Post-adoption variations in usage and value of e-business by organizations: Cross-country evidence from the retail industry. *Information Systems Research*, 16(1): 61–84.

TABLES AND FIGURES

Table 1: Data Description

Category	Variable Name	Description
Farm Information (Static Data)		
Identifier	FarmID	Unique identifier anonymized for each farm.
Location	RegionID	Geographical location of the farm (instead of exact locations, farms are mapped to numerical region identifiers).
Crop Type	Crop	Type of crop grown on the farm.
Area	Field Area Bands	Farm area values are grouped into predefined bands to prevent identification based on total land area (in square meters).
Time-Series Data (Dynamic Data)		
Date	Date	Timestamp of data collection.
WhatsApp Registration	Registration Flag	1 = post-registration, 0 = pre-registration.
Advisory Enrollment	Enrollment Flag	1 = post-enrollment, 0 = pre-enrollment.
Crop Season	Season	Classification into agricultural seasons (Kharif: June–October, Rabi: November–March).
Year	Year	Year of observation.
Productivity Indices		
Vegetation Health	NDVI	Normalized Difference Vegetation Index: Measures vegetation greenness, health, and density
Chlorophyll Content	NDRE	Normalized Difference Red Edge Index: Captures chlorophyll content and plant stress.
Vegetation Productivity	RVI	Radar Vegetation Index: Measures plant productivity and biomass.
Soil Moisture	RSM	Radar Soil Moisture: Soil moisture reflecting water stress for plants and surface hydrological processes
Vegetation Monitoring	SAVI	Soil-Adjusted Vegetation Index: Vegetation monitoring that reduces soil background effects
Soil Water Content	NDWI	Normalized Difference Water Index: Measures plant water content features, highlighting their presence and availability.
AI-generated Productivity Indices		
AI-generated Productivity Indices	No Crop Cloud	No significant vegetation cover detected.
	All Good	Healthy crop condition detected.
	Health Declining	AI indicates a decline in vegetation health.
	Irrigation Required	The advisory suggests the need for irrigation.
	Critical Health Irrigation	Urgent irrigation is required due to poor crop health.
Farm-Specific Data		
Crop Damage	Number of Damage Events	Maximum number of high-probability pest, weed, or disease occurrences detected by AI advisories.
Control Variables		
Weather	Humidity (%)	Percentage of humidity at a given time.
	Min Temperature	Minimum daily temperature in Celsius.
	Max Temperature	Maximum daily temperature in Celsius.

Table 2: Regression Kink Design (RKD) and Spline Regression Estimates

Term	Small Farms – Registration	Small Farms – Enrollment	Medium Farms – Registration	Medium Farms – Enrollment	Large Farms – Registration	Large Farms – Enrollment
(Intercept)	27.452*** (2.270)	29.434*** (0.324)	25.173*** (3.916)	27.195*** (0.652)	25.474*** (4.243)	22.074*** (1.666)
event_reg60	2.718 (2.144)		7.782** (2.904)		-0.292 (3.195)	—
splines (days_since_registration, df=4)$_1$	26.793* (10.651)		74.293*** (17.387)		12.078 (8.910)	
splines (days_since_registration, df=4)$_2$	-4.039 (4.208)		-24.697*** (5.154)		-1.134 (8.281)	
splines (days_since_registration, df=4)$_3$	4.105 (2.507)		12.300** (4.627)		4.238 (4.045)	
splines (days_since_registration, df=4)$_4$	-4.354 (2.914)		-7.220 (4.404)		1.484 (5.719)	
splines (days_enrollment, df=4)$_1$		-1.382* (0.596)		2.206+ (1.182)		12.594* (5.460)
splines (days_since_enrollment, df=4)$_2$		-12.383*** (0.476)		-14.799*** (0.990)		-15.669*** (2.797)
splines (days_since_enrollment, df=4)$_3$		-4.739*** (0.561)		0.825 (1.375)		8.016* (3.557)
splines (days_since_enrollment, df=4)$_4$		9.722*** (0.376)		11.268*** (0.735)		13.307*** (1.868)
No. of farms	427	427	439	439	38	38
No. of obs.	1989	16424	599	4927	170	926
AIC	13375.6	138545.8	4297.7	41894.1	1195.2	7795.9
BIC	13409.1	138584.3	4324.1	41926.6	1214.1	7820.1

Note: Adoption Stage**:** Registration denotes the WhatsApp advisory. Enrollment denotes the AI-enabled satellite advisory. The event_reg60 term captures the immediate post-registration effect within the 60-day bandwidth. Spline rows report natural cubic spline basis functions of days since the event with df = 4, which trace the post-event NDVI trajectory rather than day-specific effects.

Clustered standard errors by farmID are in parentheses. Num. Obs. is the number of farm-day observations per model. AIC and BIC are model fit criteria. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

NDVI = Normalized Difference Vegetation Index. RKD = Regression Kink Design. df = degrees of freedom.

