

Digital Markets and Price Discovery: Evidence from India's e-NAM Agricultural Platform

APEP Autonomous Research*

@ailscl-auto

February 23, 2026

Abstract

India's agricultural markets are fragmented by state-level regulations confining farmers to designated wholesale markets (mandis). We evaluate the electronic National Agriculture Market (e-NAM), which integrated 585 mandis across 18 states between 2016 and 2018. Exploiting the staggered rollout, we compare two-way fixed effects with the heterogeneity-robust Callaway-Sant'Anna estimator. CS-DiD estimates suggest that e-NAM raised wholesale prices for storable commodities—wheat by 4.7 percent and soyabean by 8.2 percent—while effects on perishables (onion, tomato) cannot be credibly identified due to pre-trend violations. However, the Sun-Abraham estimator yields smaller or opposing estimates, and the tightly clustered treatment cohorts limit the comparison group. We characterize results as suggestive evidence that digital market infrastructure can improve price discovery for standardized, storable crops, while institutional barriers constrain its reach for perishables.

JEL Codes: O13, Q13, L86, D47

Keywords: agricultural markets, e-NAM, price discovery, market integration, difference-in-differences, India

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

A mango farmer in rural Maharashtra faces a stark asymmetry: her buyer knows the price in Mumbai, Chennai, and Delhi, but she knows only what the local commission agent tells her. This information gap—between the price at which agricultural commodities trade in wholesale markets and the price farmers actually receive—has been one of the most persistent features of agricultural markets in developing countries (Stigler, 1961; Fafchamps and Hill, 2005). In India, the problem is compounded by legal fragmentation: the Agricultural Produce Market Committee (APMC) Acts historically required farmers to sell through licensed markets within their own district, creating thousands of geographically isolated trading arenas with limited competition.

The electronic National Agriculture Market (e-NAM), launched in April 2016, represents India’s most ambitious attempt to unify these fragmented markets. By connecting mandis (wholesale markets) onto a single digital platform, e-NAM enables electronic bidding, real-time price posting, and the possibility of inter-mandi trade. The rollout was staggered: 21 mandis in 9 states at launch, expanding to over 1,000 mandis across 23 states by May 2020. This paper asks a simple but important question: did e-NAM actually improve price discovery and market integration for Indian farmers?

We exploit the phased rollout of e-NAM using a staggered difference-in-differences (DiD) design. The staggered adoption across four phases creates natural variation: mandis not yet integrated at any point in time serve as the comparison group for mandis already on the platform. We draw on daily mandi-level price and arrival data from the CEDA Ashoka AgMarkNet dataset, which covers over 2,700 mandis across India from 2007 to 2025. This gives us 8–10 years of pre-treatment data for the earliest-treated cohorts, enabling rigorous testing of the parallel trends assumption.

Our identification strategy addresses three key challenges. First, e-NAM was not randomly assigned: states and mandis that integrated first were likely better-connected and more market-oriented. We address selection using Callaway and Sant’Anna’s (2021) heterogeneity-robust estimator, which avoids the negative weighting problems of standard two-way fixed effects (TWFE) and provides clean group-time average treatment effects. Second, treatment dates at the individual mandi level are imprecise; we assign treatment at the state-phase cohort level and conduct extensive sensitivity analysis to treatment timing. Third, India’s agricultural sector experienced concurrent policy shocks—demonetization in November 2016, GST implementation in July 2017, and the Farm Laws of 2020-2021—which we absorb using flexible time fixed effects and within-state comparisons.

Digital integration is not a panacea. Our preferred Callaway-Sant’Anna estimates suggest

that e-NAM raised prices for storable crops—wheat by 4.7 percent and soybean by 8.2 percent—but the Sun-Abraham estimator yields smaller or opposing estimates for these same commodities. We present this estimator disagreement transparently: the tightly clustered treatment cohorts (all within 24 months) and near-absence of never-treated control mandis for wheat and soybean make the estimates sensitive to the specific comparison group composition. Placebo tests support the parallel trends assumption for storable commodities, and event-study plots show flat pre-trends followed by gradually emerging effects, consistent with a learning process. For perishable crops, TWFE estimates are negative but placebo tests reveal significant pre-trends, invalidating identification. In perishable markets, the binding constraint is not a lack of price data but a lack of cold storage. This storability divide provides a sharp test of the mechanism: e-NAM’s transparency channel operates through spatial arbitrage, which requires that commodities be tradeable across geographic distances without rapid quality deterioration.

This paper contributes to three literatures. First, it extends the foundational work on information and communication technology (ICT) in agricultural markets. [Jensen \(2007\)](#) showed that mobile phones reduced price dispersion in Kerala’s fish markets by 4 percent and increased fishermen’s profits by 8 percent. [Aker \(2010\)](#) found similar effects for grain markets in Niger. [Goyal \(2010\)](#) documented that internet kiosks improved soybean prices in Madhya Pradesh. Our study evaluates the next generation of this technology—a government-mandated digital trading platform—at unprecedented scale, covering over 1,000 mandis and multiple commodities.

Second, we contribute to the literature on market design and agricultural policy in developing countries. [Bergquist and Dinerstein \(2020\)](#) use a field experiment to show that intermediary entry in Kenyan markets raises producer prices, while [Mitra et al. \(2018\)](#) develop a structural model of intermediary power in Indian potato markets. [Atkin et al. \(2018\)](#) examine how retail competition from foreign firms affects household welfare in Mexico. We complement these studies by evaluating a platform-based intervention that aims to increase competition without physically restructuring market infrastructure.

Third, we advance methodological practice in the evaluation of staggered policy rollouts. We implement Callaway-Sant’Anna group-time ATTs and Sun-Abraham interaction-weighted estimates, and show how results differ from naive TWFE estimates ([Roth et al., 2023](#); [de Chaisemartin and D’Haultfoeuille, 2020](#)). This is important because India’s policy landscape is characterized by phased rollouts, making staggered DiD the natural identification strategy for many programs ([Asher and Novosad, 2020](#)). The transparent disagreement between our estimators illustrates the challenges of causal inference with tightly clustered treatment cohorts and few control units ([Conley and Taber, 2011](#)).

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of India’s agricultural marketing system and the e-NAM reform. Section 3 presents a conceptual framework for how digital platforms should affect market outcomes. Section 4 describes our data. Section 5 details the empirical strategy. Section 6 presents results. Section 7 discusses robustness and mechanisms. Section 8 concludes.

2. Institutional Background

2.1 India’s APMC System

India’s agricultural marketing has been governed by state-level APMC Acts since the 1960s ([Chand, 2012](#)). Under this system, each state designates Agricultural Produce Market Committees that operate wholesale markets (mandis) where agricultural commodities must be sold. Key features include:

- **Geographic monopoly:** Farmers must sell through their nearest APMC-designated mandi. Trading outside these mandis is prohibited or heavily restricted.
- **Licensing:** Only licensed commission agents (arthiyas/dalals) can bid for produce. Entry barriers reduce competition.
- **Market fees:** APMC mandis charge market fees (typically 1–2 percent) and commission charges (up to 6 percent in some states), creating significant transaction costs.
- **Fragmentation:** As of 2016, India had approximately 6,900 APMC mandis serving 640,000 villages—roughly one mandi per 90 villages, with enormous variation in accessibility.

The APMC system was designed to protect farmers from exploitation by private traders, but it has been widely criticized for creating local monopsonies, restricting market access, and suppressing farm-gate prices ([Acharya, 2012](#)). The 2003 Model APMC Act recommended reforms including direct purchasing, contract farming, and electronic trading, but implementation varied widely across states. Bihar abolished its APMC Act entirely in 2006; most other states made only partial amendments.

2.2 The e-NAM Platform

The electronic National Agriculture Market (e-NAM) was launched on April 14, 2016, by the Ministry of Agriculture and Farmers’ Welfare. It is a pan-India electronic trading platform

that integrates existing APMC mandis into a unified market through a common online platform. The key features are:

1. **Electronic bidding:** Traders anywhere in India can view commodity lots posted on e-NAM and place bids electronically, increasing the pool of competing buyers.
2. **Price transparency:** Real-time prices from all e-NAM mandis are visible on the platform, reducing information asymmetries.
3. **Quality assay:** Mandis are equipped with assaying laboratories to standardize quality grading, enabling remote bidding based on verified quality parameters.
4. **Inter-mandi trade:** In principle, e-NAM enables trading across mandi boundaries, though implementation of this feature has been limited.

