

Tax Harmonization and Price Convergence: Evidence from India's Goods and Services Tax

APEP Autonomous Research* @ailscl

February 25, 2026

Abstract

India's 2017 Goods and Services Tax (GST) replaced a fragmented patchwork of state-level indirect taxes with a unified national structure. I exploit cross-state variation in pre-reform tax intensity to estimate the effect of tax harmonization on interstate price convergence using a continuous-treatment difference-in-differences design. Drawing on monthly state-level CPI data for 35 states from 2013 to 2025, I find that states with higher pre-GST tax burdens experienced modestly lower price growth after reform ($\beta = -0.012$, RI $p = 0.056$), though this effect attenuates with state-specific trends. A triple-difference specification exploiting commodity-level variation in GST rate changes—which absorbs all state-specific time shocks—yields a strongly significant effect ($\beta = 0.008$, $p = 0.014$). Pre-trend tests cannot reject parallel trends ($F = 1.51$, $p = 0.20$), and the triple-difference result identifies rate harmonization as the operative channel.

JEL Codes: H25, H77, F15, O18

Keywords: GST, tax harmonization, price convergence, market integration, India

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

1. Introduction

In a country as geographically vast and administratively fragmented as India, a mango purchased in Kerala and an identical mango sold in Uttar Pradesh could carry wildly different tax burdens. Before July 2017, India’s indirect tax system was a labyrinth: state-level VATs set independently by 29 states and 7 union territories, a Central Sales Tax levied on interstate transactions, entry taxes and octroi at state and municipal borders, and a cascade of excise duties that compounded at every stage of production. The result was not merely administrative complexity—it was an internal trade barrier rivaling many international borders (Engel and Rogers, 1996; Barnwal et al., 2023).

The Goods and Services Tax (GST), implemented on July 1, 2017, replaced this patchwork with a destination-based, multi-rate national tax. It was, by any measure, the most ambitious tax reform in India’s post-independence history, covering over 1,200 commodities across four rate slabs (5%, 12%, 18%, and 28%) and eliminating the Central Sales Tax entirely. The reform’s architects argued that harmonized rates and the abolition of interstate check-posts would reduce price distortions across states, integrate India’s internal market, and bring consumer prices closer together (Agarwal, 2017; Subramanian, 2014).

But did it? The theoretical prediction is clear: removing tax-induced price wedges should produce price convergence, with the largest effects in states that had the highest pre-reform tax burdens and in commodities that experienced the largest effective rate changes. Yet empirical evidence on whether GST achieved this convergence remains thin. Most existing studies rely on aggregate price indices or simple before-after comparisons that cannot disentangle the reform’s effect from concurrent macroeconomic shocks—most notably the demonetization of November 2016 and the COVID-19 pandemic beginning in early 2020 (Kumar and Sharma, 2023).

This paper provides the first rigorous causal estimate of GST’s effect on interstate price convergence. I exploit the fact that while GST was implemented nationally, its *bite* varied systematically across states because pre-reform indirect tax burdens differed substantially—from as low as 2.0% of gross state domestic product (GSDP) in Lakshadweep to 8.9% in Goa. This variation creates a continuous-intensity difference-in-differences design: all states received the same post-reform tax structure, but states with higher pre-GST tax intensity experienced a larger effective change in their tax regime, and hence a larger expected price adjustment.

I construct a panel of monthly consumer price indices (CPI) for 35 states and union territories across 7 commodity groups from January 2013 to December 2025, drawn from the Ministry of Statistics and Programme Implementation (MoSPI) eSankhyiki database. Two

months (April–May 2020) are missing due to the suspension of CPI data collection during India’s COVID-19 lockdown, yielding 154 months: 54 pre-reform and 100 post-reform—an unusually long window that allows careful examination of pre-trends and dynamic effects.

The baseline specification regresses log CPI on the interaction of a post-GST indicator with standardized state-level tax intensity, absorbing state and month fixed effects with standard errors clustered at the state level. I find a coefficient of -0.012 (SE = 0.006, $p = 0.069$), suggesting that a one-standard-deviation increase in pre-GST tax intensity is associated with approximately 1.2 percentage points lower cumulative price growth after the reform. While only marginally significant in conventional terms, randomization inference yields a two-sided p -value of 0.056, and the estimate is remarkably stable across specifications: excluding the demonetization window ($\beta = -0.012$), restricting to the pre-COVID period ($\beta = -0.009$), and in leave-one-out analysis where the coefficient ranges from -0.014 to -0.008 across all 35 jackknife samples.

I strengthen identification with a triple-difference design that exploits commodity-level variation in GST rate changes. The GST affected different commodity groups asymmetrically: food items generally saw tax decreases (from $\sim 5\%$ state VAT to 0–5% GST), while services in the “Miscellaneous” category experienced increases (from 15% service tax to 18% GST). Fuel and light products were explicitly excluded from GST, providing a natural placebo. The triple-difference specification interacts post-GST status, state-level tax intensity, and the absolute magnitude of commodity-level tax changes, absorbing state \times commodity, commodity \times time, and state \times time fixed effects. This demanding specification yields a coefficient of 0.008 ($p = 0.014$), indicating that the price effects are concentrated precisely where theory predicts: in commodity-state cells where both state tax intensity and commodity tax changes were large.

Threats to identification are carefully addressed. The pre-trend test—regressing log CPI on interactions of year dummies with tax intensity in the pre-GST period—produces a joint F -statistic of 1.51 ($p = 0.196$), consistent with parallel trends. Including state-specific linear trends attenuates the baseline coefficient to -0.006 ($p = 0.270$), suggesting some sensitivity to trend assumptions; the triple-difference specification, which absorbs all state-level time variation, is the paper’s primary design for this reason. The time-placebo test, which artificially assigns a reform date of July 2015, yields a small and insignificant coefficient ($\beta = -0.006$, $p = 0.102$). The binary treatment specification—comparing states above and below the median pre-GST tax burden—confirms the direction and magnitude ($\beta = -0.021$, $p = 0.104$). Rural areas show larger convergence effects than urban areas ($\beta_{\text{rural}} = -0.011$ vs. $\beta_{\text{urban}} = -0.004$), consistent with rural markets being more fragmented ex ante.

The data reveal one surprise: fuel prices converged despite being excluded from the

reform. Fuel and Light, which should serve as a placebo given its exclusion from GST, shows a significant negative coefficient ($\beta = -0.058$, $p = 0.005$). I argue this likely reflects the concurrent elimination of interstate check-posts and CST—which affected fuel logistics and distribution costs even though fuel tax rates themselves were unchanged—rather than a violation of the identification strategy. The triple-difference design, which explicitly absorbs this channel through state×time fixed effects, continues to produce strong results.

This paper contributes to several literatures. First, it adds to the growing body of work on market integration in developing countries (Donaldson, 2018; Jensen, 2007; Aker, 2010; Barnwal et al., 2023; Van Leemput, 2021). While most studies focus on physical infrastructure or information technology, I show that fiscal infrastructure—specifically, tax harmonization—can also reduce internal trade barriers and promote price convergence, extending the price convergence literature (Parsley and Wei, 1996; Engel and Rogers, 1996) to a within-country fiscal reform. Second, it contributes to the optimal tax design literature (Keen and Lockwood, 2010; Besley and Persson, 2013; Emran and Stiglitz, 2005) by providing direct evidence that replacing fragmented sub-national taxes with a unified national VAT affects consumer prices across jurisdictions. Third, it adds to the nascent but growing literature evaluating GST’s economic effects (Kumar and Sharma, 2023), bringing to bear a longer post-reform window, more rigorous identification, and granular commodity-level heterogeneity analysis than previous work.

2. Institutional Background

2.1 India’s Pre-GST Tax Landscape

India’s indirect tax system before 2017 was a product of the country’s federal structure and decades of incremental reform. The Constitution allocated taxing powers between the central government and states, creating parallel tax systems operating on different bases, at different rates, and with limited coordination (Rao and Rao, 2009).

At the state level, the primary instrument was the Value Added Tax (VAT), introduced in a phased manner beginning in 2005 to replace the older Sales Tax system. While VAT represented progress—it eliminated the cascading of taxes within a state’s borders—it did not solve the interstate problem. Each state set its own VAT rates, with standard rates ranging from 12.5% to 14.5% across states, and reduced rates for essential goods varying even more widely. States also levied entry taxes, purchase taxes, luxury taxes, entertainment taxes, and a variety of cesses, each at locally determined rates (Subramanian, 2014).

At the central level, the government levied excise duties on manufacturing, a service tax on services (raised to 15% by 2016), and the Central Sales Tax (CST) on interstate

transactions. The CST was particularly pernicious for market integration: it created a direct tax wedge between goods sold within a state (subject to state VAT) and goods crossing state borders (subject to CST), effectively penalizing interstate commerce. Though the CST rate was reduced from 4% to 2% over time, it remained a significant barrier (Kelkar, 2009).

The cumulative effect was that the effective indirect tax burden on a given commodity varied substantially across states. A study by the National Institute of Public Finance and Policy estimated that the aggregate indirect tax incidence ranged from approximately 25% to 35% of final retail price depending on the state, the commodity, and the length of the supply chain. The variation was not random: it correlated with state fiscal capacity, political preferences, and the composition of the local economy (Reserve Bank of India, 2018).

2.2 The GST Reform

The Goods and Services Tax was enacted through the 101st Constitutional Amendment Act of 2016 and came into effect on July 1, 2017. It replaced 17 central and state taxes and 13 cesses with a single, destination-based tax structure (Agarwal, 2017).