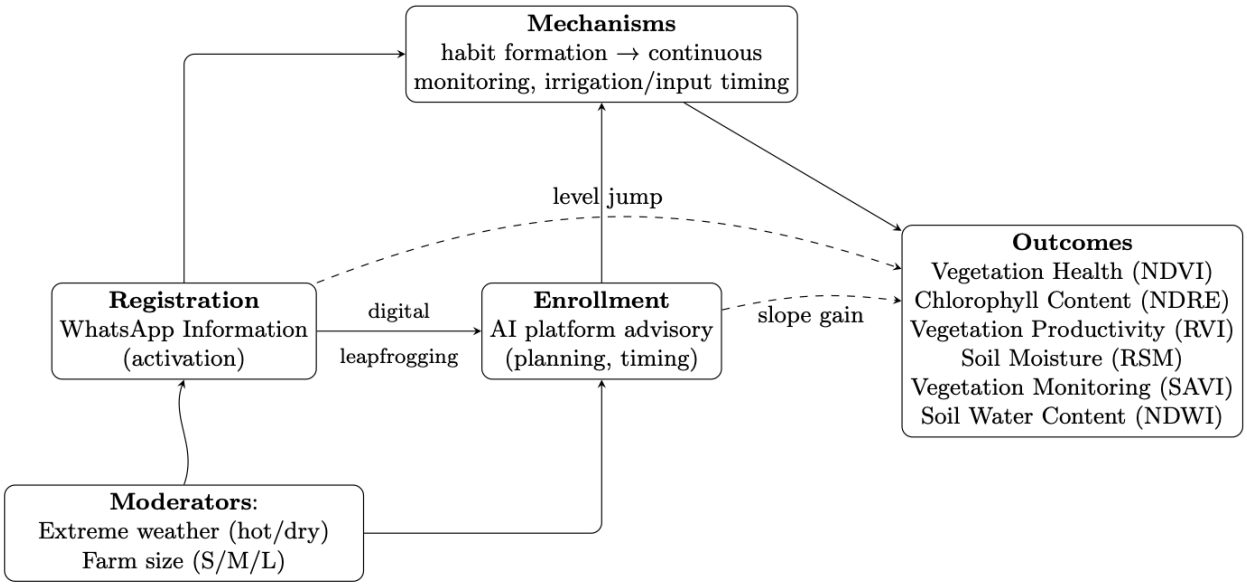


Figure 1: Process Framework for Staged digital adoption in agriculture, with moderators

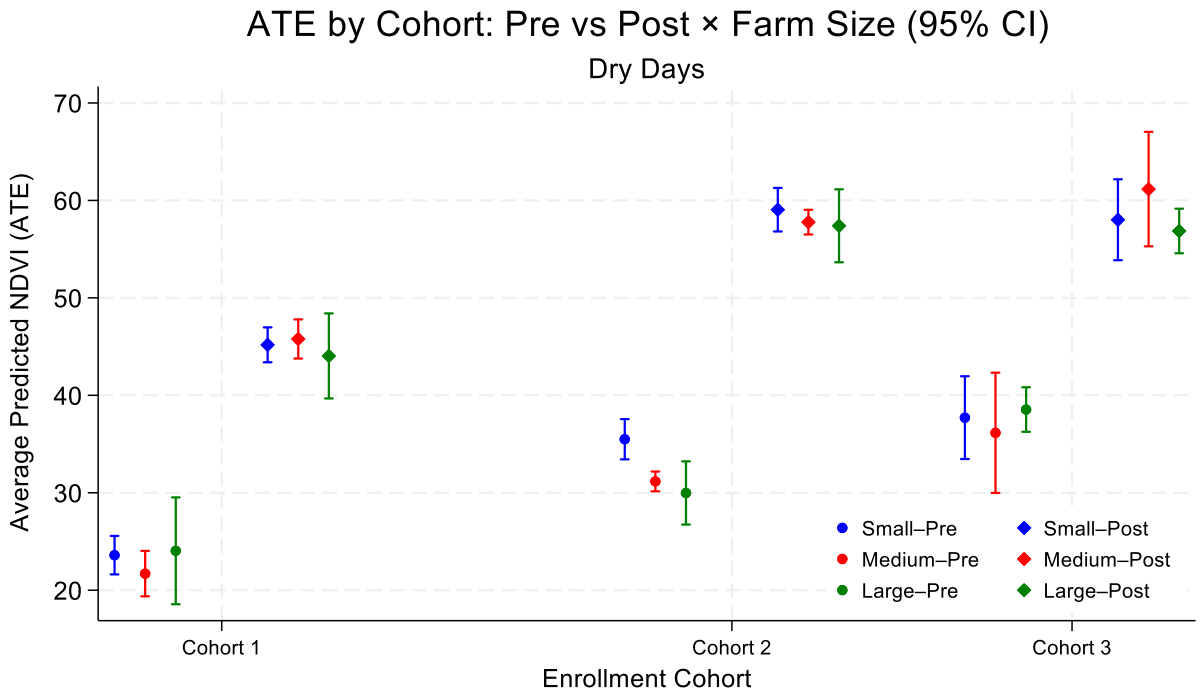


Figure 2. Average Predicted NDVI Before and After Enrollment on Extreme Weather Days

Notes: For the event-time analysis, enrollment occurs on three calendar dates that define cohorts: July 2 (cohort 23193), July 5 (23196), and July 7 (23198). Points are cohort-level average treatment effects for pre and post-periods on cloud-free dry days (precipitation ≤ 4 mm). Error bars are 95% CIs with SEs clustered by farm. Size strata use area-band midpoints: Small ≤ 3,000 m²; Medium 3,000–10,000 m²; Large > 10,000 m². Estimates use farm fixed effects with a ±101-day event window.

