The Long and Winding Road to Growth:

Terrain, History, and Economic Development in India [§]

Mohammad Asad Malik [†]

August 4, 2025

Abstract

This paper provides causal evidence on the impact of national highways on local GDP per capita in India. To isolate exogenous variation in modern highway density, I implement a dual-instrument strategy that combines terrain features with historical road networks from the pre-independence period. I extend an original spatial dataset by digitizing Indian road maps from 1928, 1969, 2011, and 2021 using supervised image classification techniques, and merge this with high-resolution grid-level GDP data.

Two-stage least squares estimates with state fixed effects show that a one percentage point increase in national highway density raises local income by 11–14 percent on average. However, these gains are highly uneven: the effects are concentrated in the richest regions, while poorer areas experience little to no benefit. The results are robust to a range of checks, including reverse causality tests and leave-one-state-out specifications.

To interpret these patterns, I develop a dynamic model of regional capital accumulation with threshold effects. In the model, infrastructure investments raise productivity only in high-capital regions, generating divergence even under equitable allocation. Together, the empirical and theoretical findings highlight how capital complementarities and historical path dependence can entrench spatial inequality—despite large-scale infrastructure expansion.

JEL Codes: O18, R11, H54, O41

[§] I am indebted to my advisor Professor *Chetan Ghate* at the Indian Statistical Institute, Delhi Center for his intellectual support and exceptional guidance. I also thank *Tishara Garg* for her mentorship during the early stages of this project and for introducing me to the road network data. Finally, I thank the *National Library of Australia* for its archival maps, without which this work would not have been empirically implementable.

[†] Economics and Planning Unit, Indian Statistical Institute, Delhi Center

1 Introduction and Motivation

India has seen one of the most rapid road expansions in the developing world. Since 2000, India has doubled the length of its road system, making it the second largest in the world.[1] Flagship programs such as the Pradhan Mantri Gram Sadak Yojana (PMGSY) and the Bharatmala Pariyojana have prioritized lagging and remote regions to promote spatially inclusive growth. Yet despite this expansion, spatial income inequality remains remarkably persistent. Districts that ranked among the poorest a decade ago—such as those in Jharkhand, Bihar, and Chhattisgarh—remain at the bottom today. This disconnect between large-scale infrastructure investment and limited convergence forms the central puzzle this paper addresses.

One explanation is that the returns to infrastructure are shaped by local conditions—fiscal capacity, institutional quality, or the presence of complementary inputs. Roads provide a particularly effective lens to study this variation: they are spatially precise, economically consequential, and relatively well-measured. But their returns vary sharply across space. In this paper, I show that a 1 percentage point increase in National Highway (NH) density raises income by 11–14% in high-income areas, but has no measurable effect in poorer regions. The same roads, built under the same program, generate sharply different outcomes depending on where they are placed.

To explain this variation and derive testable predictions, I develop a stylized dynamic model and test it with new spatial data linking roads and economic activity across India. The model posits that the marginal return to infrastructure depends on the local capital stock: roads significantly raise output only once capital surpasses a threshold, reflecting fixed costs and missing complements in low-capacity environments. Below this threshold, infrastructure has little effect—leading to persistent divergence even amid expanding connectivity. This yields a sharp empirical prediction: the elasticity of output with respect to infrastructure increases in local capital, proxied here by income.

The paper makes three main contributions on the empirical front. First, I construct a novel spatial dataset that merges GDP per capita at a $0.5° \times 0.5°$ grid level with road density disaggregated by type and time. The road dataset builds on earlier efforts to digitize Survey of India maps but extends them significantly: I incorporate previously unused maps from 1928, 1969, and 2021, increase spatial resolution, and implement a supervised machine learning pipeline. Specifically, I train Random Forest classifiers to identify road classes from map raster images, filter out cartographic artifacts, and extract road density by type. This enables consistent historical comparisons and precise disaggregation across space and road categories.[2]

Second, I implement a dual-instrument strategy to address endogenous placement of modern highways. The first instrument captures topographic suitability, derived from mean elevation, elevation variance and terrain ruggedness. The second uses historical road density from 1928 and 1969 maps, exploiting legacy infrastructure patterns that predate modern economic geography. Together, these instruments allow me to isolate exogenous variation in modern national and state highway coverage.

Third, I estimate the causal effect of national highways on grid-level GDP per capita and document substantial heterogeneity. The results hold across different spatial resolutions and identification strategies, and align closely with the model's prediction. State highways, by contrast, display no robust effect on income once endogeneity is addressed. The data thus reveal a key tension: while infrastructure may be uniformly allocated, its benefits are not.

The theoretical mechanism builds on threshold-based models of capital accumulation (e.g., Ghatak (2015), Ghatak (2017), Sachs et al. (2004)), adapted to infrastructure. In my framework, road investment raises productivity only once a region surpasses a capital threshold, leading to multiple steady states. This structure helps explain why similar road investments yield divergent outcomes across space, and why infrastructure may entrench, rather than reduce, spatial inequality.

[1] Source: Ministry of Road Transport and Highways, Annual Report 2023–24.