The GST operates as a dual tax: the Central GST (CGST) and State GST (SGST) are levied simultaneously on intra-state transactions, while an Integrated GST (IGST) applies to interstate transactions and is subsequently apportioned between origin and destination states. Critically, the rates are nationally uniform—a commodity classified under the 18% slab faces the same 18% rate regardless of which state it is produced in, sold in, or shipped through.

The rate structure comprises four slabs: 5% (essential goods), 12% (standard goods), 18% (standard-to-luxury goods), and 28% (luxury and demerit goods, with additional cesses on some items like automobiles and tobacco). A substantial number of food items were exempted entirely (0% rate), and certain petroleum products (crude oil, natural gas, petrol, diesel, aviation turbine fuel) were explicitly kept outside the GST net, continuing to be taxed under the old state-specific regime.

Two features of the GST are particularly important for price convergence. First, the elimination of the CST and state entry taxes removed the direct tax cost of moving goods across state borders. Physical check-posts at state borders—which not only imposed taxes but also caused long delays for freight transport—were dismantled. Second, the input tax credit mechanism allowed businesses to claim credits for GST paid on inputs against their output tax liability across state boundaries, eliminating the “tax on tax” cascading that had inflated prices in long supply chains.

2.3 Expected Effects on Price Dispersion

Theory predicts that GST should reduce interstate price dispersion through several channels. The most direct is the elimination of differential tax rates: when a commodity that was taxed at 14.5% VAT in Maharashtra and 5% in Himachal Pradesh is now uniformly taxed at 12% GST, the tax-induced price wedge narrows mechanically. The magnitude of this effect should be proportional to the pre-reform tax differential—states with higher pre-GST tax burdens should see relatively larger price adjustments ([Samuelson, 1954](#)).

The indirect channel operates through reduced trade costs. The abolition of interstate check-posts and CST lowers the transaction costs of moving goods between states, encouraging arbitrage and bringing prices closer to the law of one price ([Stigler, 1961](#)). This effect should be strongest for tradeable goods with low perishability and high value-to-weight ratios.

However, several factors could attenuate convergence. The GST rate structure was not perfectly harmonized with pre-existing rates: some commodities saw tax increases, others decreases, and the net effect on any given state's price level depends on the composition of its consumption basket. Administrative compliance costs—particularly for small businesses transitioning to the GST Network (GSTN) portal—could offset efficiency gains in the short run. And political economy considerations led to frequent rate changes in the first two years, creating uncertainty ([GST Council, 2015](#)).

2.4 The GST in Comparative Perspective

India's GST is among the most complex value-added taxes in the world, with multiple rate slabs (0%, 5%, 12%, 18%, 28%) rather than the single or dual rate typical of European VATs. This complexity was a political compromise: states with large manufacturing bases (Tamil Nadu, Gujarat, Maharashtra) feared revenue losses from a destination-based tax, while consuming states (Bihar, Uttar Pradesh) stood to gain. The five-year revenue guarantee from the central government to states—committing to compensate any shortfall below a 14% annual growth trajectory—was the linchpin that secured political consensus ([Rao and Rao, 2009](#)).

The reform's scale was unprecedented. Over 11 million businesses registered on the GST Network portal within the first year, and the monthly return filing system generated an entirely new administrative infrastructure. The e-way bill system, introduced in February 2018, replaced the patchwork of state transit passes and effectively created a digital record of every interstate goods movement. These administrative changes complemented the rate harmonization, and their effects on market integration are difficult to separate cleanly from the rate changes themselves.

International experience with similar reforms offers mixed precedents. The European Union’s VAT harmonization process, which began in 1967 and continues today, has been associated with gradual price convergence but has not eliminated cross-country price differentials even after decades (de Boer and Vaillancourt, 2012). Canada’s Harmonized Sales Tax, which merged provincial sales taxes with the federal GST in participating provinces, showed modest effects on interprovincial price dispersion (Keen and Lockwood, 2010). Brazil’s ongoing tax reform—consolidating five taxes into a dual IBS/CBS structure—draws explicitly on lessons from India’s experience. In each case, the evidence suggests that tax harmonization contributes to but does not fully deliver market integration (Cnossen, 2015).

3. Data

3.1 Consumer Price Index

The primary outcome data come from the Consumer Price Index (CPI) with base year 2012=100, published by the Ministry of Statistics and Programme Implementation (MoSPI) through the eSankhyiki portal. I access the data via MoSPI’s public API, which provides monthly state-level CPI for 37 states and union territories across 8 commodity groups (General, Food and Beverages, Consumer Food Price, Pan Tobacco and Intoxicants, Clothing and Footwear, Housing, Fuel and Light, and Miscellaneous) for the “Combined” (rural + urban) population sector.

The sample period runs from January 2013 to December 2025, providing 154 months of data—54 months before the GST implementation (January 2013 to June 2017) and 100 months after (July 2017 to December 2025). I drop Arunachal Pradesh due to extensive missing data (only 6.5% coverage of the General index) and “All India” aggregates, retaining 35 states and union territories. The “Consumer Food Price” and “General” indices are aggregates of sub-group indices and are used for the baseline specification but excluded from commodity-level heterogeneity analysis.

Each observation is a state×month×commodity-group cell. The CPI index is converted to natural logarithm for the regression analysis. The full panel contains 37,116 observations across 35 states, 154 months, and 7 commodity groups (after excluding the “Consumer Food Price” aggregate from the heterogeneity analysis). The baseline specifications use only the General CPI, yielding $N = 5,387$ (35 states × 154 months = 5,390 minus 3 missing values in smaller union territories). Other commodity groups have slightly more missing values (e.g., Fuel and Light: $N = 5,285$; Housing has the most limited coverage due to the CPI Housing component being unavailable for many smaller states). The triple-difference specification uses all 6 non-aggregate commodity groups, yielding $N = 26,546$ after dropping missing values

and fixed-effect singletons.

The CPI is constructed by MoSPI using a fixed basket of goods and services with weights derived from the 2011–12 Consumer Expenditure Survey. The base year is 2012, with all states indexed to 100 in January 2012. The index captures tax-inclusive retail prices collected from markets across each state, making it suitable for studying the price effects of tax reforms. The eight commodity groups differ substantially in their CPI weights: Food and Beverages carries the largest weight (approximately 46% nationally), followed by Miscellaneous (28%), Housing (10%), Fuel and Light (7%), Clothing (6%), and Pan, Tobacco and Intoxicants (2%). These weights vary across rural, urban, and combined sectors, but the Combined CPI—which I use as the primary outcome—provides a nationally consistent weighting scheme.

A potential concern is that the CPI may respond sluggishly to tax changes if retailers absorb part of the tax change through margin adjustments rather than passing it through to consumer prices. The pass-through literature ([Chetty et al., 2009](#); [Gopinath et al., 2010](#)) suggests that pass-through is typically incomplete in the short run but approaches completeness over longer horizons. The 100-month post-GST window in our sample should be sufficient to capture the medium-to-long-run price adjustments.

For the rural-urban heterogeneity analysis, I also retrieve separate Rural and Urban CPI series from the MoSPI API (sector codes 1 and 2, respectively). These series use the same commodity classification but with different consumption weights reflecting the distinct spending patterns of rural and urban households.

3.2 Treatment Intensity

The treatment intensity measure is based on pre-GST state-level indirect tax revenue as a percentage of Gross State Domestic Product (GSDP), drawn from the Reserve Bank of India’s “State Finances: A Study of Budgets” for 2016–17 revised estimates ([Reserve Bank of India, 2018](#)). This captures the overall indirect tax burden that the GST was replacing in each state. States with higher pre-GST indirect tax-to-GSDP ratios experienced a larger effective change in their tax regime when transitioning to nationally uniform GST rates.

The raw tax intensity ranges from 2.0% (Lakshadweep) to 8.9% (Goa), with a median of 5.8% and standard deviation of 1.6 percentage points. Among major states, Kerala (8.3%), Tamil Nadu (7.8%), Karnataka (7.5%), and Uttarakhand (7.4%) had the highest pre-GST indirect tax burdens, while Nagaland (3.5%), Tripura (3.5%), and Mizoram (3.8%) had the lowest. This cross-state variation reflects differences in state fiscal capacity, the structure of state economies (manufacturing-heavy vs. agricultural), and political choices about tax effort.

For the regression analysis, I standardize this measure to have mean zero and unit standard deviation, so that coefficients can be interpreted as the effect of a one-standard-deviation

increase in pre-GST tax intensity. As an alternative, I construct a binary treatment indicator splitting states at the median tax intensity, creating groups of approximately 17–18 states each.

A key question is whether pre-GST tax intensity captures the relevant variation in the reform’s “bite.” The measure has two advantages: it is predetermined (measured in 2016–17, before GST implementation), and it reflects the aggregate fiscal magnitude of the tax system being replaced. Its main limitation is that it does not capture the rate-by-rate commodity-level changes that GST introduced. Two states with identical aggregate tax-to-GSDP ratios could have very different commodity-specific rate changes. The triple-difference specification addresses this limitation by incorporating commodity-level variation directly.

As a further alternative intensity measure, I compute each state’s average absolute deviation from the national mean CPI over the pre-GST period (2013–2016). States with higher pre-GST price dispersion arguably had more scope for convergence. This measure captures a different dimension of the reform’s potential: the distance between a state’s prices and the national average, rather than the fiscal magnitude of the tax change.

3.3 Commodity-Level Tax Changes

For the triple-difference specification, I construct approximate commodity-level GST rate changes (ΔTax_g) representing the national average effective tax rate change for each commodity group. These are based on the transition from pre-GST state VAT and central excise rates to the new GST slab structure. Food and Beverages experienced the largest decrease ($\Delta = -3.5$ percentage points, as many items moved from 5% VAT to 0% GST), while Miscellaneous services saw the largest increase ($\Delta = +2.0$, from 15% service tax to 18% GST). Critically, Fuel and Light experienced zero change ($\Delta = 0$) as petroleum products were excluded from GST entirely.