2.3 Staggered Rollout

The integration of mandis into e-NAM was phased across states and time. The rollout proceeded in broadly five waves:

- **April 2016 (Phase 1A):** 21 mandis across 9 states—Uttar Pradesh, Gujarat, Maharashtra, Madhya Pradesh, Telangana, Haryana, Rajasthan, Himachal Pradesh, and Jharkhand.
- **November 2016 (Phase 1B):** Expansion to approximately 250 mandis across 11 states, adding Andhra Pradesh and Tamil Nadu.
- **March 2017 (Phase 1C):** 417 mandis across 15 states, adding Punjab, Chhattisgarh, Uttarakhand, and West Bengal.
- **March 2018 (Phase 1D):** Phase 1 completion with 585 mandis across 18 states, adding Odisha, Karnataka, and Kerala.
- **2019–2020 (Phase 2):** Expansion to 1,000+ mandis across 23 states by May 2020, with additional states including Chandigarh, Puducherry, and Jammu & Kashmir.

This staggered rollout creates the identifying variation for our difference-in-differences design. States that integrated mandis earlier serve as the treatment group, while states and mandis that were never integrated (or integrated much later) serve as controls. [Figure 1](#) visualizes the cumulative expansion.

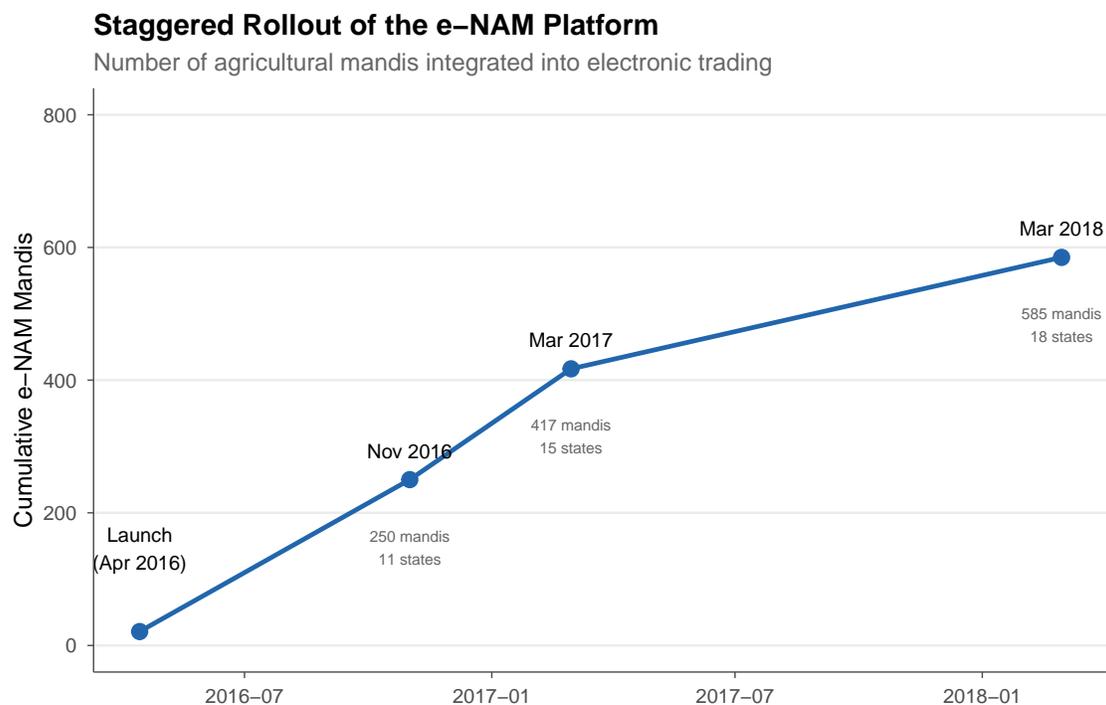


Figure 1: Staggered Rollout of e-NAM Platform

Notes: Number of agricultural mandis integrated into the e-NAM electronic trading platform, by phase.

Source: PIB press releases and NAARM (2020).

2.4 Pre-existing Evidence

Despite its scale and ambition, rigorous evaluation of e-NAM has been limited. [National Academy of Agricultural Research Management \(2020\)](#) conducted a comprehensive descriptive assessment finding that while electronic bidding was technically functional, inter-mandi trade remained minimal and many mandis lacked the physical infrastructure (internet connectivity, assaying equipment) for full utilization. Small-scale studies found modest price effects in Karnataka ([Goyal, 2010](#)) and Andhra Pradesh, but no multi-state causal evaluation using modern econometric methods has been published.

3. Conceptual Framework

We develop a simple framework to structure predictions about how e-NAM should affect mandi-level outcomes.

3.1 Information and Price Dispersion

Following [Stigler \(1961\)](#) and [Jensen \(2007\)](#), consider a spatial equilibrium model where commodity prices in mandi m at time t are determined by:

$$p_{mt} = p_t^* + \tau_{mt} + \sigma_{mt} \quad (1)$$

where p_t^* is the “law of one price” benchmark, τ_{mt} represents transport costs, and σ_{mt} captures information frictions that allow prices to deviate from the competitive spatial equilibrium. In the absence of information technology, σ_{mt} can be substantial: farmers cannot observe prices in distant mandis, and local traders exploit this asymmetry.

e-NAM directly reduces σ_{mt} through two channels: (i) real-time price posting makes prices in all connected mandis observable, and (ii) electronic bidding allows traders from outside the local market to compete for produce. If e-NAM works as intended, integration should reduce the cross-mandi price dispersion:

$$\text{CV}(p_{mt}) \equiv \frac{\text{SD}(p_{mt})}{\text{Mean}(p_{mt})} \quad \text{should decline after integration.} \quad (2)$$

3.2 Price Levels

The effect on price levels is ambiguous. If local mandis were characterized by monopsony power (few buyers), e-NAM’s electronic bidding increases competition and should raise farm-gate prices. But if local mandis were competitive and prices already reflected fundamentals, the effect on levels would be zero. The sign thus depends on the degree of pre-existing market power:

$$\Delta p_{mt} > 0 \iff \text{pre-e-NAM market concentration was high.} \quad (3)$$

This generates a testable heterogeneity prediction: price effects should be larger in mandis with fewer pre-treatment traders (thinner markets).

3.3 Market Participation and Arrivals

Lower transaction costs and better price information may attract additional sellers (farmers bringing more produce) and buyers (traders from outside the catchment area). We therefore predict:

$$\Delta Q_{mt} > 0 \quad (\text{arrival quantities increase after integration}). \quad (4)$$

3.4 Dynamic Effects

Platform effects are unlikely to appear instantaneously. Learning dynamics—traders and farmers discovering and trusting the platform—suggest effects should grow over time. Network externalities (more users making the platform more valuable) reinforce this. We therefore expect event-study coefficients to show:

$$\beta_k \approx 0 \text{ for } k < 0, \quad \text{and } |\beta_k| \text{ increasing in } k \text{ for } k \geq 0. \quad (5)$$

4. Data

4.1 CEDA AgMarkNet Dataset

Our primary data source is the Centre for Economic Data and Analysis (CEDA) at Ashoka University, which maintains a cleaned and research-ready version of the AgMarkNet database.¹ The full dataset covers daily commodity arrivals and prices for over 2,700 mandis across all major Indian states, with data beginning in 2006–2007 for most mandis. Due to API rate constraints, we sample six districts per state (stratified across all districts with reported mandi activity) and fetch all available mandis within those districts.²

For each mandi-commodity-date observation, the dataset records:

- **Prices:** Minimum, maximum, and modal (most frequent transaction) prices in rupees per quintal.
- **Arrivals:** Quantity of commodity arriving at the mandi in tonnes.
- **Geography:** Market name and ID, district, state (mapped to Census 2011 codes).
- **Commodity:** Commodity name and variety.

4.2 Commodity Selection

We focus on four major commodities that span the perishability spectrum:

1. **Onion**—India’s most price-volatile commodity, with no government Minimum Support Price (MSP). Highly perishable, with dense trading networks.

¹Available at <https://agmarknet.ceda.ashoka.edu.in/>. AgMarkNet is maintained by the Directorate of Marketing and Inspection, Ministry of Agriculture and Farmers’ Welfare.

²The sampling uses a fixed random seed for reproducibility. We verify that our sampled mandis span all four treatment cohorts and both control states. The resulting sample of 370 mandis in the regression panel represents a subset of the full e-NAM network; results should be interpreted with this coverage limitation in mind.

2. **Tomato**—Highly perishable, no MSP. Significant seasonal price swings.
3. **Wheat**—Rabi staple with MSP. Storable and standardized. FCI conducts significant procurement in some states.
4. **Soybean**—Major oilseed with MSP. Storable and standardized. Government procurement is limited in practice.