[2] Clarification on the notation: for the entire analysis, whenever the term '*National Highways*' is used, it includes Golden Quadrilateral segments and all-weather motorable National Highways. '*State Highways*', for ease of interpretation, denote all-weather motorable roads which are not classified as '*National Highways*'. Finally, there is one more class of roads, that I refer to as '*Arterial/Feeder Roads*'. Although this class of roads doesn't constitute my analysis, it includes all roads that are motorable in fair weather, those which are 'less important', or those where delay might occur. The interested reader is encouraged to refer to the latest road map on Survey of India's website for further details.

This paper contributes to three strands of literature. First, it extends work on infrastructure and development, particularly in India. While prior studies have focused on flagship programs or specific corridors (e.g., Asturias et al. (2019); Asher and Novosad (2020)), I provide nationwide evidence on heterogeneous returns, using newly digitized historical infrastructure. Second, I build on the market access tradition (Donaldson and Hornbeck (2016); Donaldson (2018)), using historical networks not just to explain persistent advantage, but to instrument for modern roads. Third, I contribute to the literature on geographical and historical persistence (e.g., Acemoglu et al. (2001); Nunn and Puga (2012); Dell (2010)), showing how terrain and colonial infrastructure jointly shape current opportunity sets.

Ultimately, this paper shows that infrastructure alone cannot equalize outcomes. Without local complements, roads deliver few gains. Where capacity is high, the same investments generate strong returns. The findings suggest that infrastructure may reinforce spatial inequality—not just because of how it's placed, but because of where it's placed. These findings speak to a broader principle: without complementary investments in absorptive capacity, infrastructure may not close spatial gaps; it may deepen them.

The remainder of the paper is organized as follows: Section 2 presents the theoretical model. Section 3 describes the data and digitization strategy. Section 4 outlines the empirical framework. Section 5 discusses the main empirical findings. Section 6 concludes.

2 Model

Model Setup

Consider a continuum of regions indexed by $i \in \mathbb{N}$, each endowed with capital $K_i > 0$, a common productivity level $A > 0$, and access to a public road network $R \geq 1$, which is treated as an exogenous input, shared across regions.

Regions differ only in their initial capital endowments; I abstract from heterogeneity in preferences, institutions, or geography.

Each region chooses between two technologies: a low productivity (baseline) technology and a high productivity (advanced) technology, the latter requiring the payment of a fixed cost $F > 0$:[3]

- *Low-productivity regime:*
$$Y_i = AK_i^{\alpha} R^{\beta}$$
- *High-productivity regime:*
$$Y_i = AK_i^{\alpha} R^{\beta+\gamma} - F$$
where $\gamma > 0$ captures stronger road complementarity under the advanced technology.

A region adopts the advanced technology if its net output exceeds that of the baseline regime, yielding a threshold level of capital $\hat{K}$ above which adoption is profitable:[4]

$$AK_i^{\alpha} R^{\beta+\gamma} - F > AK_i^{\alpha} R^{\beta}$$

The endogenous capital threshold is, therefore:

$$K_i > \hat{K} \equiv \left(\frac{F}{A(R^{\beta+\gamma} - R^{\beta})} \right)^{1/\alpha}$$

[3] Fixed in the sense that it must be paid in any period where the advanced technology is used, but does not vary with the level of output produced in this regime.

[4] I assume that strictly positive net gains are required to induce switching. At $K_i = \hat{K}$, the region remains in the low regime.

Output is thus given by:

$$Y_i = \begin{cases} AK_i^{\alpha} R^{\beta} & \text{if } K_i \leq \hat{K} \\ AK_i^{\alpha} R^{\beta+\gamma} - F & \text{if } K_i > \hat{K} \end{cases}$$

Capital Accumulation and Steady States

Capital accumulates according to the standard differential equation:

$$\dot{K}_i = sY_i - \delta K_i$$

where $s \in (0, 1)$ is the savings rate and $\delta > 0$ is the depreciation rate. I examine each regime in turn.

Low Regime: When $K_i \leq \hat{K}$, substituting for Y_i gives:

$$\dot{K}_i = sAK_i^{\alpha} R^{\beta} - \delta K_i$$

Setting $\dot{K}_i = 0$ yields the low-regime steady state:

$$\underline{K_i} = \left(\frac{sAR^{\beta}}{\delta} \right)^{\frac{1}{1-\alpha}}$$

A poverty trap arises when the steady-state capital level under the baseline technology lies below the adoption threshold $\hat{K}$. Substituting the expressions for $\underline{K_i}$ and $\hat{K}$, the condition becomes:

$$F > F^* \equiv \left(\frac{sAR^{\beta}}{\delta} \right)^{\frac{\alpha}{1-\alpha}} \cdot A \cdot (R^{\beta+\gamma} - R^{\beta})$$

This threshold F^* defines the minimum fixed cost necessary to sustain a poverty trap. The parameters are interpreted as follows:

- A higher A, R, or γ raises potential output, requiring a higher fixed cost to induce a trap.
- A low s or a high δ reduces the steady-state capital stock in the low regime. As a result, even a relatively modest F can push $\hat{K}$ above $\underline{K_i}$, thereby increasing the likelihood of a poverty trap.

In short, traps arise when the economy cannot accumulate enough capital to make the transition to the high technology regime worthwhile, even though roads are present.

