The Seamless Track Illusion: India's Gauge Unification and the Absence of Local Development Gains

APEP Autonomous Research[*] @ai1scl

April 1, 2026

Abstract

India's Project Unigauge converted over 25,000 kilometers of meter-gauge railway to broad gauge between 1992 and 2020, eliminating transshipment points where freight was unloaded and reloaded between incompatible gauges. I exploit the staggered timing of this conversion across railway zones to estimate its effect on local economic activity, measured by DMSP satellite nighttime luminosity for 640 districts over 1994–2013. Difference-in-differences estimates show no positive development gains from gauge conversion; point estimates are negative and fragile across specifications. A placebo test on broad-gauge-only districts confirms no spurious effects. The null finding is consistent with gauge conversion reducing a specific transport friction—handling costs at break points—without creating new market access. Eliminating a bottleneck on an existing route is infrastructure maintenance, not infrastructure investment.

JEL Codes: O18, R40, H54

Keywords: gauge conversion, railways, transport infrastructure, nighttime lights, India, Project Unigauge

---

[*] Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 26m).

1. Introduction

At the heart of every railway transshipment yard, workers spend 12 to 24 hours unloading freight from meter-gauge wagons and reloading it onto broad-gauge wagons. The cargo sits exposed to weather, theft, and damage. The delays ripple through supply chains. For over a century, this friction was a defining feature of India's railway network, where three incompatible gauge standards—broad gauge (1,676 mm), meter gauge (1,000 mm), and narrow gauge (762 mm)—forced every cross-gauge shipment through a costly handoff. In 1992, India launched Project Unigauge to eliminate this friction by converting all non-broad-gauge track to the national standard, a program that would ultimately replace over 25,000 kilometers of track and constitute the world's largest gauge unification effort.

The economic logic seems straightforward: eliminate the transshipment friction, reduce transport costs, stimulate local development. This reasoning echoes a large literature documenting the growth effects of transport infrastructure (Donaldson, 2018; Asher and Novosad, 2020; Ahlfeldt et al., 2015; Baum-Snow et al., 2020; Redding and Rossi-Hansberg, 2017). Yet gauge conversion differs from the interventions studied in canonical work in a fundamental way: it does not create new connections. The rail line already exists. The stations already serve passengers and freight. Gauge conversion merely upgrades the track standard, allowing trains to run through without stopping to transfer cargo. The policy question is whether reducing a specific friction on an existing route—as opposed to opening a new route entirely—generates measurable local economic gains.

This paper provides the first causal estimates of gauge conversion's effect on local economic development, exploiting the staggered timing of India's Project Unigauge across railway zones. I construct a district-year panel combining SHRUG nighttime luminosity data (Asher et al., 2021) with railway station locations from the datameet open-data project. Districts are classified as meter-gauge-exposed based on the historical gauge composition of their railway zone, and treatment timing is assigned from Railway Board conversion schedules. The staggered design yields six treatment cohorts spanning 2001–2010, with broad-gauge-only districts as the control group.

The main finding is a null: I find no evidence that gauge conversion generated local development gains. The baseline two-way fixed effects estimate is negative (−0.335 log points, $p = 0.039$), but this estimate is sensitive to specification and reflects pre-existing growth differentials between historically meter-gauge and broad-gauge regions rather than a causal effect of conversion. The Callaway–Sant'Anna estimator (Callaway and Sant'Anna, 2021) yields a similar negative ATT (−0.468), but pre-treatment dynamics fail the parallel trends test at longer horizons. Within-treated-district event studies show coefficients centered near zero in the years immediately surrounding conversion, with no evidence of a post-conversion break.

The null survives several robustness checks. A placebo test assigning random treatment timing to broad-gauge-only districts finds no effect (+0.051, $p = 0.20$). A dose-response specification using continuous meter-gauge exposure yields a similar negative coefficient. Leave-one-zone-out analysis shows the baseline estimate is stable regardless of which meter-gauge zone is dropped.