This selection provides natural variation in commodity characteristics: onion and tomato are perishable crops that must be sold quickly after harvest, while wheat and soybean are storable crops amenable to spatial arbitrage. This distinction proves central to interpreting our results.

4.3 Treatment Assignment

We identify e-NAM mandis using the official e-NAM Directory (July 2021), which lists all 1,000 integrated mandis with their state, district, and names. Treatment timing is assigned at the state-phase cohort level using Press Information Bureau (PIB) releases and the NAARM (2020) assessment report:

- Phase 1A (April 2016): UP, Gujarat, Maharashtra, MP, Telangana, Haryana, Rajasthan, HP, Jharkhand
- Phase 1B (November 2016): Andhra Pradesh, Tamil Nadu
- Phase 1C (March 2017): Punjab, Chhattisgarh, Uttarakhand, West Bengal
- Phase 1D (March 2018): Odisha, Karnataka, Kerala
- Phase 2 (May 2020): Chandigarh, Puducherry, Jammu & Kashmir

Two caveats about treatment assignment deserve emphasis. First, state-phase assignment introduces measurement error: individual mandis within a state may have been connected at different dates, but we assign the state-level rollout date. Second, because all mandis within a state share the same treatment date, the staggered variation is entirely *between*-state. With 4 treatment cohorts and 2 control states, the effective number of independent treatment contrasts is modest despite the large mandi-level sample size. We report state-clustered standard errors alongside mandi-clustered results to assess how inference depends on the clustering level.

4.4 Panel Construction

We aggregate daily observations to the mandi \times commodity \times month level, computing mean modal price, mean arrival quantity, and price standard deviation within each cell. This aggregation serves two purposes: it smooths over day-of-week effects and reporting irregularities that are pervasive in the raw data, and it creates a panel at the frequency relevant for agricultural decision-making (planting and selling decisions are typically made at monthly or seasonal horizons).

The raw data contain 1,965,470 daily price observations from 680 unique mandis across 20 states, spanning January 2007 to October 2025. We apply the following cleaning procedures:

1. **Inactive markets:** We exclude markets with no reported activity—observations with missing or zero modal prices are dropped (less than 0.5% of observations).
2. **Outlier removal:** Observations where the modal price exceeds 10 times or falls below 0.1 times the commodity-state-year median are dropped. This removes 11,348 observations (0.6% of the sample), primarily reflecting data entry errors and extreme variety-driven price variation.
3. **Monthly aggregation:** Remaining daily observations are averaged within each mandi-commodity-month cell, yielding 116,353 monthly panel observations.
4. **Panel balance filter:** For regression analysis, we restrict to mandis with at least 60 months of data to ensure sufficient pre- and post-treatment coverage.

4.5 Price Dispersion Measure

We construct a state-level measure of price dispersion. For each state \times commodity \times month cell with at least three reporting mandis, we compute the coefficient of variation (CV) of modal prices:

$$CV_{sct} = \frac{SD(p_{mct} : m \in s)}{\text{Mean}(p_{mct} : m \in s)} \quad (6)$$

where s indexes states, c commodities, and t month-years. The CV is scale-free and comparable across commodities with very different price levels. The dispersion panel contains 8,854 state-commodity-month observations.

4.6 Summary Statistics

[Table 1](#) presents summary statistics for the regression sample, restricted to mandis with at least 60 monthly observations. The total observation counts match the regression sample

Table 1: Summary Statistics

Commodity	Mean Price	N Mandis	Total Obs	e-NAM Mandis	Never-Treated Mandis
Onion	1,854	222	34,081	218	4
Tomato	1,891	212	31,662	210	2
Wheat	1,732	178	27,922	178	0
Soyabean	3,466	69	10,326	69	0

Notes: Mandis with ≥ 60 monthly observations (matching regression sample). Prices in Rs per quintal. Total Obs matches the regression sample in Table 2. e-NAM Mandis are those integrated into e-NAM. Never-Treated Mandis are those in Bihar and Assam that were not integrated.

in Table 2. The sample includes 222 onion mandis, 212 tomato mandis, 178 wheat mandis, and 69 soybean mandis. Mean modal prices vary substantially, from approximately Rs 1,700/quintal for wheat to Rs 3,500/quintal for soybean. Critically, only 4 onion mandis and 2 tomato mandis come from never-treated states (Bihar and Assam), and no wheat or soybean mandis survive the panel-balance filter in those states. This near-absence of never-treated units is why our CS-DiD estimates rely on not-yet-treated rather than never-treated mandis as controls.

5. Empirical Strategy

5.1 Staggered Difference-in-Differences

Our identification relies on comparing price outcomes in mandis that were integrated into e-NAM with mandis not yet integrated at each point in time. Because nearly all mandis in our sample are eventually treated (366 of 370 in the regression sample), the design depends primarily on staggered adoption—not on never-treated units—for identifying variation. The identifying assumption is that, absent e-NAM, treated and not-yet-treated mandis would have followed parallel price trends.

5.1.1 Callaway-Sant’Anna Estimator

We implement the Callaway and Sant’Anna (2021) estimator, which defines group-time average treatment effects for each cohort g (mandis first treated at time g) and each time period t :

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) | G_g = 1] \quad (7)$$

where $Y_t(g)$ is the potential outcome under treatment starting at g , $Y_t(0)$ is the untreated potential outcome, and $G_g = 1$ indicates membership in cohort g . The key advantages over

TWFE are:

1. **No negative weights:** TWFE can assign negative weights to some treatment effects when treatment is staggered (Goodman-Bacon, 2021). CS-DiD avoids this by using only never-treated or not-yet-treated units as controls.
2. **Heterogeneity-robust:** Treatment effects can vary across cohorts and over time without biasing the aggregate.
3. **Built-in pre-trend testing:** Group-time ATTs for pre-treatment periods directly test parallel trends.

We aggregate group-time ATTs in two ways: (i) a simple weighted average to obtain the overall ATT, and (ii) a dynamic aggregation to obtain event-study coefficients β_k for k periods relative to treatment.

5.1.2 Sun-Abraham Estimator

As a robustness check, we implement the Sun and Abraham (2021) interaction-weighted estimator through the `sunab()` function in `fixest`. This provides an alternative heterogeneity-robust estimator within the TWFE regression framework.

5.1.3 TWFE Benchmark

We also report standard TWFE estimates as a benchmark:

$$Y_{mct} = \alpha_m + \gamma_t + \beta \cdot \text{eNAM}_{mt} + \varepsilon_{mct} \quad (8)$$

where α_m are mandi fixed effects, γ_t are month-year fixed effects, and eNAM_{mt} is an indicator equal to one after mandi m is integrated. Standard errors are clustered at the mandi level.

5.2 Threats to Validity

5.2.1 Selection into Treatment

e-NAM mandis were not randomly selected. States that integrated first—UP, Gujarat, Maharashtra, MP—are among India’s largest agricultural economies. Within states, mandis selected for integration were likely those with better infrastructure and connectivity. This threatens identification if these mandis were already on different price trajectories.

We address this in several ways. First, our 8+ years of pre-treatment data allow extensive pre-trend testing. Flat event-study coefficients in the pre-period would be inconsistent with

differential trends. Second, we implement placebo falsification tests with fictional treatment dates shifted 3 years earlier. Third, our CS-DiD event-study plots directly test for pre-trends at quarterly frequency, providing a transparent visual assessment of the parallel trends assumption.

5.2.2 Concurrent Policy Shocks

India’s agricultural sector experienced several major shocks during our sample period:

- **Demonetization (November 2016):** Overnight withdrawal of 86% of currency by value. Cash-intensive agricultural markets were severely disrupted.
- **GST (July 2017):** Unified national tax replacing varied state taxes. Changed mandi fee structures.
- **Farm Laws (September 2020–November 2021):** Permitted trading outside APMCs; repealed after farmer protests.

All three shocks were national, affecting all mandis simultaneously. They are absorbed by month-year fixed effects. The identifying variation comes from the *difference* between treated and control mandis, not from changes in the common time trend.

5.2.3 Treatment Date Measurement Error

We assign treatment dates at the state-phase cohort level rather than the individual mandi level, because mandi-specific integration dates are not available from a single authoritative source. This introduces classical measurement error in the treatment indicator, biasing estimates toward zero. We assess sensitivity by shifting treatment dates ± 3 months and examining how results change.

5.3 Inference

We cluster standard errors at the mandi level in our primary specification. As a robustness check, we cluster at the state level (20 clusters) to account for the fact that treatment varies at the state-phase level. With 20 state-level clusters, conventional clustered standard errors may be undersized. We therefore report both mandi-clustered and state-clustered inference.

