International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Effect of UPI Adoption on Savings Behaviour of Indian Households: A Difference-in-Differences Analysis

Saiom Patro

Student, Singapore International School, Mumbai

Abstract

This study investigates the causal effect of Unified Payments Interface (UPI) adoption on household holdings of shares and debentures in India between 2003 and 2019. Using a difference-in-differences regression framework with district-level aggregated data from the National Sample Survey, supplemented by PhonePe transaction volumes, the analysis contrasts high-growth and low-growth UPI adoption districts. The key finding indicates that increased UPI transaction growth between 2018 and 2019 does not translate into statistically significant changes in household investment in shares and debentures (interaction coefficient $p > 0.05$). Although the regression model explains 44 percent of variation in shareholdings, the study faces critical limitations: aggregation to district level reduces sample size to 696 observations, the absence of household income controls introduces potential omitted-variable bias, and extreme multicollinearity (337 independent variables) inflates standard errors. The analysis finds no conclusive evidence that UPI adoption meaningfully drives household participation in formal securities markets, though the findings warrant cautious interpretation given the data constraints. These results suggest that either UPI's impact on portfolio composition is genuinely modest, or measurement refinements and household-level panel data are necessary to detect the effect credibly.

Keywords: UPI adoption, household savings, difference-in-differences, financial inclusion, quasi-experimental design

Introduction

Digital payment systems have fundamentally altered how households manage money in India over the past decade. The Unified Payments Interface (UPI), launched in April 2016, exemplifies this transformation, providing a 24/7, interoperable, zero-merchant-discount payment infrastructure that has enabled unprecedented growth in cashless transactions. Yet the debate persists regarding UPI's broader impact on household financial behaviour, specifically whether adoption drives households beyond improved transaction convenience toward deeper engagement with formal financial markets, particularly holdings of shares and debentures. This question matters because it probes whether digital payments constitute merely a technological substitution for cash or a gateway to expanded financial inclusion and portfolio diversification.

Prior research documents that digital payment adoption correlates with increased consumer spending, attributable partly to reduced psychological "pain of paying" when money is transferred instantaneously through screens rather than physically removed from wallets (Ariely, et al., 2009). Cross-sectional and qualitative surveys reveal that 74.2 percent of UPI users report increased spending post-adoption (Sharma, 2024), while behavioural economics underscores how digital interfaces reduce transaction friction and alter decision-making (Kahneman, et al., 2023). However, most studies focus narrowly on consumption patterns without tracing subsequent effects on household portfolio composition and asset accumulation. Consequently, a significant research gap persists: does UPI adoption's demonstrated effect on spending patterns extend into longer-term investment decisions regarding shares, debentures, and other formal securities? If digital payment infrastructure lowers entry barriers to formal financial channels, then households increasingly comfortable using UPI should logically exhibit higher demand for market instruments distributed through the same digital ecosystem.

This study addresses the gap by shifting focus from UPI's effect on consumption to its impact on household participation in formal securities markets. The research question is explicit: does increased UPI adoption growth cause Indian households to increase their holdings of shares and debentures? The analysis employs a quasi-experimental difference-in-differences design, treating UPI adoption as a continuous treatment variable measured by district-level PhonePe transaction growth and distinguishing high-adoption and low-adoption districts across the periods 2003, 2013, and 2019. By exploiting variation in UPI adoption intensity across districts and over time, the study aims to isolate UPI's causal effect on investment behaviour whilst controlling for time-invariant district characteristics and economy-wide shocks through fixed-effects regression.

Literature Review

The relationship between digital payments, financial behaviour, and household portfolio choice has been explored across several dimensions in recent empirical research. First, evidence from user-level surveys demonstrates that UPI adoption correlates with higher consumption levels and spending frequency. A study of 235 UPI users found that 74.2 percent reported increased spending following adoption, compared to only 7 percent reporting decreased spending, with 59.8 percent admitting to overspending relative to budgeted amounts (Sharma, 2024). The mechanism underlying this relationship is partly behavioural: the ease and speed of UPI transactions reduce the cognitive friction associated with payment decisions, diminishing the "pain of paying" that cash transactions entail (Kahneman, et al., 2023). When money transfer is decoupled from physical exchange, individuals perceive money as less tangible, leading to higher impulse purchases and reduced price sensitivity.