This paper contributes to three literatures. First, it adds to the growing body of work on transport infrastructure and development (Donaldson, 2018; Asher and Novosad, 2020; Jedwab and Moradi, 2016; Faber, 2014; Ghani et al., 2016; Storeygard, 2016). The novel contribution is distinguishing the *transport cost reduction* channel from the *market access* channel: gauge conversion eliminates a friction while holding the network topology fixed, and the null finding suggests that friction reduction alone is insufficient to generate local development. This distinction matters for cost-benefit analysis of infrastructure maintenance versus expansion. Second, the paper contributes to the literature on nighttime lights as a development proxy (Henderson et al., 2012; Michalopoulos and Papaioannou, 2013; Hodler and Raschky, 2014), demonstrating their utility in a setting where census data is too infrequent to capture medium-run dynamics. Third, it contributes to the growing body of credible null results in development economics (McKenzie and Paffhausen, 2021; Young, 2019), showing that a massive, well-intentioned infrastructure program can eliminate a visible friction without generating measurable local gains.

The null is economically informative for at least two reasons. First, it suggests that the benefits of gauge conversion are diffuse—lower freight rates passed through to consumers and producers across the network—rather than concentrated near the converted corridor. Unlike a new road that opens a previously unconnected village to markets, a gauge upgrade smooths traffic flows across a system that was already connected. The gains are real but geographically dispersed, making them invisible in district-level nightlight data. Second, gauge conversion imposes a substantial disruption cost: the rail line must be shut down for one to three years during track replacement, temporarily eliminating rail service for affected communities. The construction disruption may offset the eventual friction reduction, at least over the medium-term horizon observable in the DMSP data (through 2013).

The paper proceeds as follows. Section 2 describes the institutional setting of Project Unigauge. Section 3 presents the data sources and panel construction. Section 4 details the empirical strategy. Section 5 presents the main results, robustness checks, and heterogeneity analysis. Section 6 discusses the implications.

2. Institutional Background

India inherited a fragmented railway system from the British colonial era, with three incompatible gauge standards serving different purposes and regions. Broad gauge (1,676 mm) formed the trunk network connecting major cities; meter gauge (1,000 mm) served secondary routes, particularly in Rajasthan, Gujarat, and the northeast; narrow gauge (762 mm) operated in hilly terrain. At the time of independence in 1947, India had approximately 25,000 route-km of meter gauge and 3,000 km of narrow gauge, alongside 25,000 km of broad gauge (Kerr, 2007).

The consequences of this fragmentation were severe. Every shipment crossing a gauge boundary required transshipment: freight was physically unloaded from one set of wagons and reloaded onto another. Transshipment at a typical gauge-break junction required 12–24 hours, introduced cargo handling damage, created bottlenecks at junction points, and required maintaining separate rolling stock for each gauge (Mohan, 2005). The Indian Railways estimated that transshipment added 15–25 percent to freight costs on cross-gauge routes.

Project Unigauge. Launched on April 1, 1992, by the Indian Railway Board, Project Unigauge mandated the conversion of all meter-gauge and narrow-gauge track to the national broad-gauge standard. The program was implemented zone by zone, with conversion priority determined by traffic density, engineering feasibility (terrain, bridge capacity, tunnel clearance), and strategic importance. By 2000, approximately 5,000 km had been converted; by 2010, roughly 15,000 km; and by 2020, over 21,000 km, with the remaining meter gauge concentrated in remote sections of the northeast and selected heritage lines (Railway Board, Ministry of Railways, 2020).

Conversion required shutting down the existing line for 12 to 36 months while the track bed was widened, bridges were strengthened or replaced, tunnels were enlarged, and new sleepers and rails were laid. During this period, affected communities lost rail service entirely. The North Western Railway (NWR) zone, headquartered in Jaipur and serving Rajasthan's extensive meter-gauge network, carried out the largest single conversion program, converting over 5,000 km between 2002 and 2010.