High Regime: When $K_i > \hat{K}$, we have:

$$\dot{K}_i = s(AK_i^{\alpha} R^{\beta+\gamma} - F) - \delta K_i$$

Let $g(K_i) \equiv \dot{K}_i$. The high steady state $\overline{K_i}$ satisfies:

$$s(A\overline{K_i}^{\alpha} R^{\beta+\gamma} - F) = \delta \overline{K_i}$$

Unlike in the low regime, this equation does not yield a closed-form solution for $\overline{K_i}$ due to the presence of the fixed cost F. Nevertheless, I can characterize the existence (in particular, I can ensure that it exists beyond $\hat{K}$), and uniqueness of the high steady state analytically, under a mild restriction, as shown below.

Proposition 1. *(Existence and uniqueness of high productivity steady state): There exists a unique high-productivity steady state $\overline{K_i} > \hat{K}$ satisfying:*

$$s\left(A\overline{K_i}^{\alpha} R^{\beta+\gamma} - F\right) = \delta \overline{K_i}$$

provided that $\exists\, \varepsilon > 0$ *such that* $\forall\, K_i \in (\hat{K}, \hat{K} + \varepsilon)$,

$$AK_i^{\alpha} R^{\beta+\gamma} - F > \frac{\delta}{s} K_i$$

That is, capital accumulation is feasible (over a non-degenerate interval) immediately after crossing the threshold. Given the strict concavity of $g(K_i) = s(AK_i^{\alpha} R^{\beta+\gamma} - F) - \delta K_i$*, this condition ensures that* $g(K_i)$ *crosses zero exactly once for* $K_i > \hat{K}$*.*

Proof. Provided in Appendix A. □

Proposition 1 is illustrated in Figure 1. It establishes that once a region crosses the capital threshold $\hat{K}$, it can sustain the higher steady state $\overline{K_i}$ characterized by faster capital accumulation due to higher productivity of infrastructure. The existence of two distinct steady states with $\underline{K_i} < \hat{K} < \overline{K_i}$ implies that identical infrastructure may generate vastly different outcomes depending on the region's initial conditions. This raises a natural question: how does the marginal productivity of infrastructure differ across these regimes?

To answer this, I compare the output elasticity with respect to infrastructure in the two steady states. This yields the model's core empirical prediction:

Proposition 2. *(Heterogeneous Infrastructure Returns): Let* $\underline{K_i}$ *and* $\overline{K_i}$ *denote the steady states in the low- and high-productivity regimes, respectively, with* $\underline{K_i} < \hat{K} < \overline{K_i}$*. Then, the marginal effect of infrastructure on output is strictly higher in the high regime:*

$$\left.\frac{\partial Y}{\partial R}\right|_{K=\overline{K_i}} > \left.\frac{\partial Y}{\partial R}\right|_{K=\underline{K_i}}$$

Specifically,

$$\left.\frac{\partial Y}{\partial R}\right|_{K=\underline{K_i}} = A\beta \underline{K_i}^{\alpha} R^{\beta-1}, \quad \left.\frac{\partial Y}{\partial R}\right|_{K=\overline{K_i}} = A(\beta+\gamma)\overline{K_i}^{\alpha} R^{\beta+\gamma-1}$$

This result follows directly from the model's structure: infrastructure is more productive in the advanced regime $(\gamma > 0)$ *and output is increasing in capital* $(\alpha > 0)$*.*

Proof. Provided in Appendix A. □

Testable implication: *The marginal return to infrastructure is increasing in capital. Since capital stock is unobservable, I proxy it with income per capita, yielding the prediction that richer regions should exhibit higher infrastructure elasticities.This forms the basis of the empirical heterogeneity analysis that follows.*

Proposition 2 formalizes the intuitive asymmetry in road impacts: the same investment yields greater gains where productive capacity is already high. Beyond cross-sectional variation, the model also implies that aggregate returns to infrastructure depend on the share of regions above the adoption threshold—motivating the estimation of average effects before turning to heterogeneity in the results section.

From a normative standpoint, multiple equilibria create a classic allocation dilemma. In decentralized systems, low-capital regions remain trapped without external support. A social planner allocating a fixed budget must weigh two objectives: (i) maximizing aggregate output by targeting road construction toward high-capacity regions, and (ii) promoting convergence by subsidizing adoption in low-capacity regions via a reduction in the (privately paid) fixed cost component. Each allocation strategy entails a tradeoff between efficiency and equity, underscoring the normative stakes of infrastructure policy. While I do not formalize the planner's problem in this paper, doing so—e.g., by maximizing a weighted social welfare function subject to regime-switching constraints—offers a clear direction for future work.