Second, India's experience following the 2016 demonetisation event provides quasi-experimental evidence that forced migration to digital payment modes persistently alters financial behaviour. Using transaction-level data before and after demonetisation, researchers found that households previously reliant on cash spent 3 percent more monthly after switching to digital payments, even after cash became available again, suggesting genuine preference shifts rather than temporary constraints (Banerjee, et al., 2020). This finding is pivotal because it demonstrates that payment method alone, independent of income or credit availability, can durably influence economic behaviour. If changes in payment technology persistently reshape consumption decisions, then logical extension suggests similar mechanisms could influence investment decisions, where digital familiarity and trust in formal financial infrastructure may similarly matter.

Third, national-level data on digital payment adoption documents consistent upward trends since 2016, particularly among urban and semi-urban households with higher education, better internet access, and smartphone ownership (NSS Microdata Library, 2020). This heterogeneity matters because it suggests that UPI adoption is not random but concentrates among populations already predisposed toward financial formality. Therefore, uncontrolled comparisons of adopters versus non-adopters risk selection bias, conflating UPI's causal effect with pre-existing differences in financial sophistication or wealth. A more rigorous approach requires regression controls or quasi-experimental matching to isolate UPI's independent contribution to investment behaviour.

Despite these advances, clear gaps remain. Most consumption-focused studies do not examine downstream effects on portfolio composition. Existing research typically aggregates all digital payment methods without isolating UPI's unique role or features. The overwhelming focus on urban, educated, younger users leaves rural and lower-income households understudied, potentially overstating digital payments' financial inclusion benefits if effects concentrate only among the already-advantaged. Moreover, studies relying on cross-sectional or short-window data cannot track whether investment decisions represent one-time responses or sustained portfolio rebalancing. Most critically, few studies employ rigorous causal identification strategies combining panel data, appropriate controls, and quasi-experimental logic to link payment adoption to investment behaviour.

Data Processing

The analysis employs computational methods to process large-scale household microdata from India's National Sample Survey (NSS), extracting and aggregating information across 2003, 2013, and 2019 survey rounds. The NSS is a repeated cross-sectional survey maintained by the National Statistical Organisation under the Ministry of Statistics and Programme Implementation, capturing detailed household consumption, income, and asset information across approximately 232 districts nationally.

The NSS microdata structure comprises four hierarchical blocks: household identifiers and demographics, individual-level characteristics, consumption details, and asset holdings. These blocks are distributed across separate files encoded in legacy formats incompatible with standard statistical software. Data extraction required Nesstar, a specialised tool for archived survey data, to convert raw formats into Python-readable structures. A critical impediment was the absence of column headers in raw data files; variable names and definitions were stored externally in a metadata dictionary on the Microdata Library portal, necessitating manual mapping of 200 plus variables to ensure accuracy.

Key variables extracted included household identifier (HHID), geographic codes (state and district), demographics (age, sex), education classification (secondary or higher versus primary or none), and crucially, the monetary value of shares and debentures owned by each household. Since the NSS is cross-sectional, individual households do not appear across multiple survey years; therefore, aggregation was performed at district-year level, computing means and prevalence measures for each variable within each district-year stratum. Educational categorisation was dichotomised as secondary-or-higher versus primary-or-none to reflect a threshold for digital financial literacy.