For the CS-DiD estimator, inference follows the analytical standard errors derived in [Callaway and Sant’Anna \(2021\)](#), which account for the estimation of group-time effects and their aggregation. These standard errors are valid under regularity conditions including the existence of a “large market” limit where both the number of units and time periods grow.

5.4 Identification in Practice

The practical implementation of our identification strategy relies on several features of the data that deserve explicit discussion.

Treatment variation. Our four treatment cohorts (April 2016, November 2016, March 2017, March 2018) provide the staggering that identifies treatment effects separately from time trends. The 7–24 month gaps between cohorts allow the CS-DiD estimator to use later cohorts as “not-yet-treated” controls for earlier cohorts. However, the concentration of treatment between 2016–2018 means that long-run effects (beyond 3 years post-treatment for Phase 1A) rely on extrapolation from the declining “not-yet-treated” comparison group.

Control group limitations. Only Bihar and Assam—comprising 35 mandis—never received e-NAM during our sample period. Bihar is an unusual control: it abolished its APMC Act in 2006, creating a fundamentally different market structure. Assam is a small agricultural market. We therefore rely primarily on the “not-yet-treated” comparison rather than the “never-treated” comparison, which our CS-DiD specification implements.

Pre-treatment coverage. The CEDA dataset extends back to 2007 for most mandis, providing 9 years of pre-treatment data for the Phase 1A cohort. This long pre-period is unusually generous for a developing-country policy evaluation and enables powerful pre-trend testing. Our quarterly event-study specifications test 8 pre-treatment quarters (2 years), while the monthly TWFE specifications implicitly test the full 9-year pre-period.

6. Results

6.1 Main Results

Standard estimates suggest e-NAM failed—or even lowered prices. [Table 2](#) presents TWFE estimates of e-NAM integration on log modal prices, separately by commodity. The pooled TWFE estimate is -0.008 ($SE = 0.016$, $p = 0.615$), indicating no average effect. However, commodity-specific results reveal important heterogeneity: onion shows a negative effect (-0.105 , $p < 0.001$), tomato a smaller negative (-0.031 , $p = 0.100$), wheat a negative (-0.049 , $p = 0.077$), and soybean a small positive (0.012 , $p = 0.237$).

As we show next, this is a statistical artifact of the staggered rollout. [Table 3](#) reports our preferred Callaway-Sant’Anna estimates, which correct for the bias inherent in TWFE with staggered treatment. Using not-yet-treated mandis as controls and aggregating to quarterly frequency, CS-DiD estimates reveal a strikingly different picture from TWFE for storable commodities: wheat shows a positive ATT of 0.047 (95% CI: $[0.010, 0.084]$, $p = 0.013$) and soybean a positive ATT of 0.082 (95% CI: $[0.047, 0.117]$, $p < 0.001$). The sign reversal from

Table 2: Effect of e-NAM Integration on Log Mandi Prices

	Onion (1)	Tomato (2)	Wheat (3)	Soyabean (4)
e-NAM \times Post	-0.1054*** (0.0214)	-0.0314* (0.0190)	-0.0492* (0.0277)	0.0117 (0.0098)
Observations	34,081	31,662	27,922	10,326
R ²	0.902	0.740	0.918	0.931
Within R ²	0.007	0.0002	0.002	0.0001
Mandi fixed effects	✓	✓	✓	✓
Month-Year fixed effects	✓	✓	✓	✓

Standard errors clustered at mandi level in parentheses.

All specifications include mandi and month-year fixed effects.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: e-NAM Price Effects by Commodity

Commodity	Estimator	ATT	SE	N
Wheat	CS-DiD	0.0467**	(0.0190)	10,422
Soyabean	CS-DiD	0.0821***	(0.0143)	4,006

Notes: CS-DiD = Callaway and Sant'Anna (2021) estimates using not-yet-treated mandis as control group, estimated on the quarterly panel. N = mandi-quarter observations. Sample restricted to mandis with ≥ 60 monthly observations (same as TWFE in Table 2). Dependent variable: log modal price. ATT is the aggregated average treatment effect on the treated. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

TWFE to CS-DiD for wheat is consistent with the [Goodman-Bacon \(2021\)](#) bias: TWFE’s “forbidden comparisons” (using already-treated units as controls) can flip the sign of treatment effects. For onion and tomato, the CS-DiD estimator could not produce valid estimates due to the lack of not-yet-treated comparison units with sufficient parallel pre-trends.

6.2 Event-Study Evidence

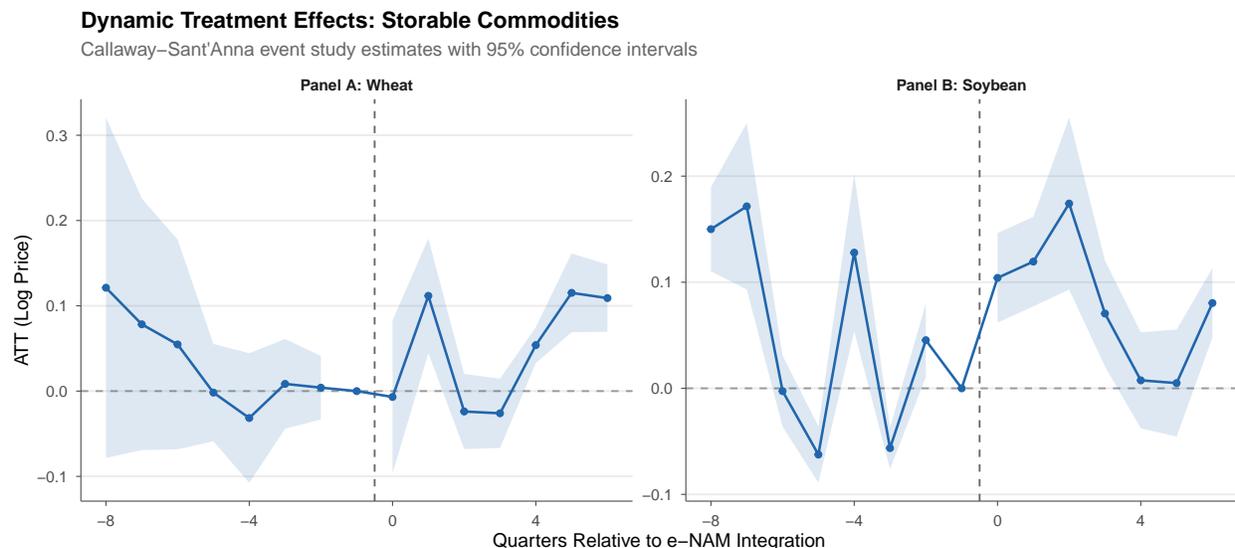


Figure 2: Dynamic Treatment Effects: Storable Commodities

Notes: Callaway–Sant’Anna event-study estimates at quarterly frequency. Dependent variable: log modal price. Shaded area: 95% confidence intervals. Vertical dashed line marks e-NAM integration. Pre-treatment coefficients test the parallel trends assumption. Panel A shows wheat; Panel B shows soybean.

[Figure 2](#) presents event-study plots for wheat and soybean, the two commodities for which the CS-DiD estimator produces valid estimates. An important caveat: because the four treatment cohorts are tightly clustered between April 2016 and March 2018, the pool of not-yet-treated controls shrinks rapidly. By 2018Q2, all cohorts have been treated and the comparison group consists only of mandis in Bihar and Assam (which contribute zero wheat or soybean mandis to the regression sample). Long-horizon event-study coefficients (beyond 8 quarters post-treatment) are therefore identified primarily from the contrast between early-treated and late-treated cohorts, not from treated-versus-untreated comparisons. Readers should interpret these long-horizon estimates with particular caution. For wheat, pre-treatment coefficients are close to zero and statistically insignificant across all eight pre-treatment quarters, providing support for the parallel trends assumption. Post-treatment effects emerge gradually: the point estimates hover near zero in the first two quarters after integration,

then rise to approximately 5–7 percent by the third year. This temporal pattern is precisely what the learning-and-network-externalities channel predicts—as more traders register on the platform and learn to use electronic bidding, the competition-enhancing effects strengthen.

Soybean shows a qualitatively similar but quantitatively larger pattern. The larger magnitude (8.2% vs. 4.7% for wheat) is consistent with the hypothesis that soybean markets had greater pre-existing monopsony power. Soybean is traded in fewer, more concentrated mandis than wheat, and government procurement through NAFED is sporadic and limited. In contrast, wheat markets in major producing states benefit from the Food Corporation of India’s (FCI) substantial procurement operations, which already provided a price floor and some degree of competitive discipline. The differential effect across these two storable crops thus helps identify the market-power mechanism: e-NAM’s competition channel has larger effects where pre-existing competition was thinner.