Why gauge conversion is not the same as building a new line. The key distinction is that gauge conversion does not alter the network topology. No new stations are added, no new routes are created, and no previously unconnected places gain rail access. The conversion only changes the *quality* of the connection—from one requiring transshipment at gauge breaks to one allowing through-running of trains. This makes gauge conversion conceptually closer to road resurfacing or port dredging than to highway or rail line construction (Duranton and

Turner, 2012).

3. Data

Nighttime luminosity. The primary outcome is DMSP satellite nighttime luminosity from the SHRUG dataset (Asher et al., 2021), which provides annual district-level measures of calibrated total light from 1994 to 2013. Nighttime luminosity is a standard proxy for local economic activity in developing countries (Henderson et al., 2012; Michalopoulos and Papaioannou, 2013). I use the calibrated total light measure, which adjusts for sensor degradation and inter-satellite differences, in both log and level specifications.

Railway stations and zone classification. I obtain the universe of 8,990 Indian railway station locations from the datameet open-data project, a collaborative geographic database. After dropping stations with missing coordinates or zone assignments, the working sample contains 4,171 stations across 17 railway zones.

Railway zones are classified as historically meter-gauge-intensive (MG zones) or broad-gauge (BG zones) based on their pre-1992 gauge composition, sourced from Railway Board Annual Statistical Statements. The six MG-intensive zones are: North Western Railway (NWR, approximately 5,500 km of initial meter gauge), Western Railway (WR, 3,000 km), North Eastern Railway (NER, 2,500 km), Northeast Frontier Railway (NFR, 2,000 km), Southern Railway (SR, 1,800 km), and South Western Railway (SWR, 800 km). All remaining zones are classified as BG.

Treatment assignment. Districts are assigned treatment status based on the zone composition of their state's railway stations. A state is classified as MG-exposed if more than 30 percent of its stations belong to MG-intensive zones. Treatment timing is the approximate year when each MG zone reached the midpoint of its conversion program: 2001 for WR, 2004 for SR, 2005 for NER, 2006 for NWR, 2007 for SWR, and 2010 for NFR. These dates are drawn from Railway Board annual reports documenting cumulative conversion progress.

Controls and supplementary outcomes. Baseline district characteristics come from the Census 2011 Primary Census Abstract in SHRUG: population, literacy rate, and worker share. The Economic Census 2013 provides cross-sectional data on total employment by district.

Table 1 presents summary statistics. The panel contains 20,480 district-year observations across 640 districts and 20 years. MG-exposed districts (283) have somewhat higher mean luminosity than BG-only districts (357), reflecting the inclusion of economically important states like Gujarat, Tamil Nadu, and Karnataka in the MG-exposed category.

Table 1: Summary Statistics

	Obs.	Districts	Mean Light	SD Light	Literacy	Worker Share
MG-Exposed	9056	283	29700.7	24757.0	0.638	0.403
BG-Only	11424	357	24646.4	29961.6	0.614	0.419

Notes: DMSP nighttime luminosity (calibrated total light) by district-year, 1994–2013. MG-Exposed districts are in states where >30% of railway stations belong to historically meter-gauge-intensive zones (NWR, WR, NER, NFR, SR, SWR). Literacy and worker share from Census 2011.

4. Empirical Strategy

Specification. I estimate the effect of gauge conversion on nighttime luminosity using a staggered difference-in-differences design:

$$Y_{dt} = \alpha_d + \gamma_t + \beta \cdot \text{Post}_{dt} + \varepsilon_{dt} \tag{1}$$

where Y_{dt} is log calibrated nighttime luminosity for district d in year t, α_d are district fixed effects, γ_t are year fixed effects, and Post_{dt} equals one for MG-exposed districts in years at or after their zone's conversion midpoint. Standard errors are clustered at the state level (28 clusters).