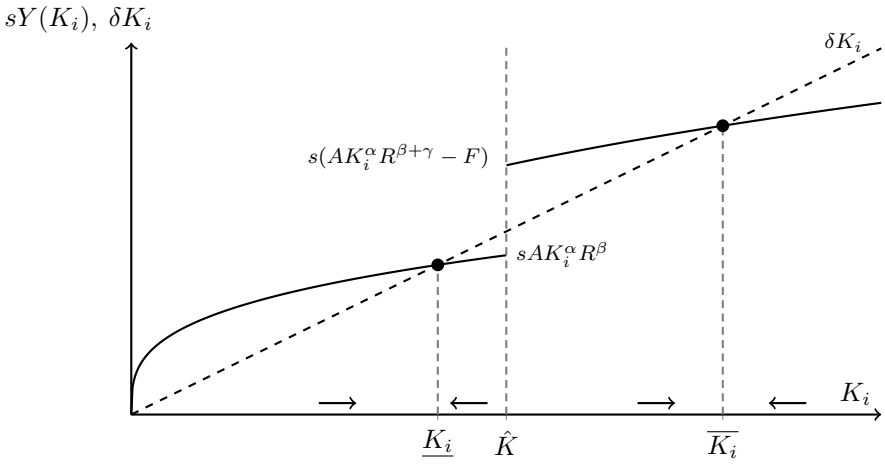

Figure 1: Capital Accumulation with Regime Shift and Threshold

3 Data

The model in Section 2 motivates a central empirical question: do the returns to infrastructure depend on a region's capital stock—specifically, whether it lies above a critical threshold? Testing this mechanism requires data on both economic development and road infrastructure exposure at fine spatial resolution. Moreover, because modern road construction is unlikely to be randomly assigned, the empirical strategy must address issues of endogeneity and selection. This section describes the dataset assembled to meet these needs.

To measure economic activity, I use gridded GDP per capita estimates developed by Rossi-Hansberg and Zhang (2025). The dataset provides annual GDP estimates from 2012 to 2021 at a 0.5° × 0.5° resolution (approximately 55 km × 55 km at the equator), covering all countries.[5] To ensure temporal alignment with the 2011 road map, I use the 2012 GDP cross-section, which precedes the highway expansion under the Bharatmala Program.

The Rossi-Hansberg–Zhang pipeline relies on a machine learning model to predict grid-level GDP shares from a rich suite of satellite and remote sensing indicators. Inputs include population density, nighttime lights, land cover, vegetation indices, impervious surface area, terrain elevation, and meteorological data, among others. The model is trained using random forest regression and calibrated against official national or sub-national GDP aggregates to scale predicted shares into absolute levels. In their validation exercises, the authors report out-of-sample R^2 values exceeding 0.9 for GDP levels and around 0.6 for annual growth rates, with good tracking of known real-world events such as the COVID shock in China.

India's grid-level estimates are extracted from this global dataset and mapped directly onto the 0.5° grid structure, resulting in 1,303 non-missing grid cells covering India's mainland and island territories. These cells span the full range of India's economic and geographic variation, from urbanized coastal zones to interior tribal districts. I use the log of GDP per capita in 2012 (const. 2017 USD) as the primary outcome of interest in all the regressions.[6] This is illustrated below in Figure 4 at the 0.5° × 0.5° resolution.

Overview of Spatial Dataset Construction for the Independent Variables

The empirical analysis in this paper is grounded in a custom-built spatial panel that combines grid-level GDP data with geo-coded road network measures, terrain characteristics, and administrative identifiers.[7] Constructing this dataset required integrating information from multiple sources—both raster and vector formats—and harmonizing them to a common geographic framework.

The first major input is road infrastructure, sourced from historical and contemporary maps of India spanning

[5] Estimates are available at three spatial resolutions: 1° × 1°, 0.5° × 0.5°, and 0.25° × 0.25°. I use the latter two for my analyses.

[6] I inflate the GDP per capita number by a scale factor of 10^{12} and add 1 to it before logging, to preserve ordinal rankings and avoid undefined logs.

[7] The construction process is described in full in Appendix C and the associated raw maps are given in Appendix D.

1928 to 2021.[8] The classification methodology builds on Allen and Atkin (2022). The maps were manually georeferenced, digitized, and their artifacts classified using a combination of supervised learning and rule-based filters. Random Forest models were trained to identify roads based on captured features. Post-classification, the networks were clipped to India's 0.5° grid structure. I extract road density (in km/km^2)—defined as total road length divided by cell area—separately for all-weather motorable national highways, other motorable roads, and the Golden Quadrilateral wherever distinguishable. The resulting road networks are illustrated in Figure 2.

These historical road datasets also form the basis of the instrumental variables (IV) strategy used to estimate the causal impact of road infrastructure on economic outcomes. Both the 1928 and 1969 maps serve as instruments for modern NH density, identifying Local Average Treatment Effects (LATE) in each case. However, the two instruments differ in their alignment with the contemporary highway network, with important implications for sample construction and interpretation, which are discussed below.

IV methods capture the effect of treatment (modern NH density) on the subpopulation of compliers—regions whose treatment status is shifted by the instrument. In the 1969 specification, historical road density exhibits strong persistence into the modern network; the share of 'impersistent' observations is minimal and thus not excluded. As a result, the exclusion restriction is plausibly satisfied across the full sample.[9] Accordingly, I retain all observations in the 1969 analysis, and the resulting LATE pertains to a wide subset of regions where historical infrastructure plausibly shaped modern access.

The 1928 specification, by contrast, presents greater challenges. A substantial share of roads classified as national highways in 1928 do not persist in the 2011 network—due to downgrading, rerouting, or disappearance. Including these 'impersistent' regions weakens the first stage and leads to a statistically significant Sargan test, raising concerns about instrument validity. In these cases, the exclusion restriction becomes tenuous: the historical classification no longer reliably predicts modern access and may correlate with endogenous factors. I therefore restrict the sample to regions where 1928 roads remain in the modern NH network, thereby improving instrument relevance and restoring overidentification validity. The resulting 1928 estimate should be interpreted as a LATE for regions with persistent historical infrastructure—a narrower complier population than in 1969, but one for which the identifying assumptions are more credible.