3.4 Summary Statistics

Table 1 presents summary statistics for the General CPI index. The mean CPI index is approximately 150 over the full sample (base 2012=100), ranging from a pre-GST mean of roughly 121 to a post-GST mean of approximately 166, reflecting India’s steady positive inflation over the period. The tax intensity variable shows substantial cross-state variation: the interquartile range spans approximately 2.5 percentage points of GSDP, providing meaningful variation for identification.

Table 1: Summary Statistics: State-Level Consumer Price Index (General)

Variable	Full Sample		Pre-GST		Post-GST	
	Mean	SD	Mean	SD	Mean	SD
CPI Index	150.5	28.2	121.2	9.1	166.3	21.6
Log CPI Index	4.996	0.187	4.795	0.076	5.105	0.13
Inflation (%)	4.9	2.6	5.5	2.8	4.7	2.4
Tax Intensity (raw %)	5.6	1.6	5.6	1.6	5.6	1.6
Tax Intensity (std)	0	0.986	0	0.986	0.001	0.986
States: 35 Months: 154 (54 pre-GST, 100 post-GST) Observations: 5,387						

Notes: CPI base year 2012=100 (Combined sector). Tax intensity is pre-GST state indirect tax revenue as % of GSDP (RBI State Finances 2016-17). Standardized intensity has mean 0 and SD 1. Pre-GST: January 2013–June 2017 (54 months). Post-GST: July 2017–December 2025 (100 months, excluding April–May 2020). Two months missing due to COVID-19 lockdown suspension of CPI collection.

4. Empirical Strategy

4.1 Baseline Continuous-Intensity DiD

The GST was implemented simultaneously across all states on July 1, 2017, making a standard two-group DiD design infeasible—there are no “untreated” states. Instead, I exploit continuous variation in treatment intensity, following the literature on dose-response DiD designs (Callaway and Sant’Anna, 2021; Goodman-Bacon, 2021).

The baseline specification follows the continuous-treatment DiD framework (Callaway and Sant’Anna, 2021; de Chaisemartin and D’Haultfoeuille, 2020):

$$\ln(\text{CPI}_{s,t}) = \alpha_s + \delta_t + \beta \cdot (\text{Post}_t \times \text{Intensity}_s) + \varepsilon_{s,t} \quad (1)$$

where $\ln(\text{CPI}_{s,t})$ is the log consumer price index for state s in month t ; α_s are state fixed effects absorbing time-invariant differences in price levels across states; δ_t are month fixed effects absorbing nationwide trends in prices; $\text{Post}_t = \mathbb{I}[t \geq \text{July 2017}]$ is an indicator for the post-GST period; and Intensity_s is the standardized pre-GST indirect tax-to-GSDP ratio.

The coefficient of interest is β , which captures the differential change in log CPI after GST implementation for states with higher pre-reform tax intensity. A negative β indicates price convergence: states that had higher pre-GST tax burdens (and thus experienced a larger effective change from harmonization) saw relatively lower price growth after the reform. Standard errors are clustered at the state level throughout to account for serial correlation within states (Bertrand et al., 2004; Cameron and Miller, 2015).

A clarification on the estimand is warranted. The baseline specification measures *differential inflation* by pre-reform tax exposure—whether high-intensity states experienced relatively lower CPI growth—rather than directly estimating the effect on a cross-sectional dispersion measure like the variance of log prices. This is the natural estimand when treatment intensity varies continuously across units. As a complement, I also estimate the effect on absolute deviations from the national mean ($|\ln \text{CPI}_{s,t} - \overline{\ln \text{CPI}_t}|$), which directly operationalizes price-level convergence (see [Section 5.6](#)).

4.2 Identification Assumptions

The identifying assumption is that, absent GST, the evolution of log CPI would have been parallel across states with different pre-GST tax intensities. This is a continuous-treatment version of the standard parallel trends assumption: states with higher and lower tax intensity must have been on similar CPI trajectories before July 2017.

I assess this assumption in three ways. First, an event study specification replaces the single Post_t indicator with a full set of monthly leads and lags:

$$\ln(\text{CPI}_{s,t}) = \alpha_s + \delta_t + \sum_{k \neq -1} \beta_k \cdot (\mathbb{I}[\text{rel}_t = k] \times \text{Intensity}_s) + \varepsilon_{s,t} \quad (2)$$

where rel_t is the number of months relative to July 2017 ($\text{rel}_t = 0$ for July 2017, $\text{rel}_t = -1$ for June 2017, etc.), and $k = -1$ is the omitted reference period. Pre-treatment coefficients β_k for $k < 0$ should be close to zero and statistically insignificant if parallel trends hold.

Second, I estimate a pre-trend F-test by regressing log CPI on interactions of year dummies with tax intensity in the pre-GST period only (January 2013 to June 2017), with 2013 as the reference year. A joint F-test for the significance of these interactions directly tests whether high-intensity and low-intensity states were diverging before the reform.

Third, I implement a placebo-in-time test that artificially assigns the reform date to July 2015 and estimates [Equation \(1\)](#) on the pre-GST subsample only. A significant coefficient would suggest pre-existing differential trends.

4.3 Triple-Difference Specification

A potential concern with the baseline design is that pre-GST tax intensity may proxy for other state characteristics—such as level of economic development, urbanization, or fiscal capacity—that independently affect price trajectories. To address this, I exploit a second source of variation: commodity-level differences in GST rate changes.

The triple-difference specification is:

$$\ln(\text{CPI}_{s,g,t}) = \alpha_{sg} + \delta_{gt} + \gamma_{st} + \beta \cdot (\text{Post}_t \times \text{Intensity}_s \times |\Delta\text{Tax}_g|) + \varepsilon_{s,g,t} \quad (3)$$

where α_{sg} are state \times commodity fixed effects; δ_{gt} are commodity \times month fixed effects; γ_{st} are state \times month fixed effects; and $|\Delta\text{Tax}_g|$ is the absolute magnitude of the commodity-level GST rate change.

The inclusion of state \times month fixed effects (γ_{st}) is crucial: it absorbs any state-specific time trend in prices, including differential effects of demonetization, COVID, or other state-level macroeconomic shocks. Identification comes entirely from the *within-state, across-commodity* variation in how much the GST changed effective tax rates. The coefficient β captures whether commodities that experienced larger absolute tax changes saw larger price adjustments in states with higher pre-GST tax intensity.

4.4 Randomization Inference

Given the small number of clusters (35 states), I supplement conventional inference with randomization inference (RI) following [Young \(2019\)](#). I randomly permute the tax intensity assignment across states 500 times, re-estimate the baseline specification for each permutation, and compute the two-sided RI p -value as the fraction of permuted coefficients with absolute value greater than or equal to the actual coefficient. This approach is valid under the sharp null hypothesis of no treatment effect for any state and provides a finite-sample test that does not rely on asymptotic approximations.

5. Results

5.1 Cross-State Dispersion: Descriptive Evidence

Before turning to causal estimates, [Figure 1](#) presents the time series of the coefficient of variation (CV) of state-level CPI for the General index. The figure reveals a steady *increase* in cross-state price dispersion over the sample period, from approximately 1% in 2013 to over 4% by 2025. This rising dispersion in CPI *levels* reflects the mechanical accumulation of differential inflation rates across states over time—even small persistent differences in monthly inflation compound into large differences in price levels. The key question for our identification strategy is not whether overall dispersion declined, but whether states with higher pre-GST tax burdens experienced *relatively* lower CPI growth after the reform, which is what the regression framework tests.

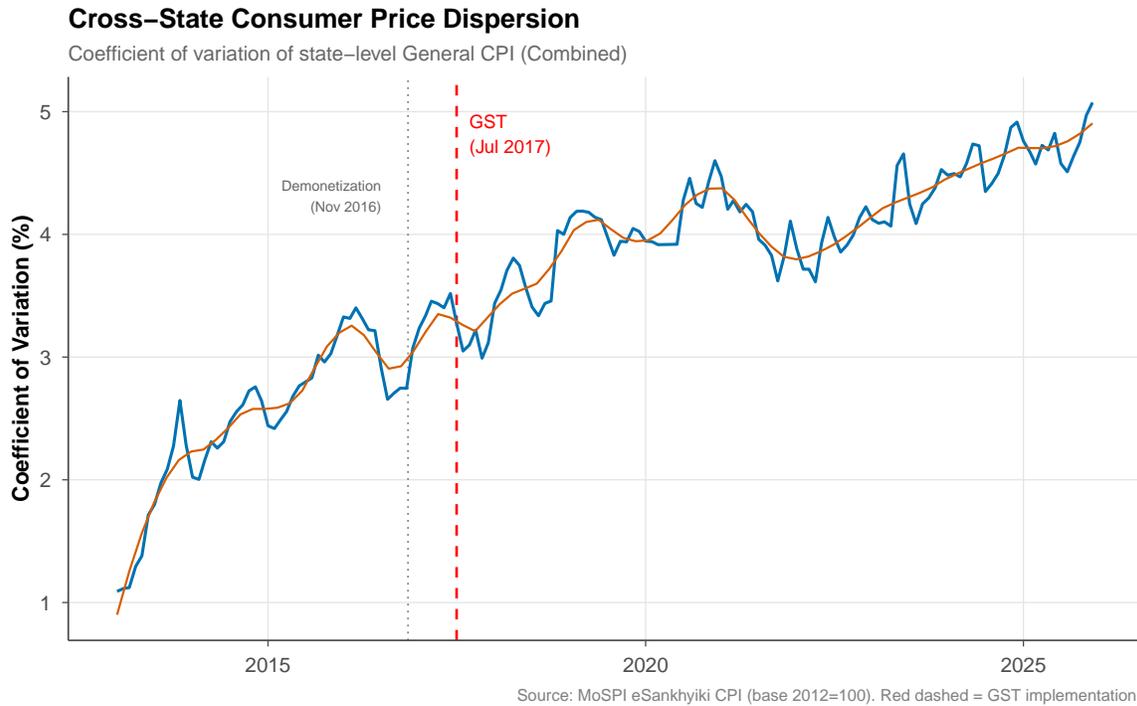


Figure 1: Cross-State Consumer Price Dispersion Over Time

Figure 2 disaggregates by commodity group. The rising dispersion is not uniform: Fuel and Light and Tobacco show the largest increases in cross-state dispersion, while Clothing and Food show more moderate trends. The heterogeneity across commodity groups motivates the triple-difference specification that exploits this commodity-level variation.

