Volume 1, Issue 1 (Inaugural & Special Issue), July-September 2026

Accessibility, Viability, and Groundwater Sustainability of Drip Irrigation in Telangana: An Integrated Micro Econometric Identification Framework

Dr.M.Rukmani

Associate Professor, Department of Economics, SVCR Government Degree College,

Palamaner, Chittoor District, Andhra Pradesh, PIN- 517408,

svcreconomics@gmail.com

Abstract:

Academic literature frequently stalls at the border of empirical observation, leaving administrative praxis untouched. The South India Journal of Economic Policy and Development (SIJEPD) directly confronts this disconnect. Its editorial mandate prioritizes persistent structural hurdles across the region. These span from uneven wealth distribution and fiscal stagnation to labor market volatility and systemic failures in public utility architecture. Research must deliver more than simple diagnostics. SIJEPD operates as a translational mechanism, converting localized data into actionable legislative frameworks. The journal aggregates cross-disciplinary intelligence from economists, civic stakeholders, and industry practitioners. This synthesis produces scholarship engineered for immediate implementation. Institutional inertia often thrives on a deficit of reliable, context-specific metrics. By anchoring public discourse in rigorous methodology, the publication demands a shift toward evidence-backed governance. It supplies policymakers with the precise theoretical and empirical tools required to navigate complex socioeconomic realities.

Keywords: drip irrigation; micro-irrigation; technology adoption; endogenous switching regression; rebound effect; Jevons paradox; groundwater; spatial econometrics; regression discontinuity; Telangana.

JEL classification: Q15, Q12, Q25, Q18, O13, C21, C31, C36, C26.

1. Introduction

Groundwater functions as the primary hydraulic engine for agriculture across Telangana. Wells currently irrigate over half the state's cropped acreage, supported by millions of active agricultural pump connections. Combined with erratic monsoon patterns and a historical affinity for water-intensive crops, this relentless extraction has pushed regional water tables down to an average depth exceeding 10 meters. The resulting over-exploitation is particularly acute across districts like Vikarabad, Nagarkurnool, and Yadadri Bhongir.

To counter this depletion, policymakers heavily promote micro-irrigation. Drip systems, specifically, have emerged as the primary tool for demand-side water management. By restricting water delivery directly to the root zone, drip irrigation substantially outperforms traditional flood methods in on-farm efficiency. It simultaneously reduces the energy required for pumping, minimizes labor demands, and cuts fertilizer waste. Through the Telangana State Micro Irrigation Project (TSMIP), the government injects massive capital subsidies into the sector. For the most disadvantaged farming demographics, these subsidies can cover the entire cost of installation.

Yet a distinct policy paradox complicates this narrative. Despite Telangana offering some of the most generous subsidies in India, regional aquifer stress continues to escalate. Adoption remains highly asymmetric across different crops, geographic pockets, and socioeconomic classes. This divergence raises three critical lines of inquiry regarding accessibility, viability, and sustainability. First, can marginalized farmers actually access this technology? Restrictive land record mandates, bureaucratic documentation, steep co-payment barriers, and poor after-sales dealer networks may inadvertently screen out the exact populations the policy intends to support. Second, is drip irrigation genuinely viable from an economic standpoint once we account for the true friction of installation costs, subsidy delays, and actual yield responses? Finally, does widespread diffusion actually protect groundwater reserves? If farmers reinvest their plot-level water savings into expanding their cultivated acreage or shifting toward thirstier crops—a phenomenon known as the rebound or Jevons effect—aggregate extraction might remain stagnant or even increase (Grafton et al., 2018; Pfeiffer & Lin, 2014).

Answering these questions requires looking past standard empirical frameworks. Technology adoption is inherently endogenous. Farmers who willingly install drip systems often possess unobservable traits—such as higher managerial acumen or different risk tolerances—that separate them from non-adopters. Any naive observational comparison between these groups inevitably conflates the actual technological impact with baseline selection bias. Beyond mere adoption, securing a subsidy is rarely an exogenous event; it is dictated by strict eligibility criteria and administrative navigation.

Groundwater adds another layer of complexity as a common-pool resource. When one farmer extracts water, the localized water table drops, directly impacting neighboring plots. This spatial dependence inherently violates the stable unit treatment value assumption (SUTVA). Credible evaluation, therefore, cannot rely on descriptive contrasts or basic single-equation models. It demands a rigorous, explicitly identified causal architecture.

Rather than reporting ex post estimates, this paper introduces a comprehensive methodological design. We develop an integrated identification framework to analyze accessibility, viability, and sustainability through a single analytical structure, establishing a pre-specified empirical strategy for future primary data evaluation.