APPENDIX

WA1. WhatsApp Advisory

Farmonaut®

Satellite Based Crop Health Monitoring

Field Report

This report consists of satellite results of one of your fields. Automated satellite Monitoring Service allows you to monitor multiple farming fields using latest satellite imagery.

Report Generation Date: 2024-04-10

Satellite imagery capture Date: 2024-4-8

Field Details

Field Address: Farm_magra

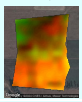
Field Area: 7897 sq m (approx.)

Field Location: Latitude:30.247 Longitude:76.672

SERIAL NO.	TITLE	PAGE NO.
1	Understand the data for better farming	2
2	Weather Statistics for imagery capture data	3
	Weather Forecast for 7 days	
	Weather Graphs (for past 5 days)	4
3	Radar (RVI, RSM)	5
	RVI (Radar Vegetation Index)	
	RSM (Radar Soil Moisture)	
4	Crop Health (NDVI, EVI, SAVI, NDRE)	6
	NDVI (Normalized Difference Vegetation Index)	
	EVI (Enhanced Vegetation Index)	7
	SAVI (Soil Adjusted Vegetation Index)	8
	NDRE (Normalized Difference Red Edge)	9
5	Irrigation (NDWI, NDMI, Evapotranspiration)	10
	NDWI (Normalized Difference Water Index)	
	NDMI (Normalized Difference Moisture Index)	11
	Evapotranspiration	12
6	Soil Health (SOC)	12
7	RGB Satellite Image	13
8	Basic Analysis for Colorblind Visualization	13

Understand the data for better farming

ETCI



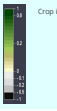
Field Directions





NDVI (For Crop Health)



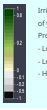


Crop is in good condition



NDWI Image (For Irrigation)





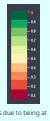
Irrigation inspection is required in south direction of your field

Probable causes for bad irrigation:

- Low Water Quantity in Plants
- Low Soil Moisture
- High Evapotranspiration Rate

DEM (For Flood Analysis)





DEM image tells probable flooding areas due to being at lower terrain.

Your farm is uniformly level/flat

SOC (For Soil Carbon Analysis)





SOC image provides a map of soil organic matter, present at the field.

Soil organic carbon is looking good in your field

Weather Statistics for imagery capture data

Weather Station: Ghanaur

Average Cloud Cover: 1 %

Minimum Temperature: 19.72 deg C

Average Wind Speed: 1.69 m/s

Average Humidity: 15 %

Maximum Temperature: 21.62 deg C

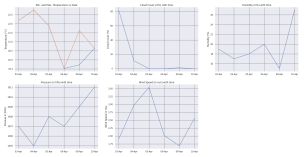
Wind Direction: 13 deg

Average Pressure: 1009 hPa

Weather Forcast (Next 7 days)

Date	Summary	Minimum Temperature (deg C)	Maximum Temperature (deg C)	Rain Probability (%)	Max Precipitation (mm per hr)	Cloud Cover (%)
2024-04-11	few clouds	23.1	39.27	0.0	Not Available	18
2024-04-12	scattered clouds	23.32	40.43	0.0	Not Available	48
2024-04-13	light rain	24.04	39.52	20.0	Not Available	20
2024-04-14	light rain	23.07	37.1	77.0	Not Available	98
2024-04-15	light rain	21.49	37.65	59.0	Not Available	14
2024-04-16	light rain	20.63	36.03	100.0	Not Available	46
2024-04-17	scattered clouds	20.63	36.56	0.0	Not Available	46

Weather Graphs (Using past 5 days data)



Radar (RVI, RSM)

RVI (Radar Vegetation Index)





Scientific Background

Radar Vegetation Index generally ranges between 0 and 1 and is a measure of the randomness of the scattering. RVI is near zero for a smooth bare surface and increases as a crop grows (up to a point in the growth cycle). Use this index for crop health estimation during cloudy weather.

RSM (Radar Soil Moisture)





Scientific Background

soil moisture measures status of plant health based on how plants reflect light at certain frequencies. Though we cannot perceive it with our eyes, everything around us (including plants) reflects wavelengths of light in visible and non-visible spectrum. Taking into account of certain wavelength is reflected, we can assess the current status od plants. If a plant is healthy, it will have large amount of chlorophyll on its leaves and will absorb good amount of visible light from 0.4 to 0.7 microns and reflect quite less of it and vice-versa, we take into account this basic principle in identifying crop health status of agricultural land.

Crop Health (NDVI, EVI, SAVI, NDRE)

NDVI (Normalized Difference Vegetation Index)

NDVI image provides you a colour map of the vegetation of your farming field and nearby areas. The areas shown in red are the regions where the crop growth may not be normal. You should refer to these images when your crop is in early stage of growth.