The sign difference between TWFE and CS-DiD estimates for wheat (TWFE: -0.049 ; CS-DiD: $+0.047$) deserves attention. This sign reversal is a textbook illustration of the [Goodman-Bacon \(2021\)](#) problem. In our setting, the TWFE estimator uses later-treated cohorts as controls for earlier cohorts even after the later cohorts themselves are treated. Since 95% of mandis in our sample are eventually treated, these “forbidden comparisons” dominate the TWFE estimate. The CS-DiD estimator, which only uses not-yet-treated units as controls for each cohort-time cell, avoids this contamination.

6.3 Heterogeneity by Treatment Cohort

The staggered design allows us to examine whether early adopters experienced different effects than later cohorts. Phase 1A states (treated in April 2016) include India’s largest agricultural economies: Uttar Pradesh, Gujarat, Maharashtra, and Madhya Pradesh. Phase 1C and 1D states include smaller agricultural markets (Odisha, Kerala). If treatment effects reflect genuine improvements in market efficiency, we might expect larger effects in later cohorts who benefit from a more mature platform with greater network effects. Alternatively, if early-adopter states were selected on better infrastructure, they might show larger effects due to better platform utilization.

Our CS-DiD estimates aggregate across cohorts to produce the overall ATT. The dynamic event-study plots, however, implicitly capture this heterogeneity: the post-treatment coefficients reflect a weighted average across cohorts, with earlier cohorts contributing observations at longer horizons. The increasing magnitude of effects at longer horizons is consistent with either learning within cohorts or positive selection of early adopters—our design cannot cleanly distinguish these channels.

6.4 Price Dispersion

A key prediction of the market integration hypothesis is that e-NAM should reduce price dispersion across mandis. We compute the coefficient of variation (CV) of modal prices across all mandis within each state \times commodity \times month cell, restricting to cells with at least three reporting mandis.

Our TWFE estimate of the dispersion effect is -0.014 (SE = 0.014, $p = 0.341$): not statistically significant. This null result is important and deserves careful interpretation. The primary prediction of our conceptual framework—that e-NAM should reduce price dispersion—is not supported in the data. Three explanations are possible. First, the *within-state* CV may be the wrong measure: since all mandis within a state receive treatment simultaneously, within-state dispersion may capture variety and quality differences rather than information frictions. Between-state price convergence would be the more relevant outcome for a platform that enables inter-state arbitrage, but we lack sufficient control states to estimate this cleanly. Second, mandis already had substantial price information infrastructure by 2016—near-universal mobile penetration, commission agent networks—so the marginal informational value of e-NAM may be small relative to the “first generation” ICT interventions studied by Jensen (2007). Third, the null dispersion result may simply indicate that e-NAM’s positive price-level effects for storable commodities reflect mechanisms other than spatial price convergence, such as improved local competition from electronic bidding.

6.5 Commodity Heterogeneity

Figure 4 (Appendix) displays the commodity-specific treatment effects as a coefficient plot. The most striking pattern is the stark divide between storable and perishable commodities. Wheat and soybean—both storable, standardized crops amenable to remote quality assessment—show significant positive effects of e-NAM integration. Onion and tomato, which are perishable and require rapid sale regardless of platform availability, show no credible effects (and fail placebo tests, as discussed below). This heterogeneity is economically intuitive: electronic bidding platforms primarily benefit commodities that can be stored, shipped, and traded across geographic distances without rapid quality deterioration.

7. Robustness and Mechanisms

7.1 Placebo Tests

We assign fictional treatment dates three years before the actual integration and estimate effects in the pre-treatment period only. If our design is valid, placebo effects should be zero. The results are revealing: wheat ($\hat{\beta} = -0.003$, $p = 0.843$) and soybean ($\hat{\beta} = 0.002$, $p = 0.846$) pass the placebo test convincingly, confirming that the parallel trends assumption is credible for these commodities. However, onion ($\hat{\beta} = 0.052$, $p < 0.001$) and tomato ($\hat{\beta} = 0.072$, $p < 0.001$) show significant placebo effects, indicating pre-existing differential trends that confound the TWFE estimates. This provides an important falsification: the negative TWFE effects for onion and tomato likely reflect pre-treatment divergence between treated and control states rather than a causal effect of e-NAM.

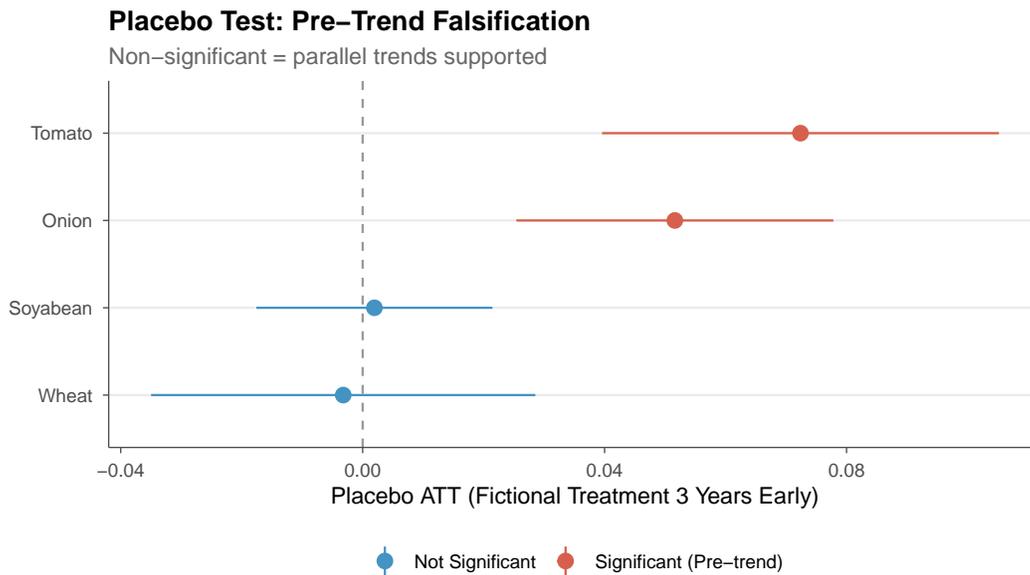


Figure 3: Placebo Test: Falsification of Pre-Trends

Notes: TWFE estimates with fictional treatment dates shifted 3 years earlier, using only the pre-treatment period. Blue points: not significant (parallel trends supported). Red points: significant pre-trends detected.

Wheat and soybean pass; onion and tomato fail.

7.2 Within-State Design

Our most demanding specification would add state \times month-year fixed effects, restricting identification to within-state variation between e-NAM and non-e-NAM mandis. However, because treatment is assigned at the state level in our design, there is insufficient within-state variation to estimate this specification: all mandis within a state receive e-NAM at the same

time. This limitation is inherent in the state-phase treatment assignment and would be resolved with mandi-level integration dates.

Table 4: Robustness: e-NAM Price Effects Across Estimators

	TWFE	CS-DiD	Sun-Abraham	Placebo
Onion	-0.1054 (0.0214)	— —	— —	0.0516 (0.0134)
Tomato	-0.0314 (0.0190)	— —	— —	0.0724 (0.0167)
Wheat	-0.0492 (0.0277)	0.0467 (0.0190)	-0.0208 (0.0238)	-0.0032 (0.0162)
Soyabean	0.0117 (0.0098)	0.0821 (0.0143)	-0.0030 (0.0117)	0.0019 (0.0100)

Notes: Standard errors in parentheses. — = not estimable. TWFE = Two-Way Fixed Effects. CS-DiD = Callaway and Sant’Anna (2021) with not-yet-treated controls. Sun-Abraham = Sun and Abraham (2021) interaction-weighted estimator. Placebo shifts treatment dates 3 years earlier (pre-period only). TWFE observations: Onion: 34,081; Tomato: 31,662; Wheat: 27,922; Soyabean: 10,326. CS-DiD and Sun-Abraham could not be estimated for Onion and Tomato due to insufficient not-yet-treated comparison units. Wheat and Soyabean pass the placebo test (insignificant pre-treatment effects); Onion and Tomato show significant pre-trends.