We frame these core themes as linked but distinct causal estimands, carefully matching each to an appropriate estimator. To evaluate viability, we utilize an endogenous switching regression. This approach corrects for unobservable selection bias and isolates treatment effects for both adopters and non-adopters—yielding the exact counterfactuals required for scaling policy. We complement this with doubly robust matching estimators and marginal treatment effect models to map underlying heterogeneity. Addressing sustainability requires shifting focus away from simple plot-level accounting. We identify the behavioral rebound effect by deploying a difference-in-differences estimator built to handle staggered adoption timelines and heterogenous impacts. A spatial Durbin model is then applied to quantify common-pool spillover dynamics, ultimately deriving a structural rebound ratio. Finally, we capitalize on the sharp cutoffs in administrative subsidy entitlement. Using a fuzzy regression-discontinuity design, we can directly link farmer-level outcomes back to the policy instrument itself. By explicitly detailing identifying assumptions, potential instruments, inference procedures, and falsification tests, we ensure the architecture remains entirely transparent and primed for pre-registration prior to data collection.

Subsequent sections delineate this proposed architecture. Section 2 establishes the conceptual framework and defines the primary estimands, modeling drip adoption as a binary choice heavily mediated by institutional subsidies. We deliberately isolate the three causal questions here, as each introduces distinct identification hurdles. Section 3 details the intended study design, data structure, and measurement protocols. Section 4 functions as the analytical core, deconstructing the specific empirical strategy tailored to each dimension of the evaluation. Following this, Section 5 dictates mechanisms for robust inference, falsification testing, and reporting standards. We critically assess potential threats to internal validity and structural limitations in Section 6, before concluding in Section 7 with an overview of the broader policy dilemmas these prospective estimates aim to resolve.

2. Conceptual framework and estimands

We treat drip adoption as a binary technology choice embedded within a heavily subsidized institutional environment. The three core causal questions are kept strictly separated, as each carries its own unique identification problem.

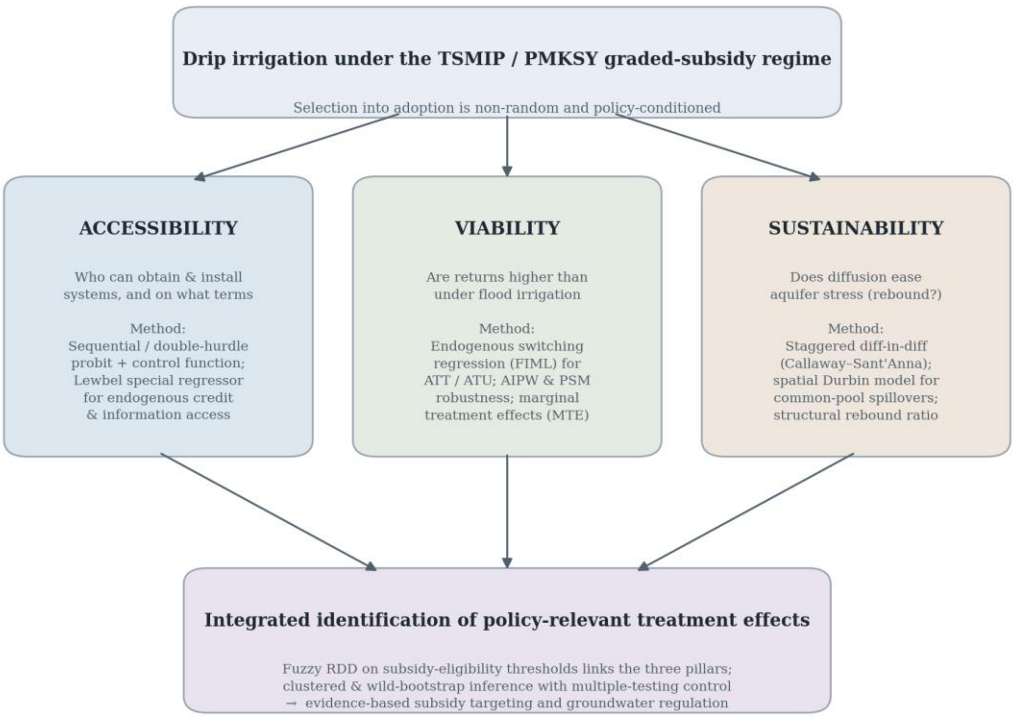

Figure 1. Integrated identification framework: research dimensions and matched econometric methods.