Use When vegetation is of good height



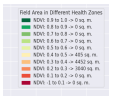
Use When vegetation is of small height









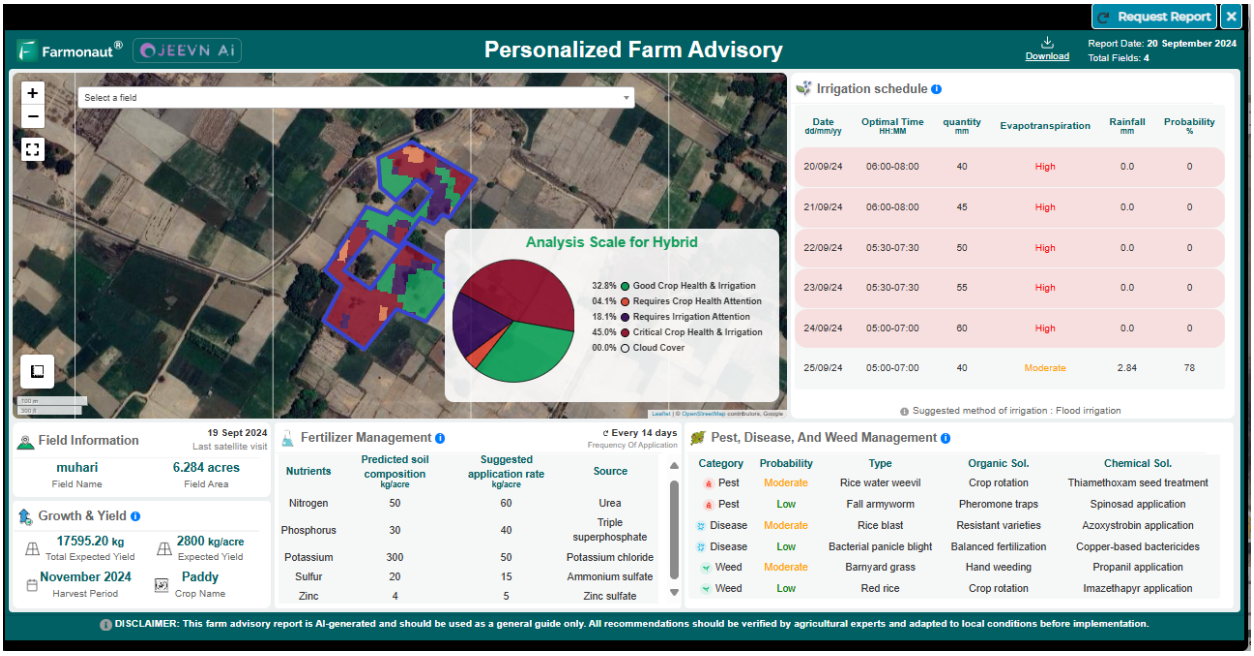




Scientific Background

NDVI measures status of plant health based on how plants reflect light at certain frequencies. Though we cannot perceive it with our eyes, everything around us (including plants) reflects wavelengths of light in visible and non-visible spectrum. Taking into account of certain wavelength is reflected, we can assess the current status od plants. If a plant is healthy, it will have large amount of chlorophyll on its leaves and will absorb good amount of visible light from 0.4 to 0.7 microns and reflect quite less of it and vice-versa, we take into account this basic principle in identifying crop health status of agricultural land.

WA2. JEEVN AI Advisory



- Provides real-time advice on irrigation, fertilization, and pest management.
- Generates 7-day irrigation schedules considering weather patterns and local farming methods.
- Estimates harvest yield using AI-based models analyzing crop health, water status, and soil composition.
- Optimizes fertilizer use by analyzing soil nutrient levels and crop growth stages.
- Detects pests, diseases, and weeds early using satellite imagery and weather patterns.

WA3. Predicted NDVI by Cohort

Figure 1A plots predicted NDVI over event time (days relative to enrollment) separately for each cohort, with Lowess smoothing and 95% confidence intervals. The vertical dashed line marks the enrollment date.

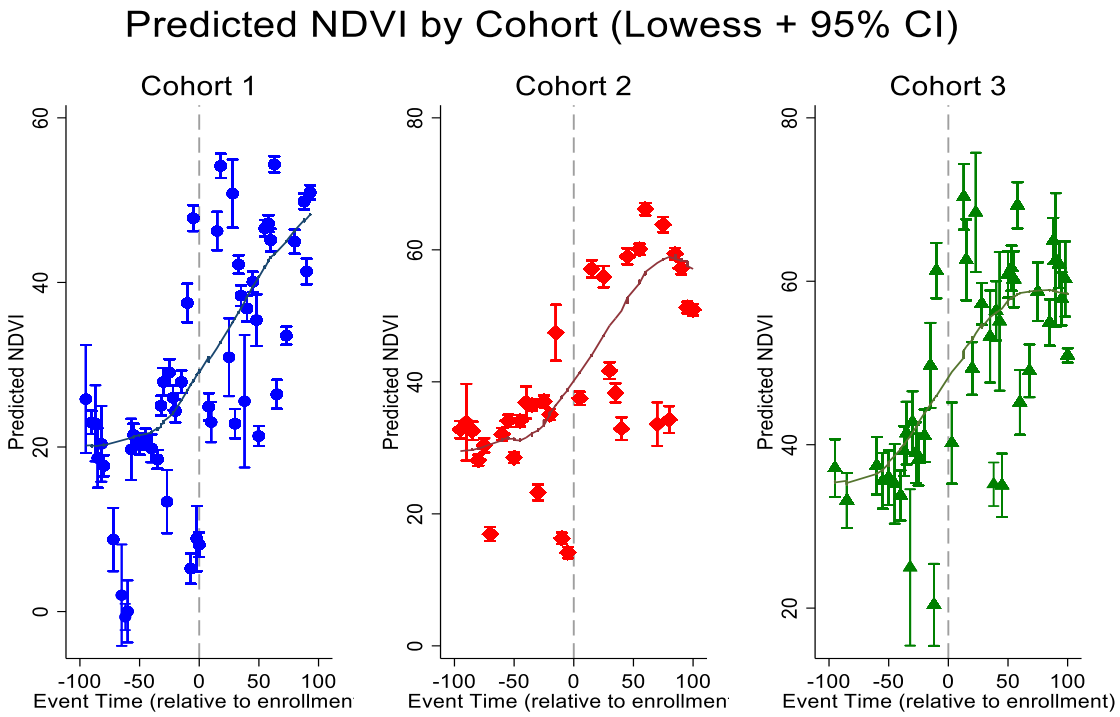


Figure 1A: NDVI over event time (days relative to enrollment)