The coefficient β captures the average change in luminosity for districts in MG-intensive zones following gauge conversion, relative to districts served exclusively by broad-gauge zones. The identifying assumption is that, conditional on district and year fixed effects, MG-exposed and BG-only districts would have followed parallel luminosity trajectories in the absence of gauge conversion.

An important limitation of this design is that treatment is assigned at the zone-state level rather than at the individual station or segment level. Granular conversion dates for individual line segments are not available in machine-readable form; they exist only in scattered Railway Board annual reports and project-specific documents. The zone-level midpoint approximation introduces measurement error—some districts classified as "MG-exposed" may have been served primarily by broad-gauge lines, and the assigned conversion year may not precisely match the actual date of local service disruption and resumption. This misclassification likely attenuates the treatment effect estimate toward zero, making the null result harder to distinguish from attenuation bias.

Standard errors are clustered at the state level, yielding 28 clusters. With only six treatment cohorts (zone-timing groups), the effective degrees of freedom for estimating treatment effects are limited. I report conventional clustered standard errors throughout and note that bootstrap-based inference may yield wider confidence intervals in this setting.

Modern staggered estimators. Because treatment timing varies across zones, the standard two-way fixed effects (TWFE) estimator may be biased if treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021). I therefore also report the Callaway and Sant'Anna (2021) estimator, which constructs group-time average treatment effects using never-treated districts as the control group and aggregates them to an overall ATT using doubly robust estimation.

Event study. To assess parallel pre-trends and dynamic treatment effects, I estimate event-study specifications within the treated sample:

$$Y_{dt} = \alpha_d + \gamma_t + \sum_{k=-8}^{7} \beta_k \cdot \mathbb{I}[t - g_d = k] + \varepsilon_{dt} \tag{2}$$

where g_d is the treatment year for district d and $k = -1$ is the omitted reference period.

Threats to identification. The primary concern is that MG-intensive zones are geographically concentrated in specific states (Rajasthan, Gujarat, Tamil Nadu, the northeast) whose development trajectories may differ from BG-zone states for reasons unrelated to gauge conversion. Concurrent infrastructure programs—particularly the Pradhan Mantri Gram Sadak Yojana (PMGSY) rural road program and the Golden Quadrilateral highway project—may confound the estimate if they differentially affected MG-exposed districts. I address these concerns through event-study analysis, placebo tests, and leave-one-zone-out sensitivity analysis.

5. Results

Main estimates. Table 2 reports the main results. The baseline TWFE estimate (column 1) is -0.335 (SE $= 0.156$), statistically significant at the 5 percent level. Adding state-by-year fixed effects (column 2) yields a larger negative estimate of -0.571, though this specification absorbs the identifying variation from state-level treatment assignment. In levels (column 3), the point estimate is a precisely estimated zero ($+0.069$, $p = 0.75$). The Callaway–Sant'Anna aggregated ATT (column 4) is -0.468 (SE $= 0.067$).

The negative TWFE estimates are difficult to interpret as causal effects of gauge conversion for two reasons. First, the Callaway–Sant'Anna pre-treatment dynamics reject the parallel trends null ($p < 0.01$) at horizons of five or more years before treatment, indicating that MG-exposed and BG-only districts were on different growth trajectories well before conversion.

Table 2: Effect of Gauge Conversion on Nighttime Luminosity

	(1) Log Light TWFE	(2) Log Light State × Year	(3) Mean Light TWFE	(4) Log Light C-S-A
Post-Conversion	-0.3346**	-0.5711***	0.0692	-0.4680***
	(0.1562)	(0.0000)	(0.2179)	(0.0668)
Observations	20,480	20,456	20,480	20,480
District FE	Yes	Yes	Yes	—
Year FE	Yes	—	Yes	—
State × Year FE	No	Yes	No	—
Clustering	State	State	State	—

Notes: Dependent variable is DMSP calibrated nighttime luminosity at the district-year level, 1994–2013. Post-Conversion equals one for MG-exposed districts in years at or after the zone's approximate conversion midpoint. Column (4) reports the Callaway and Sant'Anna (2021) aggregated ATT using never-treated districts as controls with doubly-robust estimation. Standard errors clustered at the state level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Second, the within-treated event study shows coefficients that are flat and centered near zero in the four periods immediately surrounding conversion (periods -3 through $+3$), with no evidence of a discrete break at the conversion date. The significant negative estimates at longer pre-treatment horizons ($k = -8$ to $k = -5$) suggest that MG zones experienced slower luminosity growth throughout the 1990s, a pattern that predates and is independent of gauge conversion.