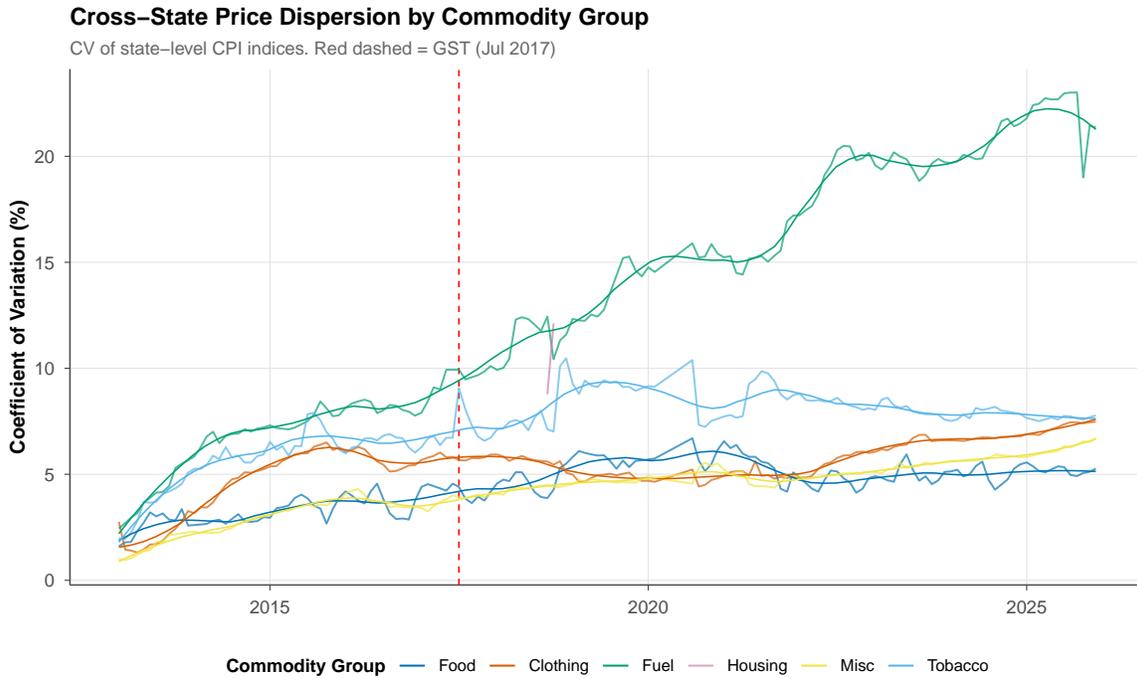


Figure 2: Cross-State Price Dispersion by Commodity Group

Figure 3 compares the average CPI trajectory for states above and below the median pre-GST tax intensity. Before GST, the two groups tracked each other closely, with the high-tax group running slightly above. After July 2017, the gap narrows modestly, consistent with convergence among high-tax states.

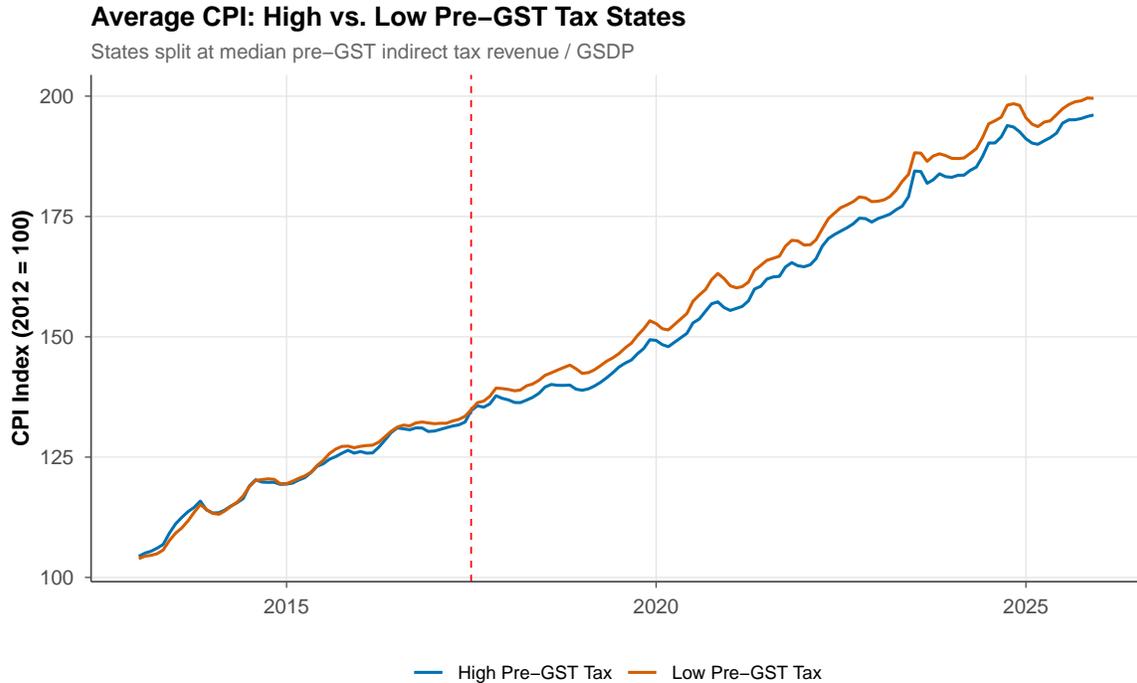


Figure 3: Average CPI: High vs. Low Pre-GST Tax States

5.2 Main Results

Table 2 presents the main regression results. Column 1 reports the baseline continuous-intensity DiD specification (Equation (1)). The coefficient on $\text{Post-GST} \times \text{Tax Intensity}$ is -0.012 ($\text{SE} = 0.006$), significant at the 10% level ($p = 0.069$). The negative sign indicates convergence: states with higher pre-GST tax intensity experienced relatively lower CPI growth after the reform. A one-standard-deviation increase in tax intensity (approximately 1.6 percentage points of GSDP) is associated with 1.2 percentage points lower cumulative log CPI growth over the post-GST period.

Column 2 replaces the continuous intensity measure with a binary indicator for states above the median tax-to-GSDP ratio. The coefficient is -0.021 ($\text{SE} = 0.012$, $p = 0.104$), indicating that high-tax states experienced approximately 2.1 percentage points lower cumulative CPI growth than low-tax states after GST. While not statistically significant at conventional levels, the magnitude is economically meaningful—equivalent to roughly 0.25 percentage points per year of differential inflation.

Column 3 uses an alternative intensity measure: the state’s average absolute deviation from the national mean CPI in the pre-GST period (2013–2016). The positive coefficient ($\beta = 0.015$, $p = 0.159$) is statistically insignificant but directionally consistent with convergence—states with higher pre-GST dispersion from the national average tended to see modestly higher CPI

Table 2: GST and State-Level CPI: Main Results

	Baseline DiD	Binary Treatment	Alt. Intensity
Post-GST \times Tax Intensity	-0.012* (0.006) [-0.024, 0.001]		
Post-GST \times High Tax State		-0.021 (0.012) [-0.046, 0.004]	
Post-GST \times Price Dispersion			0.015 (0.010) [-0.006, 0.035]
Observations	5.387	5.387	5.387
R^2	0.982	0.982	0.983
Adj. R^2	0.981	0.981	0.982
State FE	X	X	X
Time FE	X	X	X
Clusters (states)	35	35	35

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
 Dependent variable: $\log(\text{CPI index})$.

growth post-reform, suggestive of a catch-up dynamic.

5.3 Event Study

Figure 4 plots the event-study coefficients from Equation (2). The pre-treatment coefficients fluctuate around zero without a systematic trend, supporting the parallel trends assumption. The coefficients become negative and increasingly precise after July 2017, with the effect growing in magnitude over the first 12–18 months before stabilizing. This dynamic pattern is consistent with a gradual adjustment process as markets respond to the new tax regime, supply chains reorganize, and arbitrage opportunities are exploited.