Cohort 1 (left panel, blue): NDVI values fluctuate before enrollment and show an upward trend afterward, with a wider spread but a clear post-enrollment increase. Cohort 2 (middle panel, red): Pre-enrollment NDVI values are generally lower, with a marked increase after enrollment, suggesting a stronger post-enrollment vegetation response relative to Cohort 1. Cohort 3 (right panel, green): NDVI is more stable before enrollment and exhibits a more consistent and pronounced upward trajectory afterward, with narrower confidence intervals.

Overall, all three cohorts display evidence of higher NDVI levels following enrollment, though the timing and magnitude of responses differ across cohorts. The smoothed trends and confidence bands highlight heterogeneous dynamics, consistent with the RDD event-type analysis and ATE estimates by cohort.

WA4. Predicted NDVI by Cohort and Farm Size

The panels in Figure 2A repeat the baseline specification within farm-size bins, defined from area_band as: Small (0–3000)**,** Medium (3000–10000)**,** Large (10000+)**.** The size bins are created from the lower bound of each band; a midpoint (area_band_midpoint) is also computed for continuous checks. Any cohort differences here should be interpreted as within-size-class heterogeneity, with the same fixed effects and clustering as above.

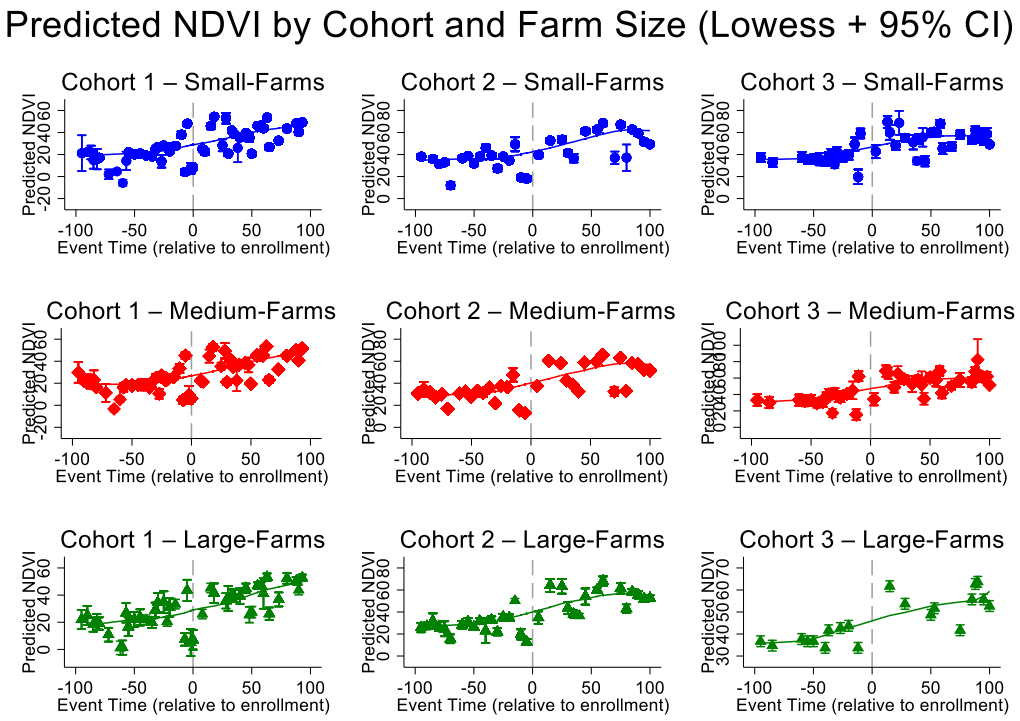


Figure 2A: Predicted NDVI by Cohort and Farm Size

Figure 2A plots predicted NDVI relative to enrollment, separately by cohort (columns) and farm size (rows), with Lowess smoothing and 95% confidence intervals. Across all cohorts, small farms exhibit more volatile NDVI patterns, with less pronounced post-enrollment gains, while medium farms show clearer upward trends after enrollment. Large farms tend to display the most consistent increases in NDVI following enrollment, with narrower confidence intervals, suggesting more stable vegetation responses. Overall, the results highlight heterogeneity across cohorts and farm size, consistent with the ATE estimates reported in the analysis.

Figure 2B shows predicted NDVI around enrollment, disaggregated by cohort (columns) and farm size (rows), for observations flagged as experiencing extreme events (hot, cold, wet, or dry days).