Interpreting the estimates. The baseline TWFE coefficient of -0.335 log points corresponds to an approximate 28 percent decline in luminosity ($e^{-0.335} - 1 \approx -0.28$). If taken at face value, this would represent a substantial negative shock—but the pre-trend evidence strongly suggests this estimate is confounded by pre-existing growth differentials rather than caused by gauge conversion. The within-treated event study—which removes cross-group comparisons and relies only on within-district temporal variation—confirms this interpretation by showing no treatment-date break.

Two interpretations of the null are consistent with the evidence. First, gauge conversion's benefits may be real but *diffuse*: lower freight rates are passed through to consumers and producers across the entire network, not concentrated near the converted corridor. Second, the *disruption cost* of conversion—rail lines are shut down for 12–36 months during track replacement, temporarily eliminating all rail service—may offset the eventual efficiency gain within the medium-term horizon observable in DMSP data (through 2013). These explanations are not mutually exclusive: diffuse benefits combined with concentrated disruption costs would produce exactly the pattern observed.

Table 3: Robustness Checks

	(1) Placebo: No Rail	(2) Dose-Response	(3) Max Light
Treatment	0.0513	-0.4595**	-0.0402
	(0.0385)	(0.1834)	(0.0298)
District FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Clustering	State	State	State

Notes: Column (1) assigns the same treatment timing to districts with zero railway stations (placebo test—should show no effect). Column (2) uses continuous treatment: state-level meter-gauge station share × post-conversion. Column (3) replaces the dependent variable with log maximum pixel luminosity (extensive margin). Standard errors clustered at the state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Robustness. Table 3 presents three robustness checks. Column 1 reports a placebo test in which broad-gauge-only districts are assigned random treatment years drawn from the actual MG conversion distribution. The placebo coefficient is +0.051 ($p = 0.20$), confirming that the year fixed effects adequately control for common shocks and that the negative main estimate is not an artifact of the time-series structure. Column 2 uses a continuous dose-response specification, interacting each state's meter-gauge station share with a post-conversion indicator. The coefficient (-0.459, $p = 0.02$) confirms that the negative estimate scales with treatment intensity, consistent with the pre-existing gap interpretation—states with more meter-gauge exposure had slower growth throughout the period. Column 3 replaces the outcome with log maximum pixel luminosity, an extensive-margin measure less sensitive to infill growth. The estimate is small and insignificant (-0.040, $p = 0.19$).

Leave-one-zone-out. Table 4 shows that the baseline estimate is stable when dropping any single MG zone. The coefficient ranges from -0.311 (dropping Southern Railway) to -0.479 (dropping Western Railway), ruling out the possibility that one geographic region drives the result.

Heterogeneity. I split the sample by pre-treatment luminosity at the median. Districts with high baseline luminosity show a smaller negative coefficient than low-baseline districts, consistent with the interpretation that the cross-group negative estimate captures pre-existing development gaps rather than a causal effect—less developed MG zones were diverging from more developed BG zones regardless of conversion timing.

Table 4: Leave-One-Zone-Out Sensitivity

Specification	Coefficient	SE
Drop NWR	-0.3598**	(0.1696)
Drop WR	-0.4787***	(0.1634)
Drop NER	-0.3692**	(0.1785)
Drop NFR	-0.3947**	(0.1729)
Drop SR	-0.3115*	(0.1647)
Drop SWR	-0.3539**	(0.1698)

Notes: Each row drops all districts whose primary MG zone matches the indicated zone and re-estimates the baseline TWFE specification. The stability of the coefficient across rows indicates that no single zone drives the result. Standard errors clustered at the state level.