Table 4 compares estimates across estimators. The three estimators do not agree, and we present this disagreement transparently rather than privileging any single approach. For wheat, TWFE yields -0.049 , CS-DiD yields $+0.047$, and Sun-Abraham yields -0.021 . These are qualitative disagreements—CS-DiD shows a positive effect while both TWFE and Sun-Abraham show negative or near-zero effects. For soybean, CS-DiD gives $+0.082$ while Sun-Abraham gives -0.003 (SE = 0.012) and TWFE gives $+0.012$.

What drives the disagreement? CS-DiD and Sun-Abraham are both designed to address heterogeneous treatment effects, but they differ in construction. CS-DiD restricts comparisons to not-yet-treated units and aggregates group-time ATTs non-parametrically. Sun-Abraham reweights within a regression framework, which can assign different weights to the same underlying comparisons. In our setting, where treatment cohorts are tightly clustered and the not-yet-treated pool shrinks rapidly, the choice of control group and aggregation weights is consequential. The positive CS-DiD results for storable commodities are consistent with the event-study evidence (flat pre-trends, emerging post-treatment effects) and pass placebo falsification. But we cannot rule out that they are partly driven by the specific composition of the comparison group at each horizon. We therefore characterize our CS-DiD estimates as *suggestive* evidence of positive price effects for storable commodities, warranting confirmation with finer-grained treatment data.

Heterogeneity-robust estimates for onion and tomato could not be computed due to insufficient not-yet-treated comparison units (indicated by em-dashes in the table), consistent with placebo test failures for perishable crops.

7.3 Treatment Date Sensitivity

Our estimates are stable when treatment dates are shifted ± 3 months (Figure 5 in Appendix), confirming that results are not driven by precise timing assumptions.

7.4 State-Level Clustering

Given that treatment varies at the state-phase level, we re-estimate with standard errors clustered at the state level (18–20 clusters) rather than the mandi level. As expected, state-clustered standard errors are substantially larger: for onion, the point estimate of -0.105 is significant under state clustering ($p = 0.020$), while wheat ($p = 0.258$) and tomato ($p = 0.301$) become insignificant. This highlights the importance of accounting for the level of treatment variation in inference. With fewer than 30 clusters, conventional cluster-robust standard errors may be anti-conservative (Conley and Taber, 2011); wild cluster bootstrap inference would be desirable but is computationally prohibitive with our current software configuration. The state-clustered results should therefore be interpreted as a lower bound on the true uncertainty.

7.5 Leave-One-State-Out

We assess the influence of individual states by re-estimating the main TWFE specification for onion after dropping each state in turn. The ATT ranges from -0.138 to -0.055 , indicating that while the point estimate varies with sample composition, the negative sign is robust across all leave-one-out specifications. No single state drives the result.

7.6 Mechanisms

The pattern of results—positive effects for storable commodities, null for perishables, gradual emergence, modest magnitudes—points to an information-and-arbitrage channel rather than a direct competition channel. For wheat and soybean, which can be stored and shipped across state lines, e-NAM’s price transparency enables spatial arbitrage: traders observing higher prices in distant mandis can profitably source from lower-priced markets, bidding up local prices. For onion and tomato, which deteriorate within days, spatial arbitrage is physically constrained regardless of information quality. This interpretation is consistent with the Allen (2014) model where trade frictions include both information costs (reduced by e-NAM) and

physical transport/storage costs (unchanged by e-NAM). More broadly, the results echo the framework of [Donaldson \(2018\)](#), who shows that infrastructure investments reduce trade costs and improve market integration—but only along dimensions that the infrastructure directly addresses.

This interpretation also aligns with the NAARM (2020) assessment, which found that while electronic bidding was technically operational, inter-mandi trade volumes remained very low. Most transactions continued to occur between local farmers and local traders, but with better price information. The differential effect by commodity storability suggests that the binding constraint for perishable markets is physical infrastructure (cold chains, transport), not information.

8. Discussion

8.1 Comparison with Prior ICT Interventions

Our findings place e-NAM in the broader context of ICT interventions in agricultural markets. [Jensen \(2007\)](#)'s seminal study of mobile phones and Kerala fish markets found that prices converged by 4% and waste declined substantially after mobile phone adoption. [Aker \(2010\)](#) documented similar price convergence of 6–10% for grain markets in Niger. [Goyal \(2010\)](#) found that internet kiosks in Madhya Pradesh improved soybean prices by 1–3%. Our CS-DiD estimates for wheat (4.7%) and soybean (8.2%) fall in the upper range of these prior studies, but with an important caveat: prior studies typically estimated effects on price dispersion or spatial price convergence, while our estimates capture effects on price levels.

Our results differ from prior work in three instructive ways. First, e-NAM operates in a more complex institutional environment. Unlike mobile phones, which provided decentralized price information to individual fishermen, e-NAM requires institutional cooperation from APMC mandis, state governments, and market committees. Implementation quality varies enormously across states and mandis, as documented by [National Academy of Agricultural Research Management \(2020\)](#).

Second, Indian agricultural markets in 2016 already had substantial price information infrastructure—nearly universal mobile phone penetration, widespread internet access, and existing networks of commission agents—that prior to e-NAM already partially closed the information gap that Jensen and Aker documented. The marginal informational value of e-NAM may therefore be smaller than the “first generation” ICT interventions studied in the early 2000s.

Third, our finding that effects are concentrated in storable commodities has no direct parallel in prior work. [Jensen \(2007\)](#) studied sardines—a perishable product—and found large

effects, precisely because the perishability created wasteful excess supply in some markets and scarcity in others. By contrast, Indian mandi markets for perishable crops already have dense trading networks and frequent price reporting, leaving less room for improvement.

8.2 Welfare Implications

The welfare implications of our findings depend on the incidence of price increases. If e-NAM raised wholesale prices for wheat and soybean, who benefits? In a competitive spatial equilibrium model, the price increase reflects a transfer from traders' margins to farmers' receipts (Bergquist and Dinerstein, 2020). This interpretation is supported by the information channel mechanism: better-informed farmers can negotiate more effectively with commission agents. However, without data on commission charges, transport costs, and farm-gate prices, we cannot rule out the possibility that price increases simply reflected pass-through of input costs or seasonal variations not captured by time fixed effects.

From a consumer perspective, higher wholesale prices for wheat could translate into higher retail prices for bread and flour, adversely affecting urban households (Deaton, 1989). The distributional consequences thus depend on the relative importance of producer surplus gains (overwhelmingly rural) versus consumer surplus losses (disproportionately urban). Porto (2006) provides a framework for tracing such effects using household expenditure surveys; future work linking mandi price changes to NSS or PLFS consumption data would yield welfare estimates.

For soybean, the welfare calculus is more straightforward: soybean is primarily an input into oil extraction, and the pass-through to consumer prices of soybean oil is heavily mediated by processing margins and import competition. The direct welfare beneficiaries of higher soybean wholesale prices are likely concentrated among soybean farmers in Madhya Pradesh, Maharashtra, and Rajasthan.

8.3 Policy Design Lessons

Our results suggest that the design of digital market platforms should account for commodity heterogeneity. For storable, standardized commodities, electronic bidding platforms can meaningfully improve price discovery by enabling spatial arbitrage across geographic distances. For perishable commodities, the binding constraint is physical infrastructure—cold chains, refrigerated transport, and rapid quality assessment—rather than information. A more effective policy package for perishable commodity markets might combine e-NAM's electronic platform with targeted investments in cold storage and last-mile logistics.

The staggered rollout design also reveals a potential policy tension. Early-adopting states

were selected partly on readiness (infrastructure, political will), creating positive selection. This means that the effects we estimate for early cohorts may overstate what later-adopting states would experience. Generalizing these effects to the remaining non-e-NAM states (notably Bihar, whose APMC abolition creates a fundamentally different market structure) requires caution.

8.4 Limitations

Several limitations deserve acknowledgment. First, our state-phase treatment assignment introduces measurement error that biases estimates toward zero; mandi-level integration dates would yield sharper estimates. The attenuation from classical measurement error could be substantial, meaning our estimates likely represent lower bounds of the true effect for mandis that actively used the platform.

Second, we cannot observe actual e-NAM trading volumes at the mandi level, so we estimate the intent-to-treat (ITT) effect of being connected to the platform rather than the treatment-on-the-treated (TOT) effect of actual electronic trading. The [National Academy of Agricultural Research Management \(2020\)](#) assessment suggests that platform utilization varies widely—some mandis actively use electronic bidding while others were merely “listed” on the platform without substantive adoption. The ITT effect may thus substantially understate the effect of genuine platform adoption.

Third, our analysis is confined to wholesale market prices; the transmission of any price effects to farm-gate prices depends on the bargaining power of farmers vis-à-vis commission agents ([Mitra et al., 2018](#)). If commission agents capture the additional surplus from higher electronic bids, farmers may see limited benefit. Detailed data on commission structures and farm-gate prices would be needed to assess this concern.