2.1 Accessibility

Accessibility concerns who can obtain and install a system and on what terms. Conceptually, it is a sequential process: a farmer must first become eligible and aware (cross the access hurdle) before deciding whether to adopt. The estimand of interest is not merely the probability of adoption but the set of barriers — documentation, land title, co-payment liquidity, dealer distance — that differentially screen out small and marginal farmers, conditional on the same agronomic incentives. Because credit and information access are themselves choices correlated with unobservables, the estimand requires a model that separates the access hurdle from the adoption decision and corrects for endogenous regressors.

2.2 Viability

Viability is a causal-effect question. Let Di ∈ {0,1} indicate adoption, and let Yi(1) and Yi(0) denote the potential outcomes (net returns per acre, water productivity, energy use) under adoption and non-adoption. The individual effect Yi(1) − Yi(0) is never observed for the same farmer. The policy-relevant estimands are the average treatment effect on the treated (ATT) and on the untreated (ATU):

$$\text{ATT} = \text{E}[Y(1) - Y(0) \mid D = 1), \qquad \text{ATU} = \text{E}[Y(1) - Y(0) \mid D = 0). \qquad (1)$$

ATT answers whether drip pays for those who choose it; ATU answers what non-adopters would gain if the program reached them — the relevant quantity for an expansion policy. Because D is self-selected, E[Y | D = 1] − E[Y | D = 0] is a biased estimator of either; identification must confront selection on unobservables.

2.3 Sustainability and the rebound effect

Sustainability asks whether diffusion reduces aggregate extraction, not merely per-hectare application. Decompose farm water use as the product of application intensity, irrigated area, and cropping intensity:

$$W_i = w_i \times A_i \times c_i, \qquad (2)$$

where W is total water extracted, w is the application per irrigated acre, A is the irrigated area, and c is the cropping intensity (including any shift to thirstier crops). Drip lowers with the engineering presumption that W falls holds only if A and C do not respond. The rebound effect is the behavioral re-expansion of A and C from the saved water and the lower marginal cost of irrigation. We summarise it by a rebound ratio.

$$R = 1 - (\Delta W^{obs} / \Delta W^{eng}), \quad (3)$$

where ΔW^{eng} is the engineering (mechanical) water saving holding area and crop mix fixed, and ΔW^{obs} is the observed change in extraction. $R = 0$ means full realization of engineering savings; $R = 1$ means complete rebound; $R > 1$ means backfire (extraction rises). Because extraction also depends on neighbors' behavior through the shared aquifer, identification must account for spatial dependence.

2.4 Exploiting the Subsidy Discontinuity

The TSMIP graded-subsidy schedule enforces sharp administrative cutoffs. Farmer demographics—specifically SC/ST or small/marginal classifications—and landholding size dictate precise subsidy rates. Farmers clustered directly around these thresholds are essentially identical in every respect except for the specific subsidy tier they qualify for. This sudden shift creates an optimal environment for quasi-experimental variation. We leverage this discontinuity as the primary analytical instrument. It acts as the anchor for three key dimensions of the study. First, it isolates accessibility by shifting the initial cost of entry. Second, it allows us to identify economic viability through exogenous shifts in technology adoption. Finally, when integrated into our extraction model, this variation reveals the broader sustainability impacts driven directly by the policy itself.

3. Methodology and Data Architecture

3.1 Sampling Design

We employ a multi-stage stratified sampling approach, prioritizing the identification of causal treatment effects over simple population description. Sample sizes were not based on fixed population proportions. Instead, we relied on rigorous power calculations designed to detect the minimum policy-relevant effect while accounting for expected clustering.

District selection was purposive. We needed to capture highly contrasting aquifer conditions. Consequently, severely over-exploited districts with deep water tables (such as Nagarkurnool and Vikarabad) were paired against regions maintaining relatively stable groundwater reserves but demonstrating active technology uptake. This deliberate contrast enables the estimation of treatment effects across vastly different hydrological realities. Downstream sampling stages focused on capturing varying adoption intensities at the mandal and village levels. At the granular farm level, participants were stratified to ensure balanced cohorts of adopters and non-adopters across large, medium, and small landholdings. To maintain statistical validity, village and mandal-level clustering informed both the sampling frame and our inferential models.

3.2 Synthesis of Primary and Secondary Data

The empirical strategy relies on integrating cross-sectional field data with longitudinal administrative records.

Primary data collection centered on a comprehensive farm-household survey. This instrument captured detailed plot rosters, input-output financial accounts, well specifications, and specialized irrigation and energy logs. To contextualize these quantitative metrics, we conducted focus group discussions alongside semi-structured interviews with key local informants, including extension workers, equipment dealers, and regional officials.