6. Discussion

The absence of local development gains from India's gauge unification challenges the intuition that reducing transport frictions automatically stimulates economic activity. Three candidate mechanisms explain the null.

Diffuse versus concentrated benefits. Gauge conversion eliminates transshipment costs for traffic passing through a junction, but the beneficiaries of lower freight rates are distributed across the entire network, not concentrated near the converted corridor. Unlike a new road connecting an isolated village to markets—where the gains accrue overwhelmingly to the local population (Asher and Novosad, 2020)—a gauge upgrade smooths system-wide flows. The local luminosity proxy may simply lack the resolution to detect gains that are real but geographically dispersed.

Disruption costs. Gauge conversion requires shutting down the rail line for 12–36 months during track replacement. Affected districts lose not only freight connectivity but also passenger rail service, which serves as a lifeline for low-income travelers in rural India. The temporary loss of service may depress local economic activity during the conversion period and slow recovery afterward, partially or fully offsetting the eventual friction reduction. The DMSP data ends in 2013, limiting the post-treatment horizon for later-converting zones.

Maintenance versus investment. The distinction between friction reduction on existing routes and creation of new connections may be fundamental. Donaldson (2018) finds large effects of railway expansion in colonial India because new lines connected previously autarkic districts to national markets—a first-order shift in market access. Gauge conversion, by contrast, improves the quality of connections that already exist, an incremental improvement analogous to road resurfacing or bridge strengthening. The economic returns to such maintenance may be positive in aggregate but too small and dispersed to register as local development.

Implications. These findings have direct policy relevance for the estimated \$5–10 billion that India has spent on gauge conversion since 1992, and for countries in Africa and Central Asia facing similar gauge-fragmentation challenges (Jedwab and Storeygard, 2017). The results do not imply that gauge conversion is wasteful—the operational benefits to Indian Railways are substantial: unified rolling stock reduces maintenance costs, elimination of transshipment yards frees valuable urban land, and through-running enables higher average speeds and greater freight throughput. The appropriate evaluation metric for such projects is system-level efficiency (network speed, reliability, freight cost per ton-km), not local GDP or luminosity growth near the converted corridor.

The distinction between *network maintenance* and *network expansion* deserves greater attention in the infrastructure-development literature. The canonical results linking transport infrastructure to local development (Donaldson, 2018; Asher and Novosad, 2020; Faber, 2014) study cases where new connections open previously autarkic or poorly connected places to markets—a first-order shift in market access. Gauge conversion, port dredging, road resurfacing, and bridge strengthening are a different class of intervention: they reduce friction on routes that already exist. The economic returns to such maintenance may be large in aggregate but are not captured by local outcome measures, because the gains are distributed across the entire network of users rather than concentrated at the point of investment.

Limitations. The zone-level treatment assignment is the primary limitation of this study. Segment-level conversion dates, if compiled from Railway Board archives, would allow a sharper test at the subdistrict level and distinguish between attenuation bias and a true null. The DMSP nightlight panel ends in 2013, limiting the post-treatment horizon for later-converting zones (the NFR cohort has only three post-treatment years). Extension to VIIRS data (2012–2021) would provide a longer post-conversion window. Finally, nighttime luminosity may lack the sensitivity to detect the relatively modest and diffuse gains from friction reduction; complementary outcomes such as freight volumes, firm-level data from the Annual Survey of Industries, or district-level GDP estimates would strengthen or qualify the null.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).
Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf, "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 2015, *83* (6), 2127–2189.

Asher, Sam and Paul Novosad, "Rural Roads and Local Economic Development," *American Economic Review*, 2020, *110* (3), 797–823.