The second major input is topography. Using 90-meter SRTM rasters, I extract the mean elevation, elevation variance, and terrain ruggedness index (TRI) for each grid cell. These variables form the basis for a topography-based instrument addressing endogenous placement of modern roads.[10] The distributions of each component, along with the resulting instrument, are shown in Figure 3.

Finally, I assign each grid cell to a single Indian state by overlaying the grid shapefile with GADM administrative boundaries and for grids that fall on the boundaries of multiple states, I retain the state with the largest area overlap with the grid. This mapping enables the inclusion of state fixed effects and clustering in all regressions.[11]

With all spatial variables—GDP, road density, and terrain—harmonized to the 0.5° grid resolution, the next section outlines the empirical framework for estimating road effects while addressing endogeneity. Grid cells with missing values in either GDP or covariates are excluded from the analysis; no further manual trimming is applied at any point.

[8] For the 1928 map in particular, I extract all roads, since the cartographic convention at the time rendered roads in the same ink with only subtle variation in stroke thickness. To isolate trunk roads, I manually delete feeder roads by cross-referencing with the 1939 map. However, it is virtually impossible to reliably distinguish state highways from arterial roads, so state highways are excluded from the 1928 analysis. I therefore caution the reader to interpret results from the 1928 instrument with care, as the final processed map is unlikely to be perfectly accurate.

[9] Estimates using a cleaned sample that excludes impersistent cells yield nearly identical results and are available upon request.

[10] The topography-based instrument, Topo PC1, is constructed using principal components analysis (PCA) applied to standardized elevation, elevation variance, and TRI. The loadings are shown in Table 3 (b).

[11] The same pipeline is also replicated for the analysis at 0.25° × 0.25° resolution. That specification uses a sample of 4881 grids and includes district fixed effects, along with the relevant level of clustering.