Secondary datasets provided the critical longitudinal and administrative backbone for the quasi-experimental models:

- Hydrological Records: Time-series data on water tables and mandal-level extraction volumes from the Central Ground Water Board and the Telangana

Ground Water Department. These series supply the panel data necessary to offset the limitations of our cross-sectional survey.

- Program Administration: TSMIP disbursement and sanction logs from the Department of Horticulture. These records establish the exact timing variation required for the difference-in-differences (DiD) framework. Crucially, they also provide the precise eligibility cutoffs and running variables needed to execute the regression discontinuity (RD) design.
- Contextual Baselines: Supplementary reports from the Agricultural Census, the Directorate of Economics and Statistics, and the PMKSY (Per Drop More Crop) initiative.

3.3 Variable Operationalization

Table 1 outlines the core variables, their dimensional utility, and specific measurement protocols. Key economic and physical metrics were operationalized as follows:

Net returns are calculated comprehensively. We subtracted both imputed and actual paid-out costs—including the amortized capital cost of drip infrastructure, net of the applied subsidy—from the gross output value. Water productivity is defined strictly as physical crop output per cubic meter of applied irrigation. Finally, we estimated farm-level groundwater extraction by multiplying recorded active pumping hours by the corresponding pump discharge, derived from either direct metering or verified engineering estimates.

Role	Variable	Measurement/notes
Treatment	D — drip adoption	1 if a functioning drip system irrigates any plot in the reference year; robustness to area share and years-since-installation.
Outcomes (viability)	Net returns/acre; water productivity; energy use; yield	Full-cost net returns (₹/acre); output per m^3; pumping kWh/acre; physical yield (kg/acre), all for comparable crops.
Outcomes (sustainability)	Farm extraction W; water-table depth	Extraction = pump discharge × hours (m^3/season); mandal water-table depth (m bgl) from secondary series for the panel design.
Covariates	X — farm & farmer	Farm size, tenure security, age, education, household labor, soil class, crop, well depth/age, assured electricity hours, distance to market.
Access/barriers	Documentation, liquidity, dealer reach, extension	Land-record completeness; ability to meet co-payment; distance to nearest authorized dealer; number of extension contacts; awareness of TSMIP.

Role	**Variable**	**Measurement/notes**
Instruments	Z — exclusion restrictions	Distance to nearest dealer/bank; historical (pre-period) extension placement; village-level peer adoption excluding own network; rainfall shock for liquidity.
RDD running variable	Landholding/category score	Distance to the subsidy-eligibility cut-off; eligibility indicator above/below threshold.

Table 1. Variables, the dimension served, and measurement conventions.

4. Empirical strategy

4.1 Accessibility: sequential discrete choice with a control function

Adopption is modeled as the outcome of a latent net-benefit index. Farmer I adopts when the index is positive:

$$A_i^* = Z_i'\gamma + u_i, \qquad D_i = 1[A_i^* > 0]. \quad (4)$$

Estimating (4) using a simple probit or logit recovers correlates of adoption. However, it conflates two stages — gaining access and choosing to adopt — and treats credit and information as exogenous, even though they are not. We therefore make two refinements. First, we model the access hurdle and the adoption decision sequentially (a double-hurdle/sequential-probit structure), so that barriers operating at the access stage are separated from preferences operating at the adoption stage; this isolates the accessibility estimand of policy interest. Second, endogenous regressors Diend (credit access, extension exposure) are handled by a control-function (Rivers–Vuong) approach: a first-stage reduced form

$$D_i^{end} = Z_i'\pi + v_i \qquad (5)$$

These yield residuals vˆi that enter the structural probit as an additional regressor; the significance of its coefficient directly tests for endogeneity, and its inclusion purges bias. When no convincing external instrument is available, identification can be supported by Lewbel's (2012) heteroskedasticity-based instruments. Candidate exclusion restrictions Z — distance to the nearest authorized dealer and bank, pre-period extension placement, and rainfall-driven liquidity shocks — are argued to shift access without independently affecting the adoption preference conditional on covariates, and their plausibility is examined in Section 5.

Hypothesis H2 (that access is not neutral across farmer categories) is tested by examining the partial effect of small/marginal status on the access-stage probability, conditional on the agronomic incentives that enter the adoption stage; a negative and significant effect indicates that higher effective barriers offset favorable subsidy provisions.