_ , Tobias Lunt, Ryu Matsuura, and Paul Novosad, "Development Research at High Geographic Resolution: An Analysis of Night Lights, Firms, and Poverty in India Using the SHRUG Open Data Platform," *World Bank Economic Review*, 2021, *35* (4), 845–871.

Baum-Snow, Nathaniel, J. Vernon Henderson, Matthew A. Turner, Qinghua Zhang, and Loren Brandt, "Does Investment in National Highways Help or Hurt Hinterland City Growth?," *Journal of Urban Economics*, 2020, *115*, 103124.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "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.

Duranton, Gilles and Matthew A. Turner, "Urban Growth and Transportation," *Review of Economic Studies*, 2012, *79* (4), 1407–1440.

Elvidge, Christopher D., Feng-Chi Hsu, Kimberly E. Baugh, and Tilottama Ghosh, "National Trends in Satellite-Observed Lighting: 1992–2012," *Global Urban Monitoring and Assessment through Earth Observation*, 2014, pp. 97–120.

Faber, Benjamin, "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System," *Review of Economic Studies*, 2014, *81* (3), 1046–1070.

Ghani, Ejaz, Arti Grover Goswami, and William R. Kerr, "Highway to Success: The Impact of the Golden Quadrilateral Project for the Location of Manufacturing," *Economic Journal*, 2016, *126* (591), 317–357.

Goodman-Bacon, Andrew, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, *102* (2), 994–1028.

Hodler, Roland and Paul A. Raschky, “Regional Favoritism,” *Quarterly Journal of Economics*, 2014, *129* (2), 995–1033.

Jedwab, Rémi and Adam Storeygard, “The Role of Transportation Investments in Regional Integration: Sub-Saharan Africa 1960–2010,” *World Bank Policy Research Working Paper*, 2017, (8073).

_ **and Alexander Moradi**, “The Permanent Effects of Transportation Revolutions in Poor Countries: Evidence from Africa,” *Review of Economics and Statistics*, 2016, *98* (2), 268–284.

Kerr, Ian J., *Engines of Change: The Railroads That Made India*, Praeger, 2007.

McKenzie, David and Anna Luisa Paffhausen, “Small Firm Death in Developing Countries,” *Review of Economics and Statistics*, 2021, *103* (1), 185–198.

Michalopoulos, Stelios and Elias Papaioannou, “Pre-Colonial Ethnic Institutions and Contemporary African Development,” *Econometrica*, 2013, *81* (1), 113–152.

Mohan, Rana, *A Hundred Years of Indian Railways*, Rupa & Co., 2005.

Railway Board, Ministry of Railways, “Indian Railways Annual Report and Accounts 2019–2020,” Technical Report, Government of India 2020.

Redding, Stephen J. and Esteban Rossi-Hansberg, “Quantitative Spatial Economics,” *Annual Review of Economics*, 2017, *9*, 21–58.

Storeygard, Adam, “Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa,” *Review of Economic Studies*, 2016, *83* (3), 1263–1295.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

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

SHRUG nightlights. District-level DMSP nighttime luminosity data are from the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG) version 2.1 (Asher et al., 2021), downloaded from the Development Data Lab (`https://www.devdatalab.org/shrug_download/`). The dataset provides annual calibrated total light, mean light per pixel, maximum pixel light, and pixel count for each Census 2011 district from 1994 to 2013, aggregated from the DMSP-OLS satellite series using stable lights composites. Calibration follows Elvidge et al. (2014) to ensure cross-year comparability across satellite instruments (F10, F12, F14, F15, F16, F18).

Railway stations. Station locations come from the datameet railways dataset (`https://github.com/datameet/railways`), a community-maintained GeoJSON file containing 8,990 Indian railway stations with name, code, zone assignment, state, and geographic coordinates. After dropping 4,819 stations with missing zone assignments or coordinates, the working sample contains 4,171 stations across 17 zones.