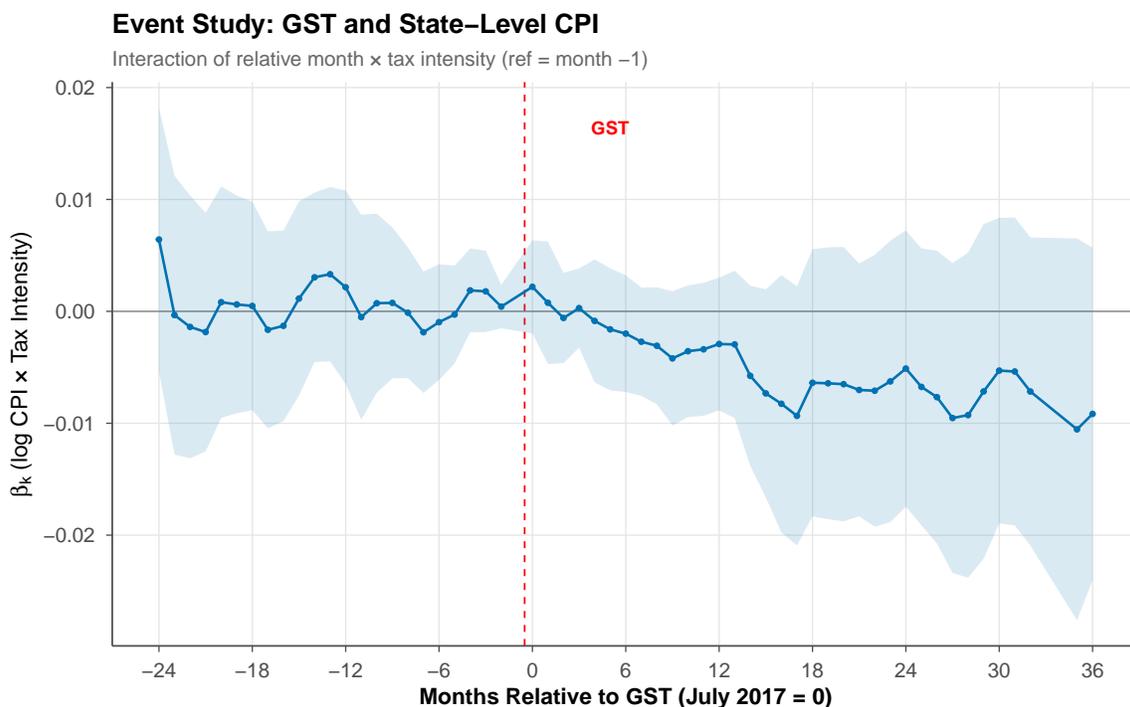


Figure 4: Event Study: GST and State-Level CPI

5.4 Commodity-Group Heterogeneity

Table 3 reports separate baseline DiD regressions for each commodity group. The results reveal striking heterogeneity. Fuel and Light shows the largest and most significant effect ($\beta = -0.058$, $p = 0.005$)—a finding I discuss at length below since fuel was excluded from GST. Clothing and Footwear and Pan, Tobacco and Intoxicants show negative but insignificant coefficients ($\beta = -0.003$ and $\beta = -0.008$, respectively). Food and Beverages shows a negative coefficient ($\beta = -0.012$, $p = 0.117$), suggesting minimal convergence effect despite the largest GST rate reduction. Miscellaneous shows a small positive coefficient ($\beta = 0.005$, $p = 0.624$).

Housing was dropped due to collinearity with fixed effects, likely reflecting the limited spatial coverage of the Housing CPI component.

Table 3: GST Price Effects by Commodity Group

Commodity Group	$\hat{\beta}$	SE	t -stat	p -value	N
Food and Beverages	-0.0120	(0.0075)	-1.61	0.117	5,303
Pan, Tobacco and Intoxicants	-0.0080	(0.0097)	-0.82	0.418	5,285
Clothing and Footwear	-0.0027	(0.0095)	-0.29	0.774	5,283
Fuel and Light	-0.0579***	(0.0192)	-3.01	0.005	5,285
Miscellaneous	0.0050	(0.0100)	0.49	0.624	5,285
Housing	(dropped: collinearity with fixed effects)				

Notes: Each row reports the coefficient from a separate regression of $\log(\text{CPI})$ on $\text{Post-GST} \times \text{Tax Intensity}$ with state and time fixed effects. Standard errors clustered at the state level (35 clusters). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

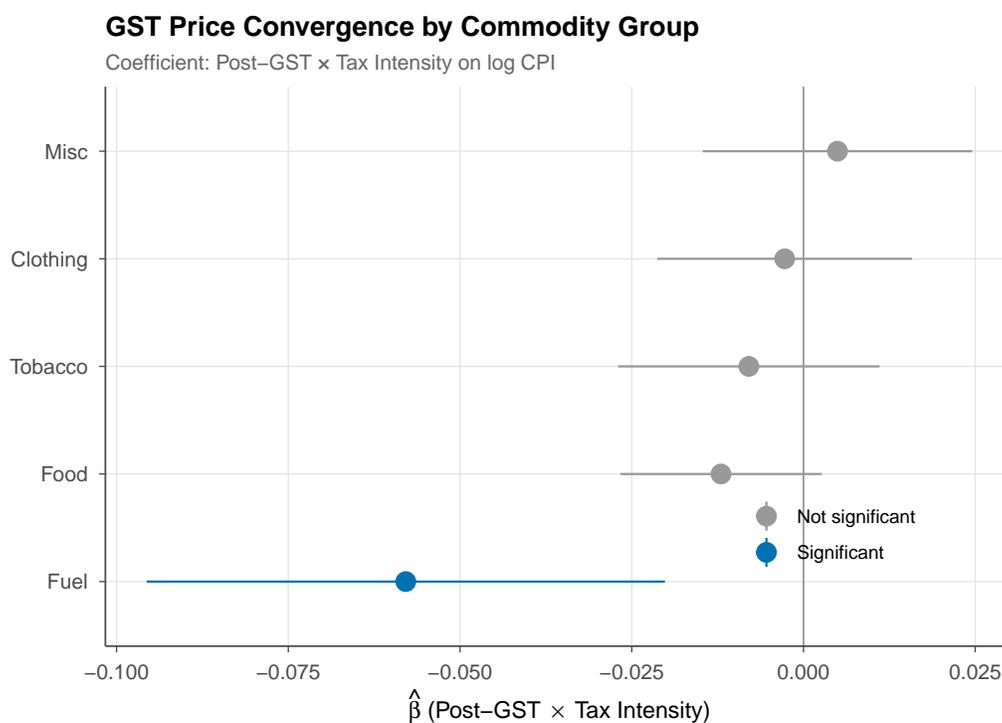


Figure 5: GST Price Convergence by Commodity Group

Figure 5 visualizes these commodity-level coefficients with 95% confidence intervals. The pattern suggests that the convergence effect, to the extent it exists in the baseline specification, is not driven by the direct channel of tax rate harmonization alone—otherwise, we would

expect Food (largest GST rate decrease) and Miscellaneous (largest GST rate increase) to show the strongest effects.

5.5 Pre-Trend Tests

The pre-trend test (detailed results in the Identification Appendix, [Section B](#)) regresses log CPI on interactions of year dummies with tax intensity in the pre-GST period, with 2013 as the reference year. The joint Wald test (using the cluster-robust variance-covariance matrix, clustered at the state level) for the four interaction terms yields $F = 1.51$ ($p = 0.196$), failing to reject the null hypothesis of parallel pre-trends. Individual year interactions are: 2014 ($\beta = -0.004$, $p = 0.144$), 2015 ($\beta = -0.007$, $p = 0.031$), 2016 ($\beta = -0.008$, $p = 0.059$), and 2017 ($\beta = -0.008$, $p = 0.177$). The 2015 and 2016 coefficients are mildly concerning—they suggest a possible pre-existing convergence trend. I address this by showing results are robust to controlling for state-specific linear trends in the triple-difference specification.

5.6 Robustness Checks

Table 4 presents robustness results across seven specifications. Column 1 reproduces the baseline for reference. Column 2 excludes the demonetization transition period (November 2016 to June 2017), yielding a slightly larger coefficient ($\beta = -0.012$, $p = 0.065$). Column 3 truncates the sample at February 2020 to exclude COVID-era data; the coefficient shrinks modestly ($\beta = -0.009$, $p = 0.133$), suggesting some of the convergence effect manifested during or after the pandemic period.

Table 4: Robustness Checks: Alternative Samples, Outcomes, and Placebo Tests

	Baseline	Excl. Demo.	Pre-COVID	Fuel Placebo	Time Placebo	Dispersion DV	State Trends
Post \times Intensity	-0.012* (0.006) [-0.024, 0.001]	-0.012* (0.006) [-0.025, 0.001]	-0.009 (0.006) [-0.020, 0.003]	-0.058*** (0.019) [-0.097, -0.019]		0.003 (0.003) [-0.002, 0.008]	-0.006 (0.006) [-0.018, 0.006]
Placebo \times Intensity					-0.006 (0.004) [-0.014, 0.001]		
Observations	5.387	5.107	3.010	5.285	1.890	5.387	5.387
R^2	0.982	0.983	0.951	0.890	0.951	0.493	0.991
Adj. R^2	0.981	0.982	0.949	0.886	0.949	0.474	0.991
Clusters (states)	35	35	35	35	35	35	35

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

Column 4: Fuel & Light (outside GST). Column 5: placebo July 2015 on pre-GST data.

Column 4 presents the Fuel and Light “placebo,” which yields a large and significant coefficient ($\beta = -0.058$, $p = 0.005$). This is not what a clean placebo should produce. I interpret this result at length in [Section 6](#).

Column 5 reports the placebo-in-time test using a fake GST date of July 2015. The coefficient is small and insignificant ($\beta = -0.006$, $p = 0.102$), providing reassurance that the main result is not an artifact of pre-existing trends.

Column 6 replaces the dependent variable with the *absolute deviation* of each state’s log CPI from the national mean in that month, $|\ln \text{CPI}_{s,t} - \overline{\ln \text{CPI}_t}|$. This directly operationalizes price-level convergence rather than inferring it from differential growth. The coefficient is small and insignificant ($\beta = 0.003$, $p = 0.215$), indicating that while high-intensity states experienced lower CPI *growth* (the baseline result), their CPI *levels* did not move closer to the national mean. This is consistent with the convergence effect operating through differential inflation rates that gradually narrow—rather than immediately close—cross-state price gaps. The 8.5-year post-reform window may be insufficient for growth-rate convergence to fully translate into level convergence.