4.2 Viability: endogenous switching regression with treatment-effect recovery

The core viability estimator is the endogenous switching regression (ESR), which jointly models the selection equation and the regime-specific outcome equations and corrects for selection on unobservable (Heckman, 1979; Lokshin & Sajaia, 2004; Di Falco et al., 2011). The selection equation is given by (4). The outcome for farmer i is observed in one of two regimes:

$$\text{Regime 1 } (D = 1): \quad Y_{1i} = X_i'\beta_1 + \sigma_1\lambda_{1i} + \eta_{1i}, \tag{6}$$

$$\text{Regime 0 } (D = 0): \quad Y_{0i} = X_i'\beta_0 + \sigma_0\lambda_{0i} + \eta_{0i}, \tag{7}$$

where $\lambda_{1i} = \varphi(Z_i'\gamma) / \Phi(Z_i'\gamma)$ and $\lambda_{0i} = -\varphi(Z_i'\gamma) / [1 - \Phi(Z_i'\gamma)]$ are the inverse Mills ratios from the selection equation (φ and Φ denote the standard normal density and c.d.f.). The terms $\sigma_1 = \rho_1\sigma_{\eta 1}$ and $\sigma_0 = \rho_0\sigma_{\eta 0}$ capture the covariances between the selection error you and the outcome errors η_1, η_0. The system is estimated by full-information maximum likelihood, which is more efficient than the two-step procedure and yields correct standard errors. A statistically significant ρ_1 or ρ_0 provides direct evidence of selection on unobservables — precisely the bias a naive comparison ignores; their signs indicate whether selection is positive (high-return farmers adopt) or negative.

From the estimated regime equations, we form the conditional expectations of the actually observed and their counterfactuals, and subtract them to obtain the treatment effects. For adopters:

$$E[Y_{1i} \mid D = 1] = X_i'\beta_1 + \sigma_1\lambda_{1i}, \quad E[Y_{0i} \mid D = 1] = X_i'\beta_0 + \sigma_0\lambda_{1i}, \tag{8}$$

$$ATT = E[Y_{1i} \mid D = 1] - E[Y_{0i} \mid D = 1] = X_i'(\beta_1 - \beta_0) + (\sigma_1 - \sigma_0)\lambda_{1i}. \tag{9}$$

The ATU is constructed symmetrically using λ_{0i}. Comparing observed and counterfactual expectations also yields transitional-heterogeneity effects that indicate whether the gains to adopters exceed what non-adopters would obtain, which speaks directly to the wisdom of broad versus targeted expansion. Hypothesis H3 (drip yields higher net returns and water productivity) is tested as ATT > 0 and ATU > 0 for comparable crops, with effects disaggregated by crop and farm-size class.

4.2.1 Robustness and complementary estimators

ESR identifies effects using both an exclusion restriction and the joint-normality functional form; we therefore triangulate. Propensity-score matching (nearest-neighbour and kernel, on common support, with covariate-balance diagnostics) estimates the ATT under selection on observables only and provides a lower-structure benchmark. Augmented inverse-probability-weighted regression adjustment (AIPW) is doubly robust — consistent if either the propensity-score model or the outcome model is correct — and is our preferred selection-on-observables estimator. Rosenbaum bounds quantify how strong an unobserved confounder would need to be to overturn the matching result. Where ESR and the selection-on-observables estimators agree, the effect is robust to the nature of selection; where they diverge, the gap itself is informative about the role of unobservables.

4.2.2 Essential heterogeneity: marginal treatment effects

If returns vary systematically with unobserved resistance to adoption, a single ATT or ATU masks policy-relevant heterogeneity. We therefore estimate marginal treatment effects (Heckman & Vytlacil, 2005) using local instrumental variables, MTE(x, p) = ∂E[Y | X = x, P = p] / ∂p, where p is the propensity to adopt. Integrating MTE with appropriate weights recovers ATT, ATU, and — most useful for policy — the policy-relevant treatment effect of a marginal subsidy expansion that attracts the next set of currently reluctant adopters. A downward-sloping MTE in resistance is the signature of selection on gains and a caution against assuming that observed adopter returns generalize to future adopters.

4.3 Sustainability: rebound identification with staggered DiD and spatial spillovers

The sustainability question cannot be answered from the cross-section alone, because the relevant object is the change in *aggregate* extraction over time and space. We combine three strategies.