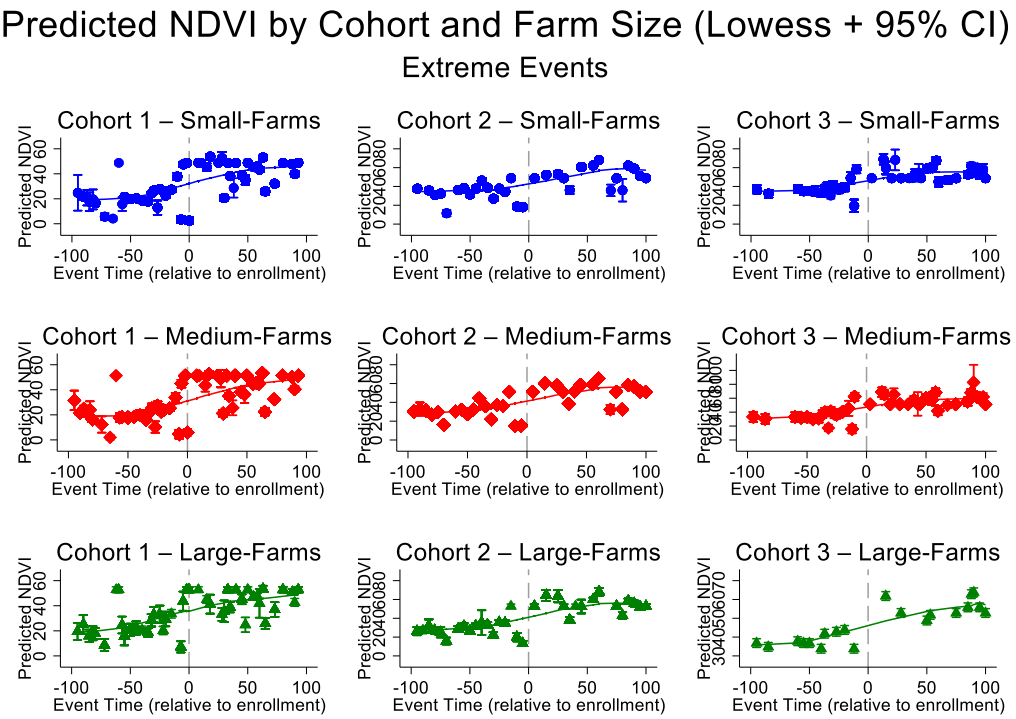


Figure 2B: Predicted NDVI by Cohort and Farm Size under Extreme Weather Events

Definitions used in the plots/controls in Figure 2B: is_hot_day = 1 if max_temp_c ≥ 35°C; is_cold_day = 1 if min_temp_c ≤ 20°C**;** is_wet_day = 1 if precip_mm ≥ 30; is_dry_day = 1 if precip_mm ≤ 4. The indicator extreme_d = 1 if any of hot/cold/wet/dry is flagged that day. Use these thresholds when interpreting any conditional or interacted lines/bars in the figure.

Extreme streaks/spells used in conditioning. Streaks mark consecutive days meeting a condition, with thresholds: Dry streak ≥ 10 dry days; Wet streak ≥ 10 wet days; Cold streak ≥ 2 cold days; Hot streak ≥ 3 hot days. The indicator extreme_s = 1 when a day lies within any of these streaks. Use these definitions whenever a panel or line is labeled by streak status.

WA5. NDVI Pre vs Post Enrollment

Figure 3A compares predicted NDVI values before (blue) and after (red) enrollment across the three cohorts.