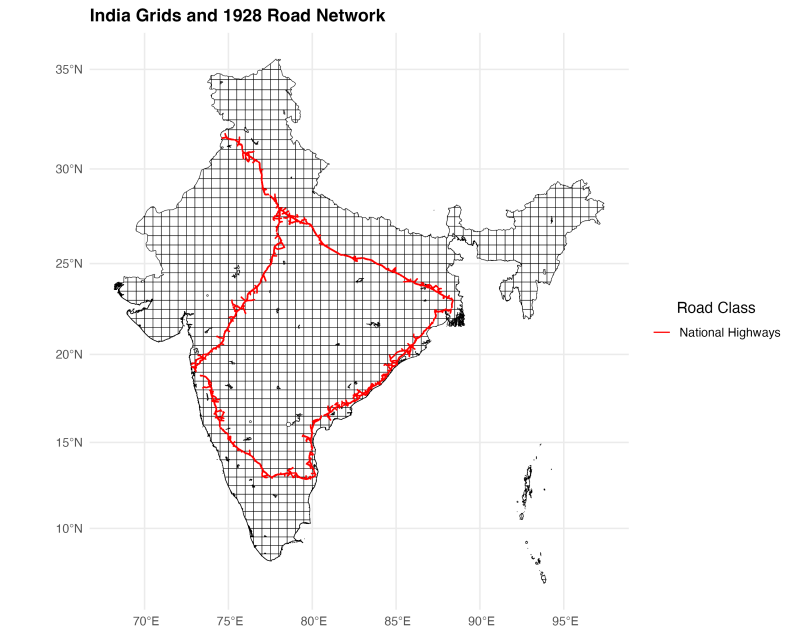

(a) 1928 Roads

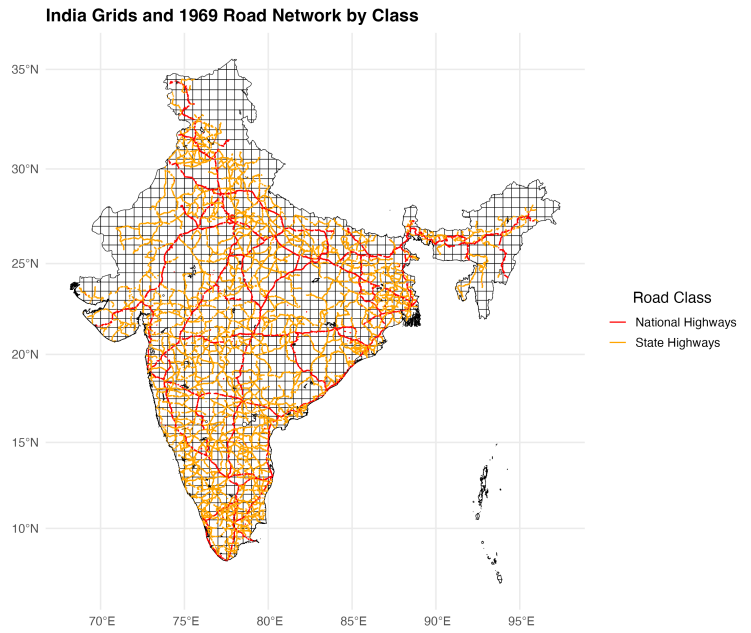

(b) 1969 Roads

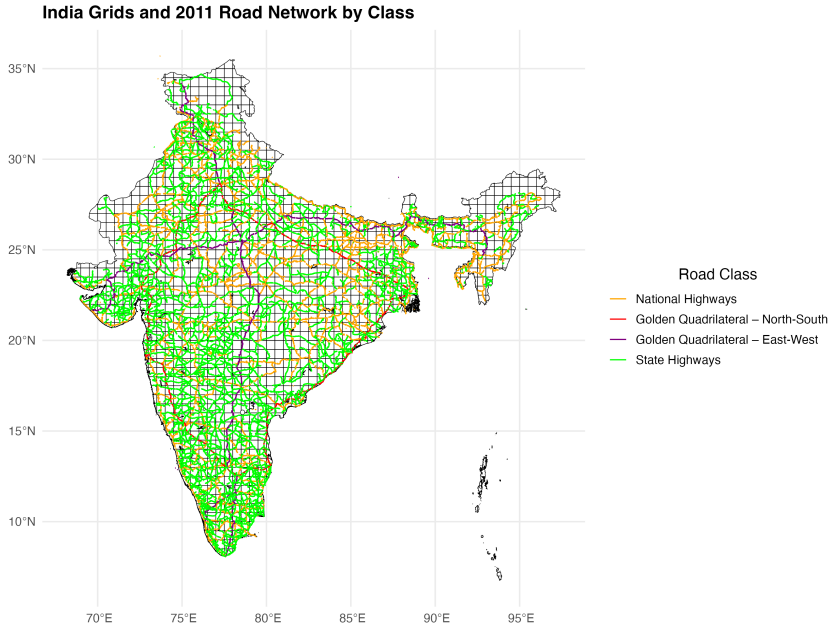

(c) 2011 Roads

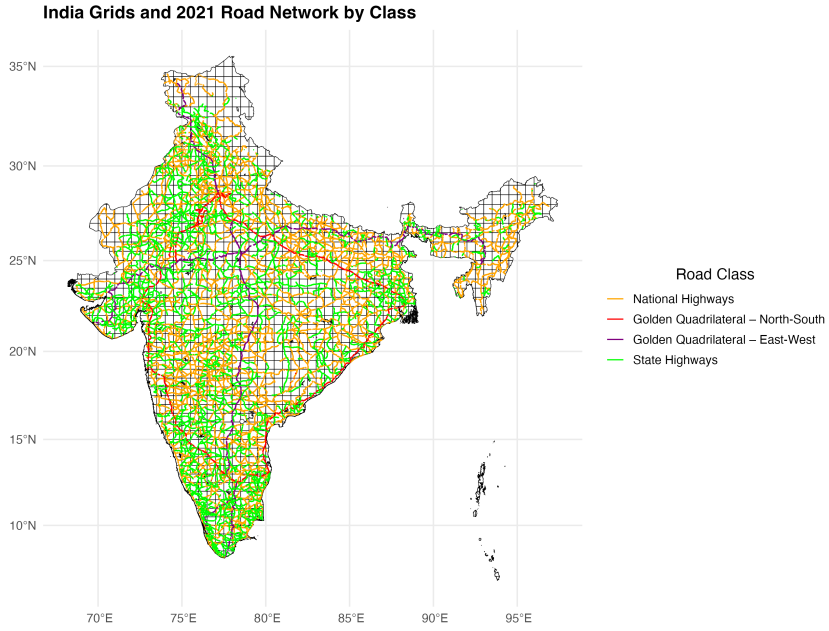

(d) 2021 Roads

Figure 2: Processed Road Network Maps (1928–2021)