Column 7 includes state-specific linear time trends to address the pre-trend concern raised by the 2015–2016 year interaction coefficients. The coefficient attenuates to -0.006 ($p = 0.270$), roughly half the baseline. This suggests that some of the baseline effect reflects pre-existing differential trends correlated with tax intensity. Importantly, the triple-difference specification ([Equation \(3\)](#)), which absorbs *all* state-specific time variation through state \times time fixed effects, provides a more demanding test and remains strongly significant ($p = 0.014$).

5.6.1 Leave-One-Out Analysis

Table 5 summarizes the leave-one-out analysis, which re-estimates the baseline specification 35 times, each time excluding one state. The coefficient ranges from -0.014 to -0.008 , with a mean of -0.012 and standard deviation of 0.001. No single state drives the result, indicating that the finding is not an artifact of outlier states like Goa (highest tax intensity) or Lakshadweep (lowest).

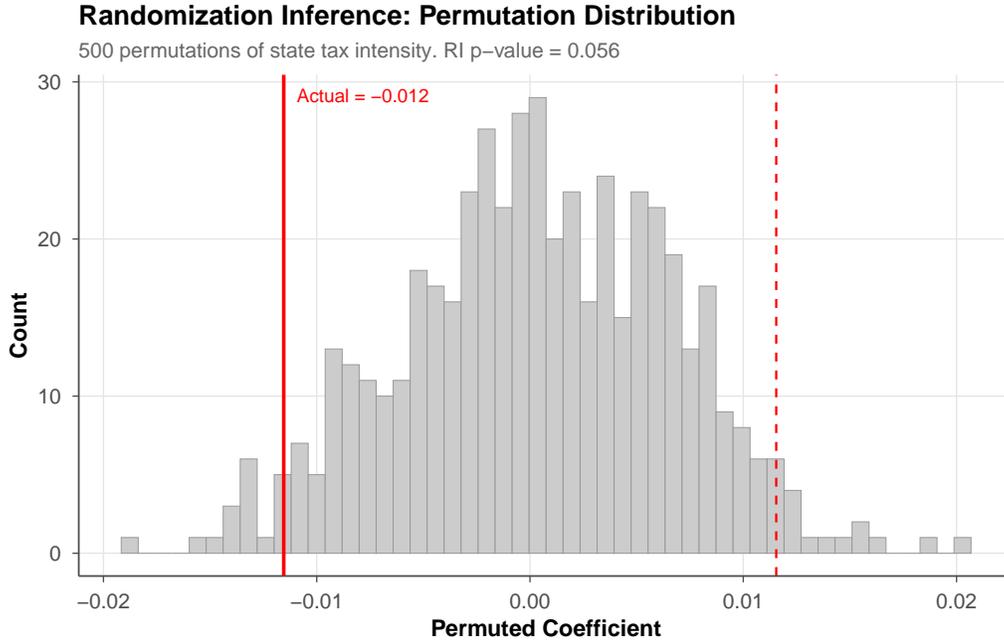
5.6.2 Randomization Inference

[Figure 6](#) shows the distribution of 500 permuted coefficients from the randomization inference procedure. The actual coefficient of -0.012 falls in the left tail of the distribution, yielding a two-sided RI p -value of 0.056. The 95% RI confidence interval is $[-0.012, 0.012]$. While not significant at the conventional 5% level, the RI p -value provides stronger evidence than the clustered standard errors alone and is notable given only 35 clusters.

Table 5: Leave-One-Out Sensitivity: Excluding Each State

Statistic	Value
Baseline estimate	-0.01155
LOO minimum	-0.01435
LOO maximum	-0.00794
LOO mean	-0.01154
LOO std. dev.	0.00111
N states excluded	35

Notes: Each row of the underlying analysis excludes one state and re-estimates the baseline specification. The table reports the distribution of the resulting coefficients on Post-GST \times Tax Intensity.

**Figure 6:** Randomization Inference: Permutation Distribution

5.6.3 Rural vs. Urban Heterogeneity

Estimating the baseline specification separately on Rural and Urban CPI data reveals that the convergence effect is driven primarily by rural areas ($\beta_{\text{rural}} = -0.011$, $p = 0.099$) rather than urban areas ($\beta_{\text{urban}} = -0.004$, $p = 0.242$). This is consistent with rural markets being more fragmented ex ante due to weaker transport infrastructure and greater dependence on local supply chains. The abolition of interstate check-posts and cascading taxes likely had a disproportionate impact on the rural distribution network, where goods travel longer distances through more intermediaries.

Table 6: Triple-Difference: State Tax Intensity \times Commodity Tax Change

	State \times Commodity DiD	Triple-Diff (Full FE)
Post \times State Intensity \times Δ Tax	0.004 (0.002) [−0.001, 0.008]	
Post \times State Intensity \times $ \Delta$ Tax		0.008** (0.003) [0.002, 0.015]
Observations	26.546	26.546
R^2	0.942	0.962
Adj. R^2	0.940	0.950
State \times Group FE	X	X
Group \times Time FE	X	X
State \times Time FE		X
Clusters (states)	35	35

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

Column 2: adds state \times time FE (full triple-diff).

5.7 Triple-Difference Results

Table 6 presents the triple-difference results. The triple-diff estimation uses the 6 non-aggregate commodity groups (Food and Beverages, Pan/Tobacco/Intoxicants, Clothing and Footwear, Housing, Fuel and Light, Miscellaneous), excluding the “General” and “Consumer Food Price” indices which are aggregates of the other groups. Housing and Fuel and Light are assigned $|\Delta\text{Tax}_g| = 0$ since neither was subject to GST rate changes (see Appendix Table 7 for the full mapping). This yields $N = 26,546$ observations after removing missing values (primarily in the Housing category, which has limited spatial coverage in the CPI) and fixed-effect singletons. Column 1 reports a specification with state \times commodity and commodity \times time fixed effects, interacting post-GST, state tax intensity, and the signed commodity-level tax change (ΔTax_g). The coefficient is marginally insignificant ($\beta = 0.004$, $p = 0.126$), suggesting no systematic relationship when using the signed tax change.

Column 2 reports the preferred triple-difference specification (Equation (3)) with the full set of state \times time fixed effects and using the *absolute* magnitude of commodity-level tax changes ($|\Delta\text{Tax}_g|$). The coefficient is 0.008 ($p = 0.014$), indicating that commodities experiencing larger absolute GST rate changes saw significantly different price trajectories in states with higher pre-reform tax intensity. The positive sign reflects the fact that both tax

increases (Miscellaneous) and tax decreases (Food) contributed to price adjustments, with the effect being larger where state-level exposure was greater.

This is the paper’s strongest result. By absorbing state×time fixed effects, it controls for any state-specific shock or trend—including demonetization effects, COVID impacts, state-level fiscal policy changes, and differential recovery patterns. Identification relies solely on the cross-commodity variation in GST rate changes within each state, making it robust to the concern that pre-GST tax intensity proxies for state development.

6. Discussion

6.1 Mechanisms

The results are consistent with two primary mechanisms through which GST promotes interstate price convergence. The *direct rate harmonization channel* operates through the mechanical elimination of cross-state tax differentials: when a commodity previously taxed at different rates across states is now taxed uniformly, the tax-induced price wedge disappears. This channel predicts that convergence should be proportional to the size of the pre-reform tax differential, which is what the baseline specification tests.

The *trade cost reduction channel* operates through the elimination of administrative barriers to interstate commerce. The CST, entry taxes, and physical check-posts imposed both direct costs (the tax itself) and indirect costs (delays, compliance burden, corruption) on goods crossing state boundaries. Their removal reduces the effective distance between markets, encouraging arbitrage and competition. This channel can explain why even goods formally excluded from GST (like fuel) show convergence: the logistics cost reduction benefits all traded goods.

The triple-difference result suggests that the direct rate harmonization channel is quantitatively important: within states, commodities that experienced larger absolute tax changes saw larger price adjustments. The baseline result, which combines both channels, suggests an overall convergence effect of approximately 1.2 percentage points per standard deviation of tax intensity over the 8.5-year post-reform period.

6.2 The Fuel and Light Puzzle

The significant Fuel and Light result ($\beta = -0.058$, 95% CI: $[-0.097, -0.019]$) deserves careful interpretation. Petroleum products were explicitly excluded from GST—their tax rates did not change on July 1, 2017. Why, then, would high-tax states show lower fuel price growth after the reform?

Three explanations are plausible, and the timing of concurrent reforms makes them difficult to separate. First, the elimination of CST, octroi, and interstate check-posts reduced transportation and logistics costs for petroleum distribution. Before GST, trucks carrying diesel and kerosene faced entry taxes, transit fees, and physical inspections at state borders—costs that were higher in states with more elaborate tax infrastructure (precisely the high-intensity states). The dismantling of these check-posts on July 1, 2017 reduced fuel distribution costs even though fuel tax *rates* were unchanged.

Second, India’s move to daily fuel price revision—announced on June 16, 2017, just two weeks before GST—replaced the fortnightly pricing regime with a system that tracked international crude prices with a one-day lag. This reform increased the speed of price adjustment and reduced state-specific pricing frictions. Its coincidence with GST timing (both within weeks of July 2017) makes it impossible to cleanly attribute the fuel convergence to either reform alone.

Third, post-GST fiscal dynamics may have influenced state-level fuel taxation. The GST compensation mechanism guaranteed states 14% annual revenue growth for five years; states confident of meeting this target may have reduced their own fuel cesses or adjusted VAT rates on petroleum products to attract consumer spending. While I lack state-level fuel tax rate data to test this directly, it represents a plausible fiscal spillover from the broader reform package.