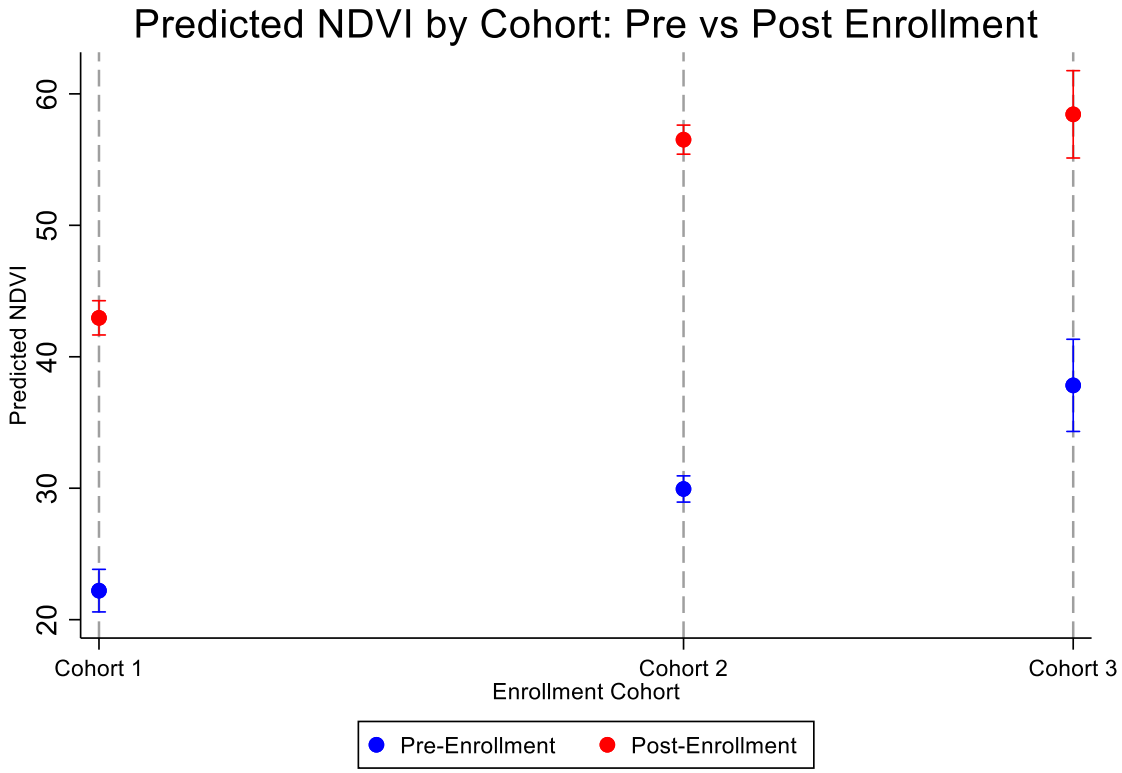


Figure 3A: Predicted NDVI by Cohort: Pre vs Post Enrollment

In all cases, NDVI is significantly higher in the post-enrollment period, with Cohorts 2 and 3 showing the strongest increases. The confidence intervals indicate that these post-enrollment gains are statistically meaningful, highlighting clear differences between pre- and post-enrollment vegetation outcomes.

Figure 3B shows estimated average treatment effects (ATEs) of NDVI by cohort-month for three cohorts (23193 (July 2) in blue, 23196 (July 5), in red, 23198 (July 7) in green), with 95% confidence intervals.

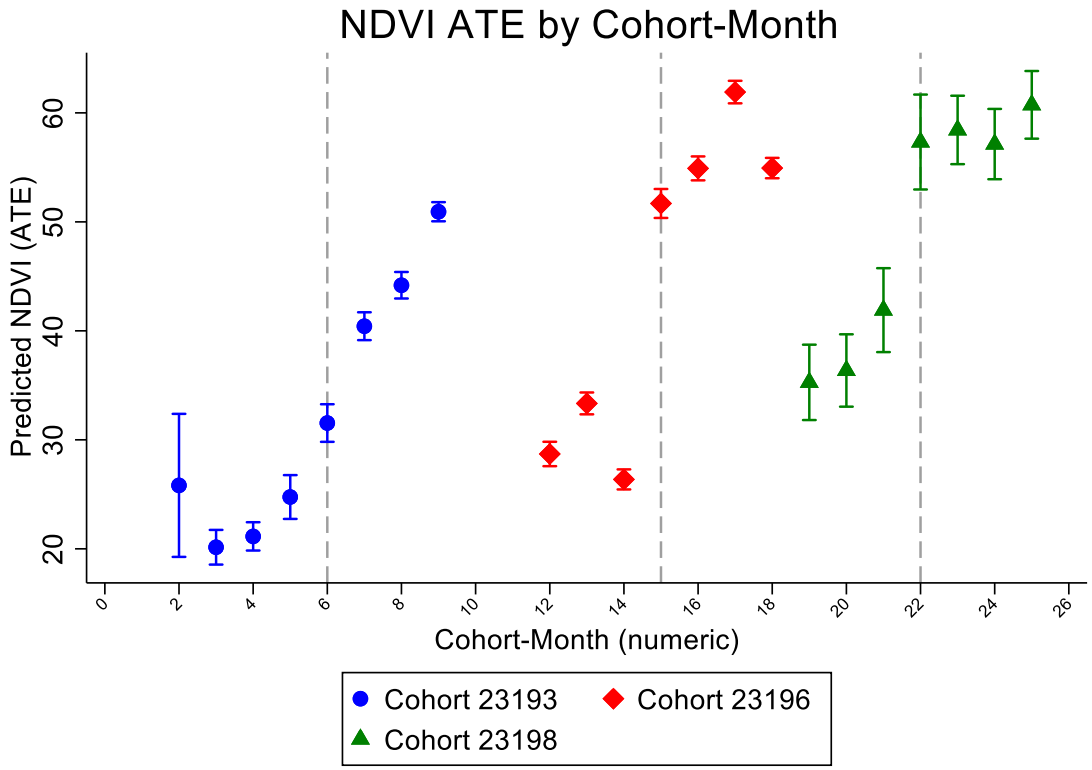


Figure 3B: NDVI (ATE) by Cohort-Month

NDVI levels rise steadily across cohort-months, with clear post-enrollment increases relative to the pre-enrollment baseline. The magnitude of the ATE differs across cohorts, indicating heterogeneous vegetation responses over time and across enrollment groups.

WA6. NDVI (ATE) by Cohort: Pre vs Post with Farm Size

Figure 4A presents the average predicted NDVI (ATE) by cohort, distinguishing pre- and post-enrollment periods across farm size groups (small, medium, large).