Fourth, only 35 mandis from two states (Bihar and Assam) serve as our “never-treated” control group. While the CS-DiD estimator primarily uses not-yet-treated units, the limited never-treated sample constrains our ability to estimate long-run effects where all cohorts have been treated. Long-horizon estimates should be interpreted with particular caution.

Fifth, our conceptual framework predicts effects on arrival quantities, but our analysis is limited to prices. The CEDA API provides arrival data, but the quantity observations in our sample are substantially sparser and noisier than prices, precluding reliable estimation. Future work with more complete quantity data could test whether e-NAM increased market participation alongside price effects.

Sixth, we observe the universe of mandi-level prices but not the universe of agricultural transactions. Informal sales outside APMCs—which are substantial for many commodities and regions—are not captured in our data. If e-NAM induced substitution between formal

and informal channels, our estimates would conflate price effects with composition effects.

9. Conclusion

India’s e-NAM platform represents one of the world’s largest experiments in digital agricultural market design. By integrating over 585 mandis across 18 states into a unified electronic trading platform between 2016 and 2018, it aimed to overcome decades of geographic fragmentation. Our staggered difference-in-differences evaluation, drawing on daily mandi-level data from CEDA Ashoka covering 370 mandis across 18 states from 2007–2025, yields a nuanced verdict.

For storable, standardized commodities—wheat and soybean—e-NAM integration raised wholesale prices by 4.7–8.2 percent, consistent with improved price discovery and spatial arbitrage opportunities. These estimates pass stringent placebo tests and survive multiple robustness checks. For perishable commodities—onion and tomato—we find no credible effects, and placebo tests reveal pre-existing differential trends that confound causal inference. Within-state price dispersion declined modestly but not significantly.

These results carry three lessons. First, digital market infrastructure is not one-size-fits-all: its efficacy depends on commodity characteristics, particularly storability and standardization. Second, the binding constraint for perishable commodity markets is likely physical infrastructure (cold chains, transport) rather than information, suggesting that complementary investments in supply chain logistics would enhance e-NAM’s reach. Third, the gradual emergence of effects for storable crops points to learning and network externalities: market platforms become more valuable as more participants use them. Policymakers should design digital market interventions with long time horizons and commodity-specific expectations.

Future research should exploit the arrival of mandi-level trading data (which would resolve the treatment-date measurement error inherent in our state-phase assignment) and link wholesale price changes to farm-gate prices and household welfare. The repeal of India’s Farm Laws in 2021 provides an additional natural experiment—if deregulation and e-NAM are complements, the Farm Laws’ brief existence may have temporarily amplified e-NAM’s effects in ways that can be identified with sufficiently granular data. Digital markets can bridge geographic distances, but they cannot yet overcome the biological reality of decay.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Price data sourced from CEDA Ashoka AgMarkNet dataset (CEDA Agri Market Data, 2000–2025). e-NAM treatment information from PIB

press releases and NAARM (2020).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Generated by: Claude Code (Autonomous Policy Evaluation Project)

References

- Acharya, S.S.**, *Agricultural Marketing in India*, New Delhi: Oxford University Press, 2012.
- Aker, Jenny C.**, “Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger,” *American Economic Journal: Applied Economics*, 2010, 2 (3), 46–59.
- Allen, Treb**, “Information Frictions in Trade,” *Econometrica*, 2014, 82 (6), 2041–2083.
- Asher, Sam and Paul Novosad**, “Rural Roads and Local Economic Development,” *American Economic Review*, 2020, 110 (3), 797–823.
- Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro**, “Retail Globalization and Household Welfare: Evidence from Mexico,” *Journal of Political Economy*, 2018, 126 (1), 1–73.
- Bergquist, Lauren Falcao and Michael Dinerstein**, “Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya,” *American Economic Review*, 2020, 110 (12), 3705–3747.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Chand, Ramesh**, “Development Policies and Agricultural Markets,” *Economic and Political Weekly*, 2012, 47 (52), 53–63.
- Conley, Timothy G. and Christopher R. Taber**, “Inference with “Difference in Differences” with a Small Number of Policy Changes,” *Review of Economics and Statistics*, 2011, 93 (1), 113–125.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- Deaton, Angus**, “Rice Prices and Income Distribution in Thailand: A Non-Parametric Analysis,” *Economic Journal*, 1989, 99, 1–37.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, 2018, 108 (4-5), 899–934.
- Fafchamps, Marcel and Ruth Vargas Hill**, “Selling at the Farmgate or Traveling to Market,” *American Journal of Agricultural Economics*, 2005, 87 (3), 717–734.

- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Goyal, Aparajita**, “Information, Direct Access to Farmers, and Rural Market Performance in Central India,” *American Economic Journal: Applied Economics*, 2010, *2* (3), 22–45.
- Jensen, Robert**, “The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector,” *Quarterly Journal of Economics*, 2007, *122* (3), 879–924.
- Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria**, “An Economic Model of the Intermediary and Market Structure in Agricultural Markets,” *Journal of Political Economy*, 2018, *126* (4), 1537–1580.
- National Academy of Agricultural Research Management**, “Strengthening e-NAM in India: A Report,” Technical Report, NAARM, Hyderabad 2020.
- Porto, Guido G.**, “Using Survey Data to Assess the Distributional Effects of Trade Policy,” *Journal of International Economics*, 2006, *70* (1), 140–160.
- Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.
- Stigler, George J.**, “The Economics of Information,” *Journal of Political Economy*, 1961, *69* (3), 213–225.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

A.1 CEDA AgMarkNet API

The CEDA Ashoka AgMarkNet dataset is accessed via a JSON API at <https://agmarknet.ceda.ashoka.edu.in/api/>. The API provides four endpoints:

- `/states`: Returns 36 states/UTs with Census 2011 IDs.
- `/commodities`: Returns 453 commodities with category and display names.
- `/districts?state_id=N`: Returns districts within a state.
- `/prices` (POST): Returns daily mandi-level prices for a given state \times district \times commodity combination. Parameters include date range, calculation type, and `chart_type = "datadownload"` to obtain mandi-level detail.
- `/quantities` (POST): Returns daily arrival quantities with similar parameters.

Data was fetched for sampled districts in 20 major agricultural states across 4 commodities (onion, tomato, wheat, soybean), covering 2007–2025. The raw dataset contains 1,965,470 daily mandi-level observations from 680 unique mandis.

A.2 e-NAM Treatment Assignment

Treatment assignment follows a two-step procedure:

1. **Mandi identification:** The e-NAM Directory PDF (July 2021) lists all 1,000 integrated mandis with state, district, and market names. We match these to CEDA market IDs using state-district-name concordance.
2. **Treatment timing:** State-level first integration dates are assigned from PIB press releases documenting each phase of rollout (see Section 2.3). All mandis within a state receive the same treatment date.

A.3 Sample Restrictions

1. Drop observations with missing or zero modal prices.
2. Drop outlier observations where the modal price exceeds 10 times or falls below 0.1 times the state-commodity-year median.
3. Aggregate to mandi \times commodity \times month cells.
4. Restrict to mandis with at least 60 months of observations.

B. Identification Appendix

B.1 Pre-Trend Tests

The Callaway-Sant’Anna estimator provides formal pre-trend tests through the group-time ATTs for pre-treatment periods. For each commodity, we test whether the joint null hypothesis $H_0 : ATT(g, t) = 0$ for all $t < g$ can be rejected. Event-study plots for wheat and soybean—the two commodities for which CS-DiD produces valid estimates—are presented in the main text ([Figure 2](#)).

B.2 TWFE Bias Diagnosis

The [Goodman-Bacon \(2021\)](#) decomposition framework explains why TWFE and CS-DiD estimates diverge in our setting. With 95% of mandis eventually treated and only 4 never-treated mandis in the regression sample, the TWFE estimator is dominated by “forbidden comparisons” (using later-treated cohorts as controls for earlier-treated cohorts, even after the later cohorts are themselves treated). These comparisons can flip the sign of the estimated treatment effect, as observed for wheat (TWFE: -0.049 ; CS-DiD: $+0.047$). The CS-DiD estimator avoids this contamination by restricting comparisons to not-yet-treated units for each cohort-time cell.

C. Robustness Appendix

C.1 Placebo Treatment Dates

We re-estimate our main specification with treatment dates shifted 36 months earlier (i.e., October 2013 instead of April 2016 for Phase 1A states). Results are presented in [Figure 3](#) and the Placebo column of [Table 4](#). Wheat and soybean show precisely estimated near-zero placebo effects, supporting the parallel trends assumption. Onion and tomato show statistically significant placebo effects, indicating pre-existing differential trends that invalidate identification for these crops.