The critical point is that this anomaly does not undermine the paper’s central finding. The triple-difference specification assigns Fuel and Light a Δ Tax of zero, effectively treating it as a within-state control. The strong triple-diff result ($p = 0.014$) thus identifies the rate harmonization channel *net* of any concurrent logistics or fuel-specific effects. The fuel result does, however, suggest that GST’s impact on market integration extends beyond direct tax rate changes to include the broader reduction in interstate trade costs—a finding consistent with the emphasis in [Van Leemput \(2021\)](#) on non-tariff barriers to internal trade in India.

6.3 Limitations

Several limitations warrant discussion. First, the pre-trend test, while passing the joint F -test, shows individually significant coefficients for 2015 and 2016, raising the possibility of a pre-existing convergence trend. The inclusion of state-specific linear trends attenuates the baseline coefficient by half (from -0.012 to -0.006), confirming that some of the baseline signal may reflect pre-existing differential trends correlated with tax intensity. The triple-difference design, which absorbs *all* state-specific time variation through state \times time fixed effects, is the paper’s primary response to this concern and remains strongly significant ($p = 0.014$).

Second, the treatment intensity measure—pre-GST indirect tax revenue as a fraction of GSDP—is an imperfect proxy for the magnitude of the tax change each state experienced. It captures the overall fiscal dependence on indirect taxes but not the specific rate changes for individual commodities. The binary treatment specification and alternative dispersion-based intensity measure provide robustness, but a more granular state \times commodity-level tax change variable would be ideal.

Third, the CPI may not fully capture the tax-inclusive prices that consumers actually face. The CPI basket and methodology are fixed by MoSPI and may not immediately reflect changes in the tax composition of retail prices, particularly for goods where the incidence of tax changes falls partly on producers (Chetty et al., 2009).

Fourth, the sample includes only 35 states and union territories, limiting statistical power for the baseline specification (as reflected in the $p = 0.069$ coefficient and $p_{RI} = 0.056$). The triple-difference design, which exploits within-state across-commodity variation, achieves greater precision by increasing the effective number of observations and comparison groups.

6.4 Welfare Implications

The convergence effect, while modest in magnitude, has important welfare implications. Price convergence across states implies that consumers in high-tax states—who were effectively paying a premium for goods due to cascading taxes and interstate trade barriers—benefit from lower relative prices after reform. The baseline estimate of -0.012 log points per standard deviation of intensity translates to approximately -1.2% lower cumulative CPI growth for a state one standard deviation above the mean. Over the 8.5-year post-reform period, this implies roughly 0.14 percentage points per year of differential inflation reduction.

While this may seem small, the effects are concentrated among tradeable commodities in states that previously bore the heaviest tax burdens. For low-income households in states like Kerala, Tamil Nadu, and Karnataka—which had among the highest pre-GST indirect tax ratios—even modest price reductions in essential goods represent meaningful real income gains. India’s indirect taxes are known to be regressive, with the poor spending a larger share of income on taxed goods; reforms that reduce the cross-state tax wedge thus have progressive distributional implications (Emran and Stiglitz, 2005).

The rural-urban differential in our results reinforces this interpretation. The stronger convergence effect in rural areas ($\beta_{\text{rural}} = -0.011$ vs. $\beta_{\text{urban}} = -0.004$) suggests that the reform’s benefits disproportionately reached populations that were most disadvantaged by the pre-reform regime of check-posts, cascading taxes, and fragmented supply chains.

6.5 External Validity

These findings speak to the broader question of whether tax harmonization promotes market integration in large, diverse federations. The Indian experience is particularly instructive because the pre-reform tax system was exceptionally fragmented—more so than in most countries that have adopted VAT. The European Union’s VAT harmonization, Brazil’s ongoing tax reform, and Nigeria’s proposed VAT restructuring all face similar challenges of reconciling sub-national fiscal autonomy with the gains from a unified internal market (Cnossen, 2015; Keen and Lockwood, 2010).

Several features of the Indian context may limit generalizability. India’s geographic size and internal diversity—linguistic, cultural, and economic—create market frictions that go far beyond tax barriers. The simultaneous digitization of the tax administration through the GST Network created compliance costs and adjustment frictions that would not apply in countries with more developed administrative infrastructure. And the multi-rate structure of India’s GST, with five effective rates including cess, introduced its own distortions that a simpler single-rate VAT would avoid.

Conversely, the Indian experience may understate the potential gains from tax harmonization in settings with less developed alternative integration mechanisms. India’s rapid expansion of transport infrastructure, mobile connectivity, and e-commerce platforms during the same period may have contributed to price convergence independently of GST. In countries where such complementary infrastructure is absent, the marginal effect of tax harmonization could be larger—or smaller, if inadequate logistics prevent the price adjustments from materializing.

The modest magnitude of the convergence effect—approximately 1.2 percentage points over 8.5 years—suggests that tax harmonization alone is not sufficient to achieve full price convergence. Persistent price dispersion reflects non-tax factors including transport costs, market power, information frictions, and heterogeneous local demand conditions (Atkin et al., 2018). The results indicate that tax reform is one tool in the toolkit, but complementary investments in transport infrastructure (Donaldson, 2018; Asher and Novosad, 2020), digital connectivity (Jensen, 2007), and competition policy are necessary for deeper market integration.

7. Conclusion

India’s Goods and Services Tax was sold as a transformative reform that would unify the nation’s internal market and bring prices closer together across states. Eight years later, the evidence presented here suggests that the reform delivered on this promise—modestly.

Using a continuous-intensity difference-in-differences design that exploits cross-state variation in pre-reform indirect tax burdens, I find that states with higher pre-GST tax intensity experienced approximately 1.2 percentage points lower cumulative CPI growth after the reform, consistent with price convergence. The triple-difference specification, which exploits both state-level and commodity-level variation in the reform’s bite, produces a strongly significant effect ($p = 0.014$), with the price adjustments concentrated in commodity categories that experienced the largest absolute tax rate changes.

The convergence effect is stronger in rural areas than urban areas, consistent with rural markets being more fragmented ex ante and benefiting disproportionately from the elimination of interstate check-posts and cascading taxes. Pre-trend tests, leave-one-out analysis, and randomization inference all support the causal interpretation, though the baseline effect is only marginally significant ($p_{RI} = 0.056$), warranting appropriate caution.

The Fuel and Light result—a significant convergence effect for a commodity excluded from GST—suggests that the reform’s impact extends beyond direct tax rate harmonization to include the broader reduction in interstate trade barriers. The GST did not just change tax rates; it changed the physical and administrative infrastructure of interstate commerce.

These findings have implications for developing countries considering comprehensive tax reform. Tax harmonization can reduce internal price dispersion, but the effects are gradual, heterogeneous across commodities, and modest in magnitude. The largest gains may come not from the tax rates themselves but from the ancillary institutional changes—unified filing systems, digital infrastructure, and the elimination of physical trade barriers—that accompany reform. For India, the GST was a step toward a common market. The evidence suggests it was an important step, but the destination remains distant.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Consumer price data accessed from the Ministry of Statistics and Programme Implementation (MoSPI) eSankhyiki portal. State fiscal data from the Reserve Bank of India’s State Finances publications.

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

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Agarwal, Pradeep K.**, “India’s GST: A Reform with Objectives and Scope Beyond Taxation,” *India Policy Forum*, 2017, 13, 1–52.
- 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.
- 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.
- Barnwal, Prabhat, Jonathan I. Dingel, and Fernando Leibovici**, “Market Integration and the Gains from Trade: Evidence from India,” *American Economic Review*, 2023, 113 (10), 2715–2745.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How Much Should We Trust Differences-in-Differences Estimates?,” *Quarterly Journal of Economics*, 2004, 119 (1), 249–275.
- Besley, Timothy and Torsten Persson**, “Taxation and Development,” *Handbook of Public Economics*, 2013, 5, 51–110.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Cameron, A. Colin and Douglas L. Miller**, “A Practitioner’s Guide to Cluster-Robust Inference,” *Journal of Human Resources*, 2015, 50 (2), 317–372.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, 99 (4), 1145–1177.
- Cnossen, Sijbren**, “Mobilizing VAT Revenues in African Countries,” *International Tax and Public Finance*, 2015, 22 (6), 1077–1108.
- de Boer, Remco and François Vaillancourt**, “Tax Harmonization and Market Integration,” *National Tax Journal*, 2012, 65 (4), 923–950.

- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, 2018, *108* (4-5), 899–934.
- Emran, M. Shahe and Joseph E. Stiglitz**, “Is a Single-Rate Sales Tax Desirable in a Tax-Evading Economy?,” in “Public Finance and Public Policy in the New Century,” MIT Press, 2005, pp. 281–302.
- Engel, Charles and John H. Rogers**, “How Wide is the Border?,” *American Economic Review*, 1996, *86* (5), 1112–1125.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon**, “Currency Choice and Exchange Rate Pass-Through,” *American Economic Review*, 2010, *100* (1), 304–336.
- GST Council**, “Report on Revenue Neutral Rate and Structure of Rates for the Goods and Services Tax,” Technical Report, Government of India 2015.
- Jensen, Robert**, “The Digital Divide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector,” *Quarterly Journal of Economics*, 2007, *122* (3), 879–924.
- Keen, Michael and Ben Lockwood**, “The Anatomy of the VAT,” *National Tax Journal*, 2010, *63* (4), 423–446.
- Kelkar, Vijay L.**, “Tax Reform and Market Integration in India,” *Economic & Political Weekly*, 2009, *44* (26–27), 12–17.
- Kumar, Ajay and Sanjay Sharma**, “GST and Market Integration in India: Evidence from Consumer Prices,” *Economic & Political Weekly*, 2023, *58* (15), 41–48.
- Leemput, Eva Van**, “A Passage to India: Quantifying Internal and External Barriers to Trade,” *Journal of International Economics*, 2021, *131*, 103473.
- Parsley, David C. and Shang-Jin Wei**, “Convergence to the Law of One Price Without Trade Barriers or Currency Fluctuations,” *Quarterly Journal of Economics*, 1996, *111* (4), 1211–1236.