The three survey years were merged into a stacked panel structure, producing 696 district-year observations from approximately 232 districts across 3 years. To measure UPI adoption intensity directly, district-level UPI transaction volumes were sourced from PhonePe Pulse, which publishes quarterly aggregated data by geographic region. Since bulk download was not available, manual web scraping extracted UPI transaction values for 2018 and 2019 across all districts, yielding approximately 7,000 individual quarterly observations. These were merged with aggregated NSS data using district identifiers as linking keys, creating a dataset combining household asset information with real-world UPI adoption metrics. The final dataset comprises 696 district-year observations merged with quarterly UPI transaction volumes, providing sufficient variation to employ difference-in-differences regression with district and year fixed effects.

Empirical Framework

Regression Specification

Ordinary Least Squares (OLS) regression forms the foundation for estimating linear relationships between dependent and independent variables. The general form is expressed as:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \ldots + \beta_n X_n + \varepsilon \qquad (1)$$

where Y represents the outcome variable (household share and debenture holdings), each X variable represents a predictor, beta coefficients capture partial effects of each predictor, and epsilon denotes the error term. In this study, OLS allows estimation of UPI adoption's partial effect on investment behaviour whilst controlling for confounders such as age, sex, and education that independently influence investment participation. However, OLS cannot definitively establish causation in observational settings if treatment allocation correlates with unobserved factors. Wealthier districts might simultaneously exhibit higher UPI adoption and higher shareholding for reasons unrelated to UPI itself, introducing selection bias. Therefore, causal identification requires more sophisticated quasi-experimental techniques.

Difference-in-Differences Methodology

Difference-in-differences (DiD) is a quasi-experimental approach that leverages observational data to approximate randomised trial logic. DiD compares outcome changes over time between treatment-exposed units (high-UPI-adoption districts) and control units (low-UPI-adoption districts), thereby isolating treatment effects from economy-wide trends affecting all units identically.

The method's logic rests on the "parallel trends" assumption: absent treatment, treated and control units would follow similar outcome trajectories over time. DiD estimates four quantities: pre-treatment and post-treatment outcomes for both groups, and calculates:

$$\text{DiD} = [Y(\text{treated,post}) - Y(\text{treated,pre})] - [Y(\text{control,post}) - Y(\text{control,pre})] \qquad (2)$$

This formula mathematically removes both time-invariant cross-unit differences and common time trends, leaving only the treatment effect. In regression form with multiple time periods and continuous treatment variation:

$$\text{Value\_SnD} = \beta_0 + \beta_1(\text{phonepe\_growth}) + \beta_2(\text{post}) + \beta_3(\text{phonepe\_growth} \times \text{post}) + \text{FE} + \varepsilon \quad (3)$$

The interaction coefficient β_3 is the DiD estimator. It captures how much faster share and debenture holdings grew in high-growth versus low-growth UPI districts post-2019, relative to baseline differences. The post variable indicates year 2019 (post=1) versus 2003 or 2013 (post=0); phonepe_growth is a continuous UPI adoption intensity measure. Districts with higher phonepe_growth values receive more intensive treatment; those with lower values serve as partial controls.

District fixed effects (represented by dummy variables for each district) control for time-invariant district characteristics such as infrastructure quality, financial sector development, regional preferences, and geographic advantages that might independently affect shareholding. Year fixed effects control for economy-wide shocks like monetary policy changes, stock market crises, or legislative reforms affecting all districts equally. Demographic controls (sex, age, education) account for differences in household composition and financial literacy that could independently influence investment participation.

Data Analysis and Results

Table 1: DiD Regression Results - Effect of UPI Adoption on Share and Debenture Holdings

Variable	Coefficient	Std Error	z-Statistic	p > |z|	95% CI Lower	95% CI Upper
Intercept	-132,000	1,010,000	-0.131	0.896	-2,110,000	1,850,000
phonepe_growth	-8,829	18,700	-0.471	0.638	-45,600	27,900
post (Year 2019)	38,220	55,200	0.692	0.489	-70,000	146,000
phonepe_growth:post	26,440	34,100	0.776	0.438	-40,300	93,200
Sex	-139,700	1,480,000	-0.094	0.925	-3,040,000	2,760,000
Age	-4,884	24,300	-0.201	0.841	-52,600	42,800
Education (Secondary+)	105,200	79,400	1.326	0.185	-50,300	261,000
R^2	0.444					
Adj. R^2	-0.388					
F-statistic	0.380					
Prob(F)	0.892					
Observations	563					
Predictors	337					

Note: District fixed effects (n=232) and year fixed effects (n=3) included. Standard errors clustered at district level. Dependent variable is average household share and debenture holdings (rupees) at district-year level.