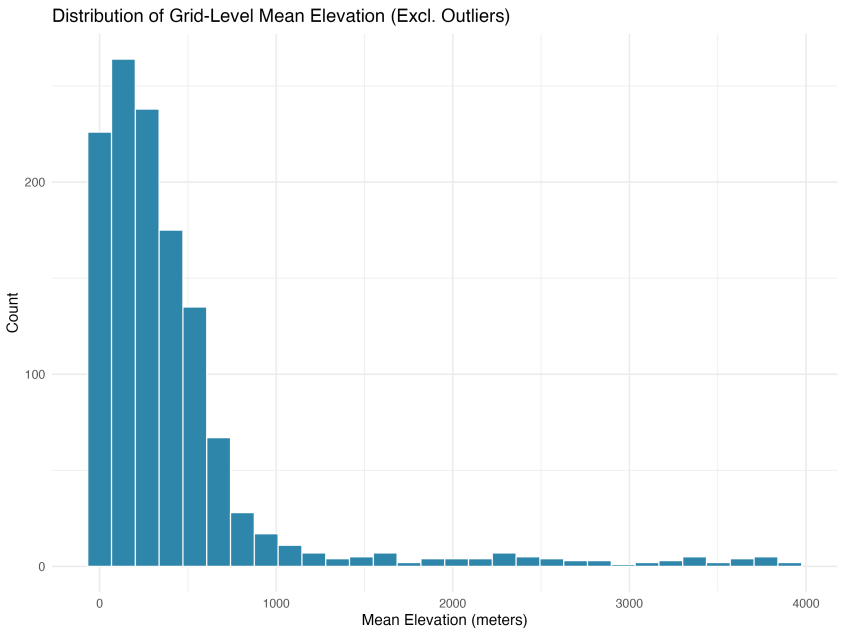

(a) Mean Elevation

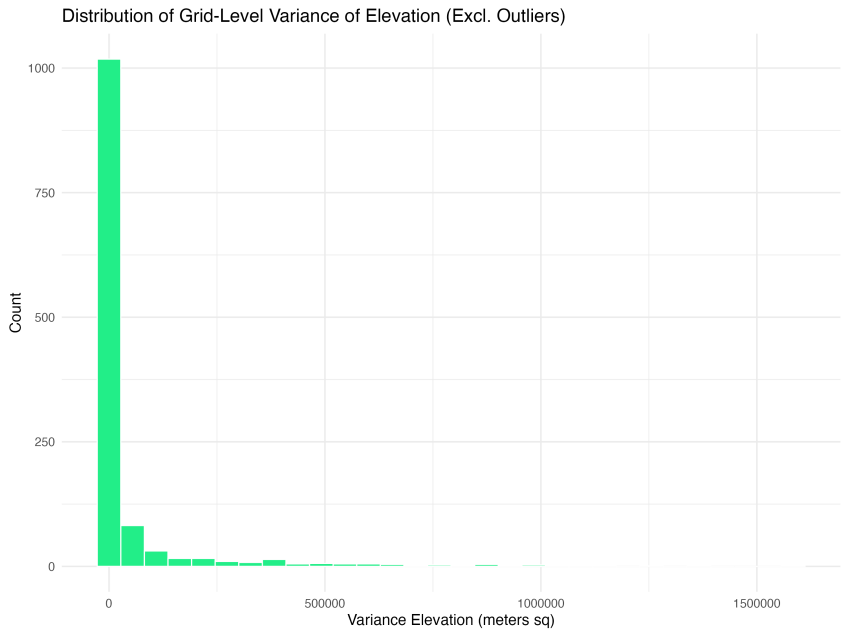

(b) Elevation Variance

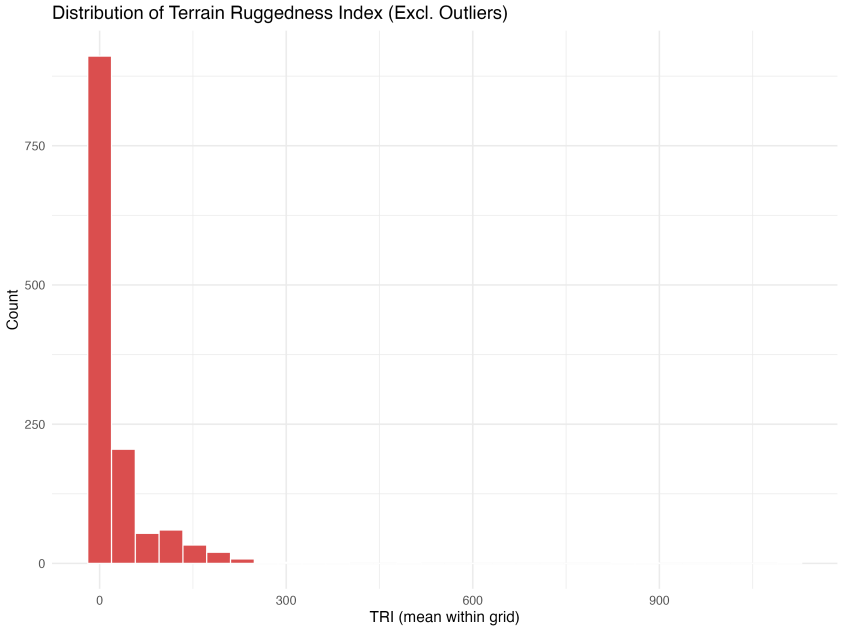

(c) Mean Terrain Ruggedness Index (TRI)

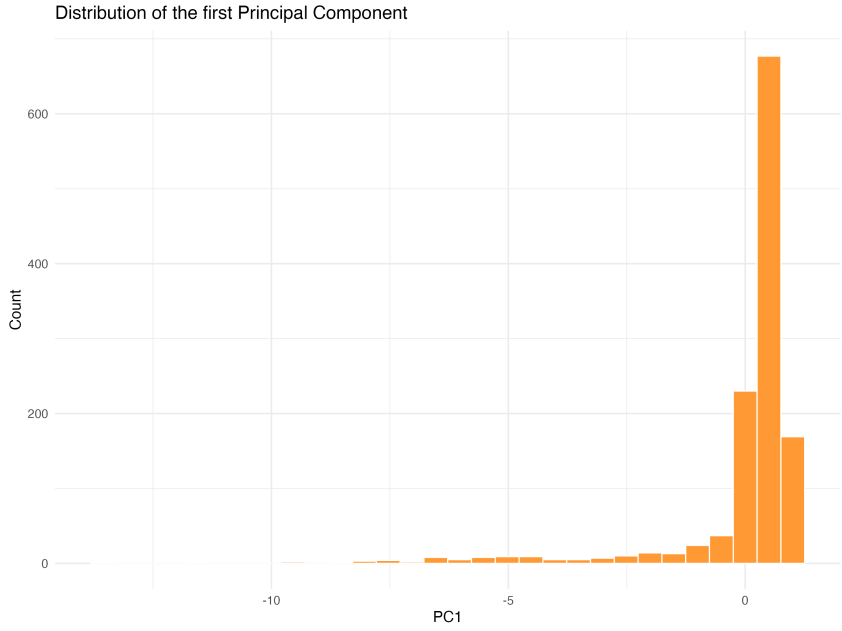

(d) Topography-Based Instrument (PC1)