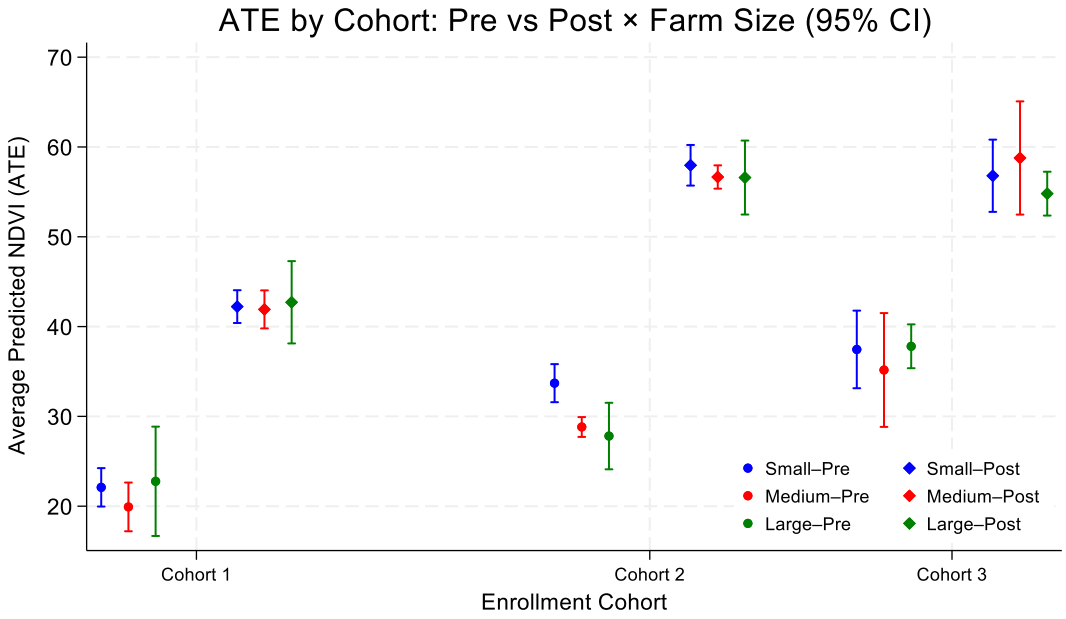


Figure 4A: NDVI (ATE) by Cohort: Pre vs Post with Farm Size

Across all cohorts, NDVI is consistently higher after enrollment, with medium and large farms showing the most significant and stable gains. In contrast, small farms exhibit greater variability and comparatively modest improvements. The results emphasize heterogeneity across cohorts and farm size, confirming that program effects differ depending on farm characteristics and timing relative to enrollment.

Figure 4B presents average predicted NDVI (ATE) by cohort, comparing pre- and post-enrollment periods across farm size groups (small, medium, large) during spells of extreme events (hot, cold, wet, or dry streaks).

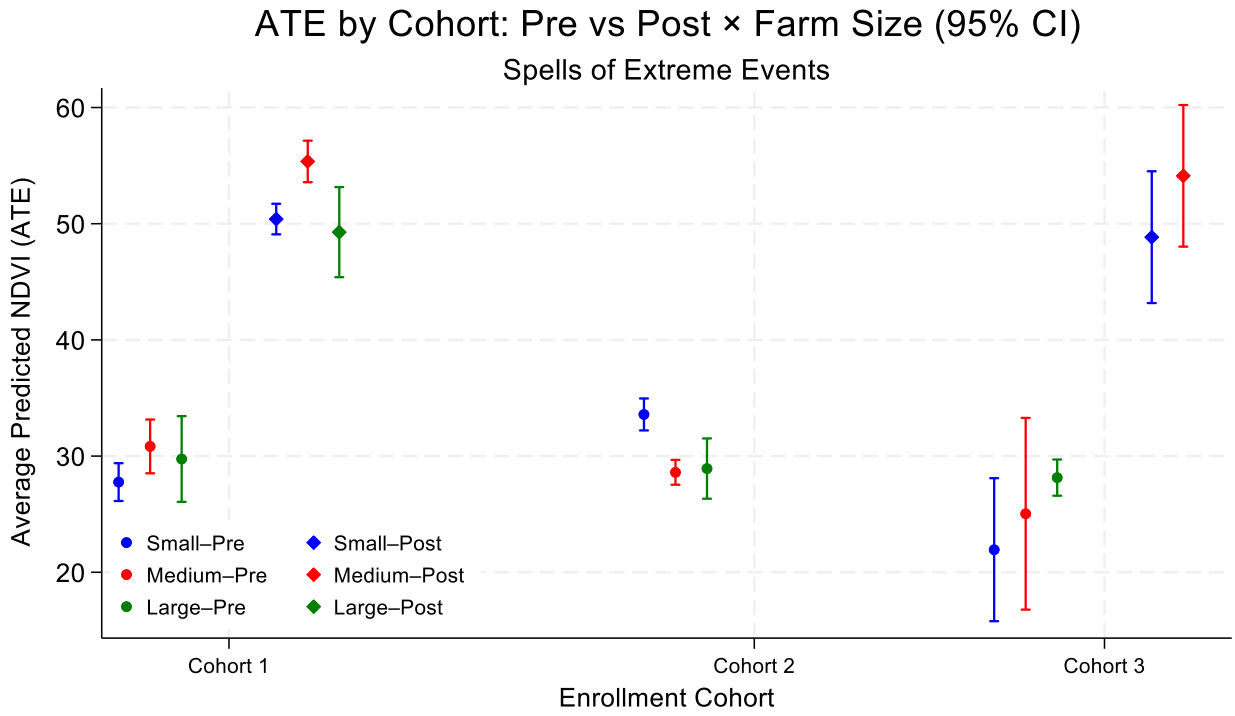


Figure 4B: NDVI (ATE) by Cohort: Pre vs Post with Farm Size under Extreme Weather Events

Across cohorts, post-enrollment NDVI is generally higher, though responses are more variable under extreme conditions. These patterns suggest that farm size and exposure to extreme weather are essential in shaping vegetation outcomes around enrollment.