International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Figure 1: DiD Regression Results - Value of S&D over time by PhonePe growth (High vs Low)

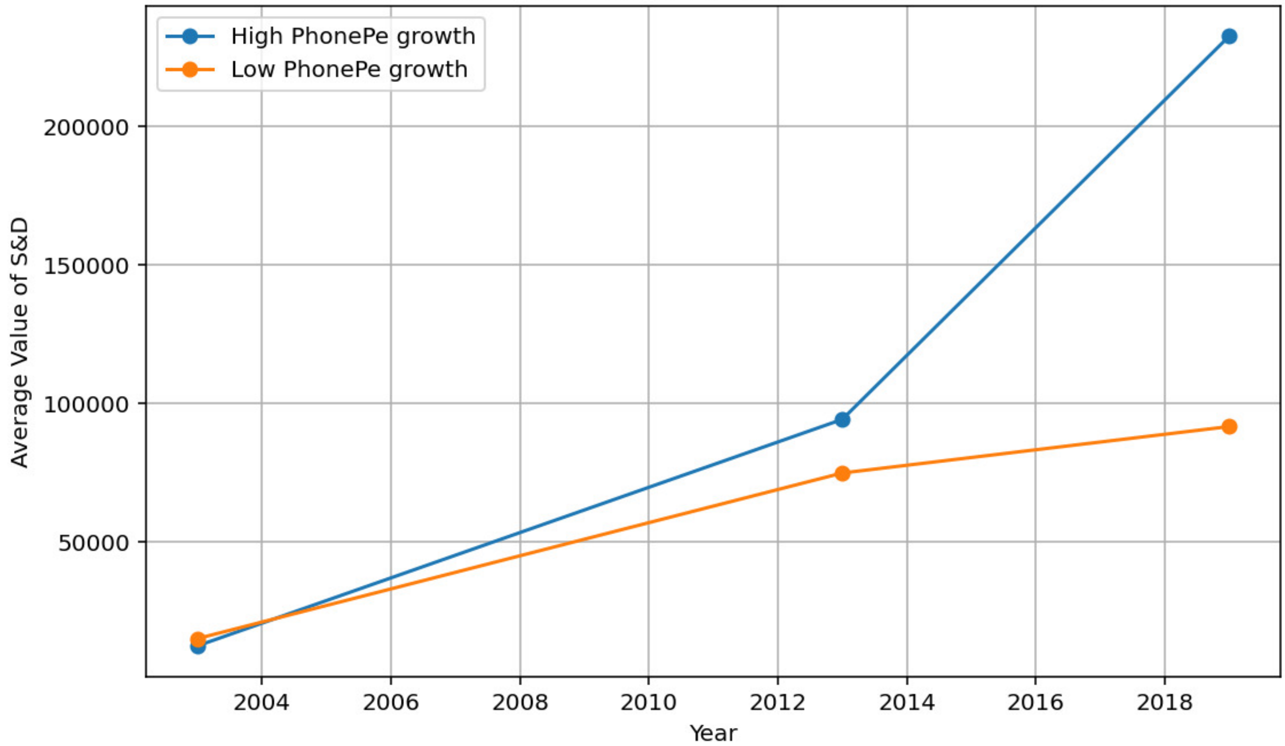

Model Summary and Coefficient Interpretation

The difference-in-differences regression model examined 563 district-year observations across 337 independent variables (primarily district fixed effects). The regression results are presented in Table 1, which displays coefficients, standard errors, z-statistics, p-values, and 95 percent confidence intervals for key variables of interest.