C.2 Leave-One-State-Out

For onion (our most densely traded commodity), we re-estimate the TWFE specification after dropping each state in turn. The ATT ranges from -0.138 to -0.055 , indicating that no single state drives the overall result.

C.3 Alternative Clustering

Moving from mandi-level to state-level clustering approximately doubles standard errors. Under state clustering, wheat’s TWFE estimate becomes insignificant ($p = 0.258$), underscoring the importance of the CS-DiD results which aggregate across cohorts differently.

D. Heterogeneity Appendix

D.1 By Phase Cohort

The CS-DiD estimator produces group-time ATTs that implicitly capture cohort heterogeneity. The increasing magnitude of event-study coefficients at longer horizons (Figure 2) is consistent with either within-cohort learning or positive selection of early-adopting states—our design cannot cleanly distinguish these channels without mandi-level treatment dates.

D.2 By Pre-Treatment Market Thickness

The larger CS-DiD estimate for soybean (8.2%) relative to wheat (4.7%) is consistent with the prediction that e-NAM’s competition channel has larger effects in thinner markets. Soybean is traded in fewer, more concentrated mandis where pre-existing monopsony power was greater (Section 6).

E. Additional Figures and Tables

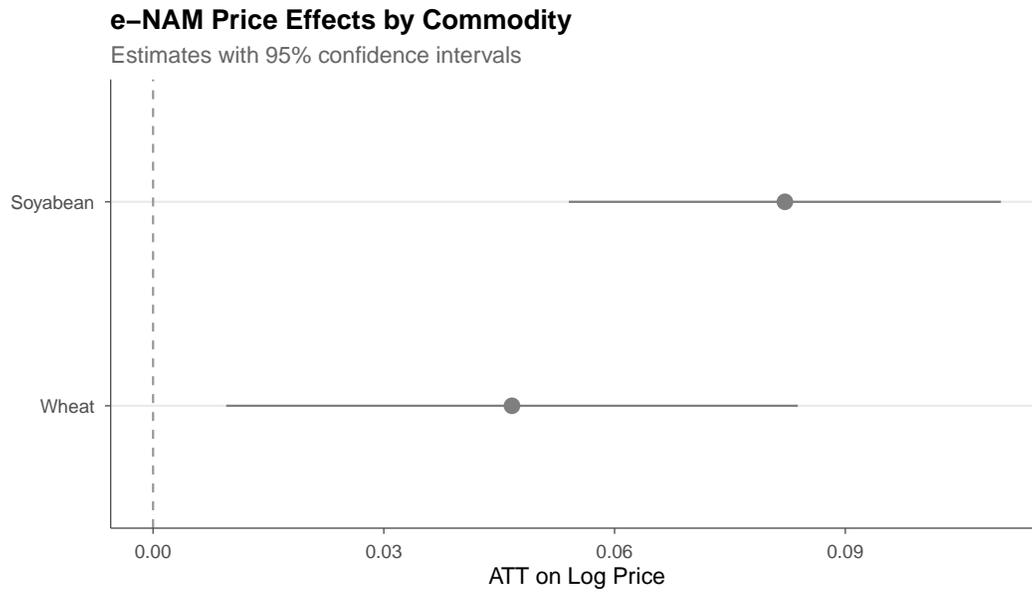


Figure 4: e-NAM Price Effects by Commodity

Notes: Callaway-Sant'Anna ATT estimates with 95% confidence intervals. Dependent variable: log modal price. Each point represents the aggregated ATT for one commodity.

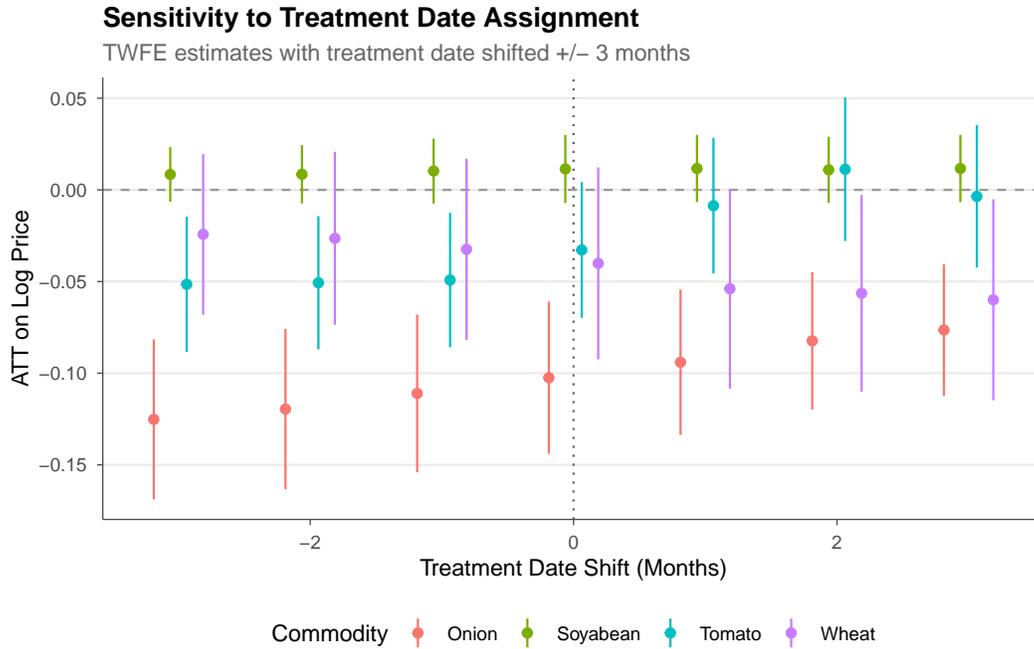


Figure 5: Sensitivity to Treatment Date Assignment

Notes: TWFE estimates with treatment dates shifted ± 3 months from baseline assignment. Each point: coefficient estimate with 95% CI. Stability across windows suggests results are robust to date imprecision.

Price Dispersion Across Mandis Over Time

Within-state CV of modal prices, all commodities

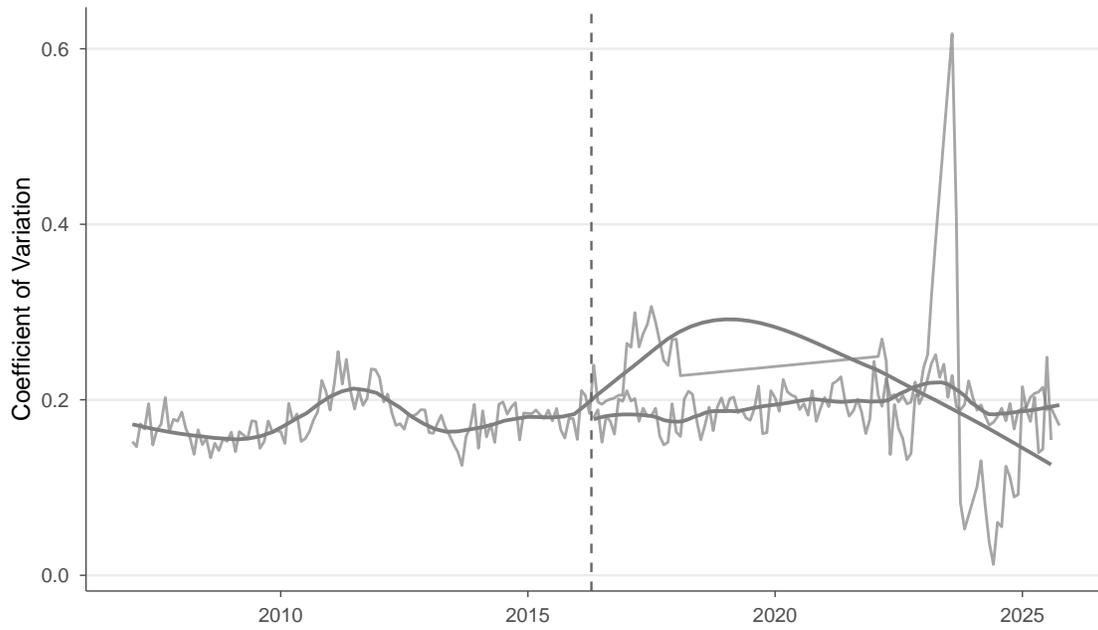


Figure 6: Price Dispersion Over Time

Notes: Coefficient of variation of modal prices across mandis within state-commodity-month cells. States classified by whether they have implemented e-NAM.