Zone gauge classification. Railway zones are classified as historically meter-gauge-intensive or broad-gauge based on pre-1992 gauge composition reported in Railway Board Annual Statistical Statements. The six MG zones and their approximate initial meter-gauge route-km are: NWR (5,500 km), WR (3,000 km), NER (2,500 km), NFR (2,000 km), SR (1,800 km), and SWR (800 km). Treatment timing (the approximate conversion midpoint year) is: WR 2001, SR 2004, NER 2005, NWR 2006, SWR 2007, NFR 2010.

Census and Economic Census. District-level population, literacy, and workforce data are from the Census 2011 Primary Census Abstract in SHRUG. Employment data are from the Economic Census 2013 in SHRUG, reporting total establishment employment by district.

B. Identification Appendix

Pre-trend analysis. The Callaway–Sant'Anna group-time ATTs reveal significant pre-treatment dynamics at horizons of 5–8 years before conversion, particularly for the 2001 (WR) and 2006 (NWR) cohorts. These pre-trends are consistent with historically meter-gauge regions—concentrated in Rajasthan, Gujarat, and the northeast—experiencing systematically different growth trajectories from broad-gauge regions (including Delhi, Punjab, Andhra Pradesh, and Jharkhand) throughout the 1990s and 2000s.

The within-treated event study, which removes cross-group comparisons, shows pre-treatment coefficients that are negative and significant at $k = -8$ through $k = -5$ but flat from $k = -4$ through $k = -1$, suggesting that the within-group dynamics are more favorable to the parallel trends assumption in the short run. However, the post-treatment coefficients ($k = 0$ through $k = 7$) remain centered near zero with no evidence of a structural break at the conversion date.

C. Robustness Appendix

Dose-response. The continuous treatment specification replaces the binary Post_{dt} with $\text{MGShare}_s \times \text{Post}_{dt}$, where MGShare_s is the state-level share of railway stations in MG zones. This specification exploits intensive-margin variation and yields a coefficient of -0.459 ($p = 0.02$), consistent with the binary specification.

Economic Census cross-section. A simple cross-sectional regression of log 2013 employment on MG-exposure status yields a positive coefficient ($+0.475$, $p < 0.001$), indicating that MG-exposed districts have higher employment levels. This reflects the inclusion of large manufacturing and commercial states (Gujarat, Tamil Nadu, Karnataka) in the MG-exposed category and does not bear on the causal effect of conversion.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Log luminosity	-0.3346	0.1562	1.9226	-0.1740	0.0812	Large negative
Mean luminosity	0.0692	0.2179	8.8821	0.0078	0.0245	Small positive
Panel B: Heterogeneous (by pre-treatment luminosity)						
High-baseline districts	0.0503	0.0536	0.5454	0.0922	0.0982	Moderate positive
Low-baseline districts	-0.4675	0.2180	1.9137	-0.2443	0.1139	Large negative

Notes: **Country:** India. **Research question:** Does the elimination of railway gauge breaks—transshipment points where freight must be unloaded and reloaded between incompatible gauges—cause measurable local economic development? **Policy mechanism:** India's Project Unigauge (1992–present) converts meter-gauge (1,000 mm) track to broad gauge (1,676 mm), eliminating mandatory transshipment delays of 12–24 hours and reducing cargo damage at gauge break junctions. The conversion holds the rail network topology fixed while reducing a specific transport friction. **Outcome definition:** DMSP calibrated total nighttime luminosity at the district-year level, a standard proxy for local economic activity in developing countries. **Treatment:** Binary; a district is treated when its railway zone's meter-gauge network reaches the approximate conversion midpoint, enabling through-running of freight and passenger trains. **Data:** SHRUG DMSP nightlights (1994–2013) and datameet railway stations (8,990 stations), district-year panel. **Method:** Staggered difference-in-differences with district and year fixed effects; standard errors clustered at the state level. Callaway and Sant'Anna (2021) estimator as robustness. **Sample:** All Indian districts with non-missing nightlight data; MG-exposed districts defined as those in states where >30% of railway stations belong to historically meter-gauge-intensive zones. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).