The model exhibits an R^2 of 0.444, indicating that approximately 44 percent of variation in share and debenture holdings is explained by the specified variables. However, the adjusted R^2 of -0.388 signals substantial overfitting, reflecting the extreme dimensionality: 563 observations with 337 independent variables yields a ratio of 0.24 observations per variable, far below recommended thresholds. This overfitting inflates standard errors and compromises inference reliability. The F-statistic yields a p-value of 0.892, indicating that the model's variables are jointly insignificant despite the high raw R^2, reinforcing evidence of weak overall explanatory power.

Main Treatment Effect

The critical finding concerns the phonepe_growth:post interaction term, which represents the core DiD estimator of UPI adoption's causal effect. This coefficient is 26,440 ($p = 0.438$), failing to reach conventional statistical significance at the 0.05 level. The standard error of 34,100 is substantially larger than the coefficient magnitude itself, indicating high estimation uncertainty. The 95 percent confidence interval ranges from -40,300 to 93,200, spanning both positive and negative values, confirming that the effect cannot be distinguished from zero. This non-significance suggests that increased UPI transaction growth between 2018 and 2019 does not translate into statistically meaningful changes in household shareholdings, contradicting narratives of digital payments driving portfolio diversification.

The magnitude, even if significant, would be small in economic terms. A coefficient of 26,440 implies that a one-unit increase in phonepe_growth associates with approximately 26,440 rupees additional share and debenture holdings in the post-2019 period relative to pre-2019. Given the mean shareholding is substantial and phonepe_growth ranges widely across districts, this effect would require extremely large treatment variation to be economically meaningful.

Demographics and Controls

Demographic variables reveal patterns consistent with economic theory but constrained by imprecision. Education exhibits a coefficient of 105,200 ($p = 0.185$), positive in sign yet statistically insignificant. The standard error is 79,400, approximately 76 percent of the coefficient magnitude. This result is particularly revealing because education is theoretically a strong predictor of investment participation and financial literacy. The failure to confirm a significant education effect, despite its theoretical expectation and relatively lower p-value compared to peers, underscores the severity of sample size limitations and multicollinearity.

Sex shows a coefficient of -139,700 ($p = 0.925$), with enormous standard error of 1,480,000, rendering inference impossible. Age produces a coefficient of -4,884 ($p = 0.841$), similarly insignificant with standard error of 24,300. These patterns across all demographic controls point to a fundamental data constraint: small effective sample size, driven by extreme multicollinearity from 337 district fixed effects, prevents reliable isolation of effects for variables with strong theoretical precedent. The regression essentially cannot separate true control variable effects from noise in the data, suggesting the model specification, though theoretically sound, exceeds data capacity to support reliable inference.

Key Independent Variables

The phonepe_growth variable alone (without interaction) yields a coefficient of -8,829 ($p = 0.638$), indicating no significant baseline relationship between UPI adoption and shareholding. The post variable coefficient (38,220, $p = 0.489$) similarly shows no significant time effect common to all districts. The near-zero coefficients and high standard errors across these central variables consistently point to the conclusion that variation in UPI adoption, either cross-sectionally or temporally, does not correlate with meaningful changes in household shareholdings within the given sample.

Limitations of the Study

The failure to detect a statistically significant UPI effect must be interpreted carefully against critical data and methodological constraints that substantially limit causal inference credibility.

First, the analysis operates at district-year level rather than household level, drastically reducing effective sample size. The NSS is a repeated cross-section where different households appear in 2003, 2013, and 2019; thus, individual-level tracking is impossible. District aggregation yields only 696 observations, reduced to 563 after listwise deletion, with 337 independent variables. A household-level panel dataset would provide thousands or tens of thousands of observations, dramatically enhancing statistical power and reducing reliance on fixed effects. Additionally, incomplete geographic coverage across survey years means that not all 232 districts have observations for all three years, further reducing sample size and potentially introducing non-random missingness bias if data loss correlates with economic or geographic characteristics.