4.3.1 Staggered difference-in-differences on the groundwater panel

TSMIP sanctions arrive at different times across mandals, resulting in staggered treatment timing in the secondary water table and extraction series. The conventional two-way fixed-effects (TWFE) event-study regression

$$Y_{mt} = \alpha_m + \delta_t + \beta D_{mt} + \varepsilon_{mt} \quad (10)$$

is now known to be biased when effects are heterogeneous across adoption cohorts or vary over time, because already-treated units serve as controls (Goodman-Bacon, 2021; de Chaisemartin & D'Haultfœuille, 2020). We therefore estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator, ATT(g, t), and aggregate them into overall and event-time (dynamic) effects. We also report the Sun and Abraham (2021) interaction-weighted event study as a cross-check. The outcome is the mandal water-table depth and, where available, total extraction; treatment is drip diffusion intensity. A pre-trend (placebo) test on event-time leads is the key identifying check: parallel pre-treatment paths support a causal interpretation, whereas their absence flags confounding by the very water stress that drives both adoption and decline.

4.3.2 Common-pool spillovers: the spatial Durbin model

Groundwater violates the stable-unit-treatment-value assumption: a farmer's extraction depends on neighbours through the shared aquifer, so adoption generates spillovers that a unit-level regression misattributes. We therefore estimate a spatial Durbin model (LeSage & Pace, 2009),

$$W = \rho\, W_* W + X\beta + W_* X\theta + \varepsilon, \qquad (11)$$

where W* is a spatial-weights matrix constructed from hydrogeological contiguity or inverse distance (row-standardized), ρ is the spatial-autoregressive parameter, and θ is the coefficient on spatially lagged covariates, including neighbors' adoption. Crucially, the model is interpreted via its decomposition into direct and indirect (spillover) impacts rather than its raw coefficients: the direct impact is a farmer's own adoption effect on own extraction, and the indirect impact is the effect of neighbours' adoption on a farmer's extraction — the common-pool externality. Spatial dependence is first established using Moran's I and Lagrange-multiplier tests; the model choice (lag, error, or Durbin) follows from those diagnostics.

4.3.3 The structural rebound ratio

Combining the plot-level engineering savings (from agronomic norms and metered application) with the estimated behavioural responses in irrigated area and crop mix, we compute the rebound ratio R in equation (3) and decompose ΔWobs into the application-intensity, area, and cropping-intensity channels in equation (2). Adoption is instrumented in the area- and extraction-response equations using subsidy eligibility and dealer distance, because area expansion and adoption are jointly determined. Hypothesis H4 — that drip cuts per-acre extraction but the rebound from area expansion and crop-mix shifts moderates, and may overturn, the aggregate effect — is tested by the sign and magnitude of R and by whether the indirect spatial impacts and the dynamic DiD effect on mandal water tables are non-negative despite negative per-acre effects.

4.4 The subsidy lever: a fuzzy regression-discontinuity design

When the TSMIP subsidy rate changes discontinuously at an eligibility threshold (a landholding or category cut-off), farmers just above and just below the threshold are comparable but face different entry prices. Because eligibility raises but does not deterministically fix adoption, the design is fuzzy: eligibility instruments for actual subsidized adoption. Let ri be the running variable and c the cut-off. The local average treatment effect at the threshold is

$$\tau_{RD} = \lim_{r\downarrow c} \mathrm{E}[Y|r] - \lim_{r\uparrow c} \mathrm{E}[Y|r] \quad \div \quad \lim_{r\downarrow c} \mathrm{E}[D|r] - \lim_{r\uparrow c} \mathrm{E}[D|r]\ . \quad (12)$$

Estimation uses local polynomial regression with robust, bias-corrected confidence intervals and data-driven bandwidth selection (Calonico, Cattaneo, and Titiunik, 2014). Validity rests on the continuity of potential outcomes at the cut-off and the absence of manipulation of the running variable, which we test using the McCrary (2008) density test and covariate-continuity (balance) checks at the threshold; estimates are reported across a range of bandwidths and polynomial orders. This design answers Hypothesis H5 — how subsidy structure shapes adoption and, through the extraction model, sustainability — with the cleanest available identification, at the cost of a locally defined effect.

4.5 Summary: questions, methods, and expected signs

Hyp.	**Question/estimand**	**Primary method (robustness)**	**Expected sign/test**
H1	Determinants of adoption (farm size, education, credit, extension, subsidy info)	Sequential probit with control function (Lewbel IVs)	Positive on size, education, credit, extension; endogeneity test on $\hat{\nu}$
H2	Access is not neutral across categories; barriers to small/marginal	Access-stage hurdle equation; FGD/KII triangulation	Negative effect of small/marginal status at the access stage

Hyp.	Question/estimand	Primary method (robustness)	Expected sign/test
H3	Viability: net returns and water productivity vs flood	Endogenous switching regression — ATT, ATU (AIPW, PSM, MTE)	ATT > 0 and ATU > 0; crop- and size-specific
H4	Per-acre saving vs aggregate effect (rebound)	Callaway–Sant'Anna DiD; spatial Durbin; rebound ratio R	Per-acre w ↓ but 0 < R ≤ 1 (or > 1); indirect impacts ≥ 0
—	Causal role of the subsidy lever	Fuzzy RDD at eligibility cut-off (bandwidth sensitivity)	τ(RD) > 0 on adoption; sign on extraction via model

Table 2. Mapping of hypotheses to estimands, methods, and pre-specified expected signs. Cells report directions to be tested, not estimates.