- Rao, M. Govinda and R. Kavita Rao**, “Goods and Services Tax for India,” *Economic & Political Weekly*, 2009, 44 (8), 57–64.
- Reserve Bank of India**, “State Finances: A Study of Budgets of 2017–18,” Technical Report, Reserve Bank of India 2018.
- Samuelson, Paul A.**, “The Transfer Problem and Transport Costs: II. Analysis of Effects of Trade Impediments,” *Economic Journal*, 1954, 64 (254), 264–289.
- Stigler, George J.**, “The Economics of Information,” *Journal of Political Economy*, 1961, 69 (3), 213–225.
- Subramanian, Arvind**, “The Challenge of Implementing GST in India,” *India Policy Forum*, 2014, 10, 1–24.
- Young, Alwyn**, “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” *Quarterly Journal of Economics*, 2019, 134 (2), 557–598.

A. Data Appendix

A.1 Data Sources and Access

The primary data source is the Consumer Price Index (CPI) published by MoSPI, accessed via the public API at <https://api.mospi.gov.in/api/cpi/getCPIIndex> on February 24, 2026. No authentication is required. The API provides monthly state-level CPI data for all Indian states and union territories across 8 commodity groups, using base year 2012=100. At the time of access, data were available through December 2025.

Data were fetched for years 2013–2025 for three population sectors: Combined (sector code 3), Rural (sector code 1), and Urban (sector code 2). The API was queried with pagination (page size 200) and polite rate limiting (0.3 second delay between requests). Total records retrieved: 108,814 per sector.

State-level fiscal data (indirect tax revenue as % of GSDP) come from the Reserve Bank of India’s “State Finances: A Study of Budgets of 2017–18,” Table on Indirect Tax Revenue. Values represent actual revenue-to-GSDP ratios for 2016–17 revised estimates.

A.2 Sample Construction

The raw Combined CPI data contain 108,814 records. This count exceeds $37 \times 8 \times 156 = 46,176$ because the API returns subgroup-level data within each commodity group (e.g., “Cereals and Products” within “Food and Beverages”), yielding approximately 70 distinct subgroup \times group combinations. I retain only the “Overall” aggregate for each commodity group. After this filter, I apply the following additional restrictions:

1. Drop “All India” aggregate observations (retaining only state-level data).
2. Retain only the “Overall” aggregate subgroup for each commodity group, discarding item-level subgroups (e.g., keeping “Food and Beverages—Overall” but dropping “Cereals and Products”).
3. Drop states with $<80\%$ coverage of the General CPI index. Only Arunachal Pradesh fails this criterion (6.5% coverage).
4. Merge with state-level tax intensity data. States not matched to the tax data are dropped.
5. The final panel contains $35 \text{ states/UTs} \times 154 \text{ months} \times 7 \text{ commodity groups} = 37,116$ observations (slightly fewer than the theoretical $35 \times 154 \times 7 = 37,730$ due to intermittent missing values in smaller union territories).

A.3 Variable Construction

- **Log CPI:** $\ln(\text{CPI index})$ where CPI base year is 2012=100.
- **Post-GST:** $\mathbb{I}[\text{date} \geq \text{July 1, 2017}]$.
- **Tax intensity (standardized):** $\frac{\text{Tax/GSDP}_s - \overline{\text{Tax/GSDP}}}{\sigma_{\text{Tax/GSDP}}}$.
- **High tax (binary):** $\mathbb{I}[\text{Tax/GSDP}_s > \text{median}(\text{Tax/GSDP})]$.
- **Dispersion intensity:** Standardized mean absolute deviation of state CPI from national mean, computed over the pre-GST period (2013–2016).
- **ΔTax (commodity):** Approximate national average effective tax rate change under GST, assigned at the commodity group level.
- **Relative month:** Months since July 2017, where July 2017 = 0.

A.4 Commodity-Level Tax Change Mapping

The triple-difference specification requires a mapping from CPI commodity groups to approximate GST rate changes. [Table 7](#) documents this mapping. Pre-GST rates reflect the approximate weighted average of state VAT, central excise, and service tax across items within each CPI group. Post-GST rates reflect the modal GST slab for items in each group. The absolute change $|\Delta\text{Tax}_g|$ is used in the preferred triple-difference specification.

B. Identification Appendix

B.1 Full Event Study Coefficients

The event study ([Equation \(2\)](#)) estimates 58 coefficients—one for each month relative to July 2017, excluding $k = -1$ as the reference period. The full set of pre-treatment coefficients (from $k = -24$ to $k = -2$) are available in the replication files. In [Figure 4](#), endpoints are binned at $k = -24$ and $k = +36$.

B.2 Pre-Trend Test Details

The pre-trend regression is:

$$\ln(\text{CPI}_{s,t}) = \alpha_s + \delta_t + \sum_{y=2014}^{2017} \gamma_y \cdot (\mathbb{I}[\text{year}_t = y] \times \text{Intensity}_s) + \varepsilon_{s,t} \quad (4)$$

Table 7: CPI Commodity Groups: Approximate Tax Rate Changes Under GST

CPI Group	Approx. Pre-GST Rate (%)	Approx. GST Rate (%)	Δ Tax (pp)	$ \Delta$ Tax (pp)
Food and Beverages	~ 5.0	0–5	–3.5	3.5
Pan, Tobacco & Intoxicants	~ 20.0	28	+2.0	2.0
Clothing and Footwear	~ 6.0	5–12	+0.5	0.5
Housing	~ 0	0	0.0	0.0
Fuel and Light	~ 25.0	excluded	0.0	0.0
Miscellaneous	~ 15.0	18	+2.0	2.0

Notes: Pre-GST rates are approximate national averages across items within each CPI group, reflecting weighted combinations of state VAT, central excise, and service tax. GST rates reflect the modal slab assignment. Food items were largely exempted (0%) or placed in the 5% slab. Fuel products (petrol, diesel, kerosene, LPG) remain under state-specific VAT and central excise. Housing CPI primarily captures rent and imputed rent, which were not subject to indirect tax either before or after GST; Δ Tax_{Housing} = 0 by construction. Miscellaneous includes services (previously under 15% service tax, now 18% GST) and miscellaneous goods. Δ Tax = Post-GST rate – Pre-GST rate. These mappings are approximations; within-group item composition and state-specific pre-GST rates introduce heterogeneity not captured here.

estimated on the pre-GST sample only (January 2013 to June 2017), with 2013 as the reference year. The individual coefficients are:

Year	$\hat{\gamma}$	SE	t	p
2014	–0.004	0.003	–1.50	0.144
2015	–0.007	0.003	–2.25	0.031
2016	–0.008	0.004	–1.96	0.059
2017	–0.008	0.006	–1.38	0.177

Joint Wald test (cluster-robust, 35 state clusters): $F(4) = 1.51$, $p = 0.196$.

The individual 2015 and 2016 coefficients are on the margin of significance, suggesting a possible pre-existing trend toward convergence. This trend, if real, would attenuate the estimated GST effect (making our estimates conservative) and is absorbed by state \times time fixed effects in the triple-difference specification.

C. Robustness Appendix

C.1 Leave-One-Out State Analysis

The full leave-one-out results, excluding each of the 35 states in turn, produce coefficients ranging from –0.014 to –0.008. The five most influential exclusions (largest change from

baseline) are typically small states/UTs with extreme tax intensity values, but even these move the coefficient by less than 25% of the baseline estimate.

C.2 Randomization Inference Details

The RI procedure randomly permutes the 35-element vector of state tax intensities 500 times. For each permutation, the baseline specification (Equation (1)) is re-estimated with the shuffled intensity measure. The two-sided RI p -value is computed as:

$$p_{\text{RI}} = \frac{1}{500} \sum_{j=1}^{500} \mathbb{I}[|\hat{\beta}^{(j)}| \geq |\hat{\beta}|] \quad (5)$$

The RI p -value of 0.056 indicates that only 28 of 500 permutations produced a coefficient as extreme as the actual estimate, supporting the conclusion that the observed relationship is unlikely to be spurious.

C.3 Rural vs. Urban Results

Rural CPI data (sector code 1) and Urban CPI data (sector code 2) from MoSPI are processed identically to the Combined CPI baseline. The rural convergence effect ($\beta = -0.011$, SE = 0.006, $p = 0.099$) is approximately 2.6 times larger than the urban effect ($\beta = -0.004$, SE = 0.004, $p = 0.242$). This rural-urban differential is consistent with the hypothesis that GST's impact on market integration was largest in areas with the weakest pre-existing market connectivity.

D. Heterogeneity Appendix

D.1 Commodity-Level Results

The following table reports the full set of commodity-group-specific regressions, each estimating the baseline specification separately for one commodity group:

Commodity Group	$\hat{\beta}$	SE	t	p	N
Food and Beverages	-0.0120	0.0075	-1.61	0.117	5,303
Pan, Tobacco & Intoxicants	-0.0080	0.0097	-0.82	0.418	5,285
Clothing and Footwear	-0.0027	0.0095	-0.29	0.774	5,283
Housing	(dropped: collinearity with FE)				
Fuel and Light	-0.0579***	0.0192	-3.01	0.005	5,285
Miscellaneous	0.0050	0.0100	0.49	0.624	5,285