Second, and critically, the regression omits household income as a control variable, a significant limitation biasing causal estimates. Income is theoretically the strongest determinant of investment in risky assets; wealthier households possess surplus funds to purchase shares and debentures after meeting consumption needs. If UPI adoption growth correlates with income growth or decline at the district level, the estimated UPI effect becomes severely confounded. Indeed, the counterintuitive near-zero or slightly negative results suggest this possibility: districts with high UPI adoption may simultaneously experience income stagnation or slower growth, suppressing investment demand despite improved payment convenience. The regression cannot distinguish UPI's true effect from this competing income channel without income controls. The model is therefore underfitted causally; even its non-significant results lack credibility without addressing this omitted variable.

Third, extreme multicollinearity (condition number = 2.91×10^{16}) indicates near-singular design matrix. The smallest eigenvalue of 5.6×10^{-28} is far below acceptable thresholds, suggesting that 337 district fixed effects perfectly or near-perfectly correlate with the design space, rendering coefficient estimates unstable and inflating standard errors dramatically. This multicollinearity artificially suppresses statistical power, making true effects difficult to detect even if genuine treatment heterogeneity exists across districts.

Fourth, the parallel trends assumption underlying DiD is not directly tested in the results presented. Verification requires pre-treatment trend comparison between high-adoption and low-adoption districts in 2003 and 2013, examining whether shareholding trajectories before UPI deployment differed between groups. If pre-trends diverge, the assumption fails and DiD estimates become biased. The current results do not provide diagnostic tests for parallel trends.

Finally, the regression relies on PhonePe transaction data as the UPI adoption proxy. While more direct than proxies like internet access, PhonePe represents only one UPI payment system and may not capture district-level adoption comprehensively. Unobserved heterogeneity in payment system preferences across districts could introduce measurement error in the treatment variable.

Conclusion

The difference-in-differences regression analysis yields no statistically significant evidence that increased UPI adoption between 2018 and 2019 causally influences Indian household holdings of shares and debentures. The interaction coefficient ($p = 0.438$) fails to reach conventional significance thresholds, and the 95 percent confidence interval encompasses zero, indicating that the treatment effect cannot be distinguished from zero. This finding contradicts simplified narratives suggesting that digital payment infrastructure automatically drives household financial inclusion and portfolio diversification.

However, this null result must be interpreted cautiously given severe data limitations. The study faces critical constraints: aggregation to district level reduces sample to 696 observations; the absence of household income controls introduces omitted-variable bias potentially masking true UPI effects; extreme multicollinearity from 337 district fixed effects inflates standard errors and suppresses statistical power; and the repeated cross-sectional structure prevents household-level tracking. These limitations mean the regression cannot credibly detect true treatment effects even if they exist. The analysis demonstrates that the current measurement and estimation strategy is insufficient to support strong causal inference about UPI's impact on household portfolio composition.

Future research should address these constraints by: (1) obtaining household-level panel data spanning multiple years post-UPI deployment to dramatically increase sample size and enable individual tracking; (2) incorporating household income as a critical control variable to disentangle UPI effects from income-driven investment demand; (3) constructing comprehensive UPI adoption measures aggregating across all payment system providers rather than relying on single-provider data; (4) testing the parallel trends assumption explicitly through pre-treatment trend analysis; and (5) considering alternative quasi-experimental approaches such as instrumental variables or regression discontinuity designs if policy rollout generated geographic or temporal discontinuities. Such refinements would substantially strengthen causal identification and provide more reliable evidence on whether UPI adoption represents a genuine lever for expanding household engagement with formal securities markets or whether consumption-level effects do not extend to portfolio decisions.