5. Inference, Falsification, and Reporting

Let us first address the mechanics of standard error clustering. Because treatment assignment occurs strictly at the village and mandal levels, relying on asymptotic cluster-robust standard errors risks severe overrejection—a well-documented hazard when the cluster count remains limited. To counteract this, primary inference for all headline effects relies on wild-cluster bootstrap p-values (Cameron, Gelbach, & Miller, 2008). Multiple hypothesis testing, an inevitable concern when evaluating diverse subgroups and outcomes, is managed by controlling both the family-wise error rate and the false-discovery rate via Romano-Wolf stepdown procedures and sharpened q-values. We also mitigate the risk of selective reporting by strictly pre-registering a concise set of primary outcomes: net returns, water productivity, and mandal-level water-table depth.

The identification strategy hinges on a rigorous falsification framework. Depending on the specific model, we deploy instrument-relevance and over-identification statistics, alongside theoretical defenses of the exclusion restrictions required for both the control-function and endogenous switching regression (ESR) stages. For the difference-in-differences (DiD) components, we run standard pre-trend and placebo-lead tests. The spatial models undergo Moran's I and Lagrange Multiplier diagnostics, heavily supplemented by sensitivity checks on the spatial-weights matrices. Meanwhile, validating the regression-discontinuity (RDD) design requires McCrary density checks and covariate-balance testing across varying bandwidths and polynomial specifications. Rosenbaum bounds are utilized to evaluate how vulnerable the matching estimates might be to unobserved confounding. Where feasible, the core estimands, primary outcomes, and predicted parameter signs are lodged in pre-analysis registries, effectively stripping the exercise of exploratory flexibility.

Results will populate three pre-formatted reporting matrices. A viability table juxtaposes the dual ESR regime equations, detailing the selection-correlation parameters (ρ_1, ρ_0), the resulting ATT and ATU with bootstrap intervals, and benchmark estimates derived from AIPW and PSM models. A separate sustainability table captures the dynamic Callaway–Sant'Anna DiD parameters. This table isolates direct and indirect spatial impacts while decomposing the structural rebound ratio. The third table focuses on policy, detailing the fuzzy-RDD first stage and local treatment effects under varying bandwidths. These matrices remain populated with placeholders until field data collection formally concludes.

6. Threats to Validity and Limitations

A few empirical constraints require acknowledgment. Foremost is the cross-sectional nature of the primary survey. Consequently, farmer-level treatment effects hinge entirely on selection-on-observables and structural correction assumptions rather than actual within-subject variation. While we do substantiate causal claims regarding aquifer dynamics, this longitudinal dimension relies on secondary, mandal-level administrative data. This means the household and hydrological data layers, though structurally linked, operate at mismatched spatial resolutions.

Second, isolating the specific impact of drip irrigation on groundwater levels is notoriously difficult. Fluctuating crop prices, erratic rainfall, canal seepage, and overlapping extraction regimes all simultaneously perturb the water table. The spatial and DiD estimators drastically reduce this omitted variable bias, but they cannot eliminate it entirely; hence, any hydrological conclusions warrant interpretive caution. Third, the micro-data rely on self-reported inputs, yields, and pumping hours, inherently inviting recall bias. Volumetric extraction is reconstructed using assumed pump discharge rates. We intend to validate these figures against a metered subsample and will report how sensitive the main results are to the chosen discharge calibration. Finally, the RDD estimates represent a Local Average Treatment Effect (LATE) localized strictly at the subsidy eligibility threshold. Extrapolating these behavioral responses to the broader farming population is risky. This limitation justifies our choice to triangulate the LATE with ESR and matching models engineered to recover broader population-average estimands.

7. Conclusion and Policy Relevance

Telangana channels massive public capital into drip irrigation subsidies, yet regional aquifers continue their rapid decline. Conventional evaluation tools fail to resolve whether this divergence is causal or merely coincidental. Our framework untangles the core mechanisms driving this policy dilemma through three distinct identification strategies. We model subsidy access via a sequential-choice framework featuring endogeneity corrections. Viability is assessed through an endogenous switching regression, augmented by doubly robust estimators and marginal treatment effects. Sustainability is evaluated by coupling a heterogeneity-robust DiD with a spatial common-pool model, allowing us to isolate a structural rebound ratio. A fuzzy RDD ties these disparate elements back to the primary policy lever: the subsidy threshold itself.

Ultimately, the resulting empirical output is designed for structural reform rather than superficial advocacy. Analyzing the access stage reveals whether the bureaucracy systematically filters out marginal farmers, pinpointing exactly where the subsidy apparatus fails. Comparing the ATU to the ATT dictates whether broad statewide expansion or stringent targeting maximizes the marginal return on public funds. Most critically, evaluating the rebound ratio establishes whether underwriting irrigation efficiency actually conserves groundwater. If the ratio exceeds unity, the subsidy essentially accelerates aquifer depletion under the guise of conservation, illustrating the Jevons paradox in the absence of hard extraction caps. These empirical parameters must dictate agricultural spending in a water-stressed state. By rendering them estimable and rigorously falsifiable, this framework provides a transparent blueprint for modern groundwater policy.

References

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

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. Econometrica, 82(6), 2295–2326.

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.

Central Ground Water Board (CGWB). (2024). Dynamic Ground Water Resources of India. Ministry of Jal Shakti, Government of India.

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review, 110(9), 2964–2996.

Di Falco, S., Veronesi, M., & Yesuf, M. (2011). Does adaptation to climate change provide food security? A micro-perspective from Ethiopia. American Journal of Agricultural Economics, 93(3), 829–846.

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

Government of Telangana. (2024). Telangana State Micro Irrigation Project (TSMIP) Guidelines. Department of Horticulture.

Grafton, R. Q., Williams, J., Perry, C. J., Molle, F., Ringler, C., Steduto, P., Udall, B., Wheeler, S. A., Wang, Y., Garrick, D., & Allen, R. G. (2018). The paradox of irrigation efficiency. Science, 361(6404), 748–750.

Heckman, J. J. (1979). Sample selection bias as a specification error. Econometrica, 47(1), 153–161.

Heckman, J. J., & Vytlacil, E. (2005). Structural equations, treatment effects, and econometric policy evaluation. Econometrica, 73(3), 669–738.

Kumar, D. S., & Palanisami, K. (2010). Impact of drip irrigation on farming system: Evidence from southern India. Agricultural Economics Research Review, 23(2), 265–272.

Kumar, D. S., & Palanisami, K. (2011). Can drip irrigation technology be socially beneficial? Evidence from southern India. Water Policy, 13(4), 571–587.

LeSage, J., & Pace, R. K. (2009). Introduction to Spatial Econometrics. Chapman & Hall/CRC.

Lewbel, A. (2012). Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. Journal of Business & Economic Statistics, 30(1), 67–80.

Lokshin, M., & Sajaia, Z. (2004). Maximum likelihood estimation of endogenous switching regression models. The Stata Journal, 4(3), 282–289.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. Journal of Econometrics, 142(2), 698–714.

Mohan, G., et al. (2024). Assessing determinants, challenges, and perceptions of adopting water-saving technologies among agricultural households in semi-arid states of India. Current Research in Environmental Sustainability, 7, 100255.

Namara, R. E., Upadhyay, B., & Nagar, R. K. (2005). Adoption and impacts of micro irrigation technologies (Research Report 93). International Water Management Institute.

Narayanamoorthy, A. (2004). Drip irrigation in India: Can it solve water scarcity? Water Policy, 6(2), 117–130.

Palanisami, K., Mohan, K., Kakumanu, K. R., & Raman, S. (2011). Spread and economics of micro-irrigation in India: Evidence from nine states. Economic and Political Weekly, 46(26–27), 81–86.

Pfeiffer, L., & Lin, C.-Y. C. (2014). Does efficient irrigation technology lead to reduced groundwater extraction? Empirical evidence. Journal of Environmental Economics and Management, 67(2), 189–208.

Rivers, D., & Vuong, Q. H. (1988). Limited information estimators and exogeneity tests for simultaneous probit models. Journal of Econometrics, 39(3), 347–366.

Romano, J. P., & Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. Econometrica, 73(4), 1237–1282.

Rosenbaum, P. R. (2002). Observational Studies (2nd ed.). Springer.

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

Telangana Ground Water Department. (2025). Groundwater Scenario of Telangana. Government of Telangana.

Wooldridge, J. M. (2010). Econometric Analysis of Cross-Section and Panel Data (2nd ed.). MIT Press.