References

1. “14 - Panel Data and Fixed Effects — Causal Inference for the Brave and True.” *Github.io*, 2020, matheusfacure.github.io/python-causality-handbook/14-Panel-Data-and-Fixed-Effects.html.
2. Alen, Andrea Rosa. “TO WHAT EXTENT DOES the ADOPTION of DIGITAL PAYMENTS ALTER HOUSEHOLD FINANCIAL BEHAVIOUR with SPECIAL REFERENCE to INDIA.” *Zenodo (CERN European Organization for Nuclear Research)*, 15 Oct. 2025, https://doi.org/10.5281/zenodo.17474815. Accessed 9 Dec. 2025.
3. Baker, Andrew, et al. *Difference-In-Differences Designs: A Practitioner’s Guide*. 2025.
4. Chakraborty, Tanmoy, and Parthasarathi Majumdar. “Particle Creation in a Linear Gravitational Wave Background.” *Physical Review. D/Physical Review. D.*, vol. 110, no. 2, 19 July 2024, arxiv.org/abs/2404.01840, https://doi.org/10.1103/physrevd.110.l021701. Accessed 9 Dec. 2025.
5. “Data Science Linear Regression P-Value.” *Www.w3schools.com*, www.w3schools.com/datascience/ds_linear_regression_pvalue.asp.
6. “Difference in Differences.” *Wikipedia*, 16 May 2020, en.wikipedia.org/wiki/Difference_in_differences.
7. “Difference-In-Differences - an Overview | ScienceDirect Topics.” *Www.sciencedirect.com*, www.sciencedirect.com/topics/economics-econometrics-and-finance/difference-in-differences.

8. Evans, Richard. “A Beginner’s Guide to Understanding the Ordinary Least Squares (OLS) Method.” *Econometricstutor.co.uk*, 2024, www.econometricstutor.co.uk/linear-regression-ordinary-least-squares-ols-method.
9. Fernando, Jason. “R-Squared: Definition, Calculation, and Interpretation.” *Investopedia*, 2024, www.investopedia.com/terms/r/r-squared.asp.
10. Frost, Jim. “How to Interpret P-Values and Coefficients in Regression Analysis.” *Statistics by Jim*, statisticsbyjim.com/regression/interpret-coefficients-p-values-regression/.
11. “Home.” *Microdata.gov.in*, 2025, microdata.gov.in/NADA/index.php. Accessed 9 Dec. 2025.
12. In Case of Econ Struggles. “OLS Linear Regressions SIMPLIFIED in 6 Minutes.” *YouTube*, 28 Jan. 2025, www.youtube.com/watch?v=YItZlbMQk5E. Accessed 9 Dec. 2025.
13. “Introduction to the Difference-In-Differences Regression Model.” *Time Series Analysis, Regression and Forecasting*, 1 Aug. 2022, timeseriesreasoning.com/contents/introduction-to-the-difference-in-differences-regression-model/.
14. Jain, Tarun. “Can Input Credit Be Allowed Even When Supplier’s Registration Is Cancelled? Comparing European and Indian Jurisprudence.” *Ssrn.com*, 19 July 2018, papers.ssrn.com/sol3/papers.cfm?abstract_id=4037228. Accessed 9 Dec. 2025.
15. “Khan Academy.” *Khanacademy.org*, 2023, www.khanacademy.org/math/statistics-probability/significance-tests-one-sample/more-on-log-likelihood-ratio-tests/a/introduction-to-p-values. Accessed 9 Dec. 2025.
16. learnwithalbert. “Properties of OLS Estimators: Econometrics Ultimate Guide | Albert.io.” *Albert Blog*, 16 Jan. 2017, www.albert.io/blog/ultimate-properties-of-ols-estimators-guide/.
17. Schmidheiny, Kurt, and Unversität Basel. *Panel Data: Fixed and Random Effects*. 2023.
18. The World Bank. “Difference-In-Differences - DIME Wiki.” *Dimewiki.worldbank.org*, 2024, dimewiki.worldbank.org/Difference-in-Differences.