Figure 3: Distribution of Topographic Features and Instrument

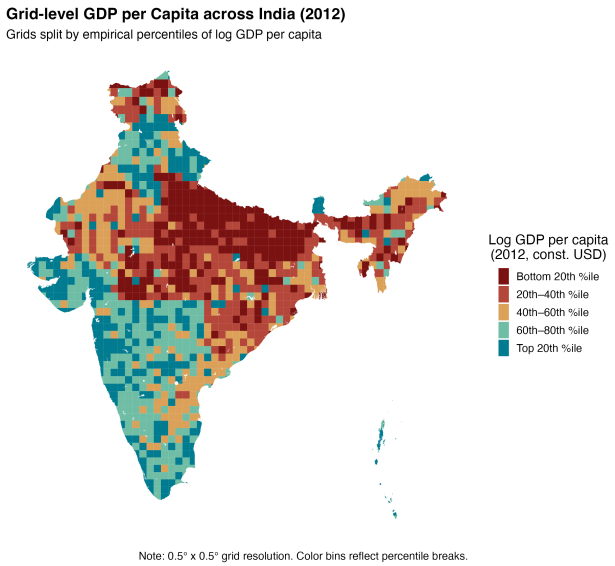

Figure 4: Spatial Distribution of log(GDP per capita), 2012

4 Empirical Strategy

This section outlines the empirical framework for estimating the causal effect of road infrastructure on local economic outcomes. I begin with a simple OLS regression of GDP per capita in 2012 on road density across grid cells. Recognizing that road placement is unlikely to be exogenous—governments often prioritize regions with high economic or political importance—I then implement an IV strategy that leverages geographic constraints and historical road networks to isolate plausibly exogenous variation.

Importantly, the empirical design is guided by the model in Section 2, which predicts that the marginal return to infrastructure is increasing in local capital stock. While the IV strategy identifies average treatment effects, I also extend the empirical framework by interacting road density with income proxies to assess whether infrastructure yields higher gains in capital-rich regions. This allows me to empirically test the model's central prediction about threshold-based heterogeneity.

Baseline OLS Specification

The basic OLS model is as follows:

$$\log(\text{GDP per cap})_i = \alpha + \beta \cdot \log(\text{Road Density}^{\tau})_i + \delta_s + \varepsilon_i$$

where i denotes the grid cell and $\tau \in \{\text{National}, \text{State}\}$ denotes highway type. Each regression is estimated separately for each τ.[12] The term δ_s captures state fixed effects, which absorb all variation common across grids within a state such as regional governance quality, state-level economic policy, shared developmental history, or fiscal transfers. The specification does not include any grid-level covariates due to the absence of consistent, high-quality data at this spatial resolution; these are subsumed under the state dummies. This approach leverages variation in road density across grids within the same state to estimate the relationship with local income, net of state-level confounders.

While straightforward to estimate, the OLS specification relies on strong identifying assumptions and is vulnerable to bias. One concern is reverse causality: roads may be targeted toward poorer areas for redistribution or political visibility. For example, the PMGSY program explicitly aimed to connect backward and remote districts. A second concern is omitted variable bias: unobserved factors such as local governance, land rights, or

[12]This study focuses on highways alone, as the classification of lower-tier roads in the 2011 map is complicated by their ink shade overlapping with non-road map artifacts; future work could extend the analysis to these networks as data quality improves.

historical conflict may influence both road construction and economic outcomes. Without credible instruments, OLS is likely to misstate the true causal effect.

IV Specification

To address these endogeneity concerns, I use an IV strategy that instruments for modern road density using two sources of plausibly exogenous variation:

1. **Topographic Suitability (Topo PC1)**: This is the first principal component of three terrain-based features—mean elevation, elevation variance, and terrain ruggedness (TRI)—measured at the grid level. The logic here is that topography shapes the feasibility and cost of road construction but, conditional on road access, does not directly influence income.
2. **Historical Road Networks (1928 and 1969):**[13] strongly predict present-day road density, as seen in Figure 6. These are constructed from digitized historical road maps and reflect the density of trunk roads and highways in each grid during the pre-liberalization era. Historical placement was largely driven by administrative, military, or colonial priorities, rather than contemporary economic development potential. This provides a natural source of variation that predicts modern roads, but is unlikely to be correlated with present-day economic shocks or unobserved determinants of income.

The general two-stage least squares framework is summarized below:

First Stage:

$$\log(\text{Road Density}^{\tau})_i = \gamma + \zeta \cdot \log(\text{Historic Road Density}^{\tau})_i + \lambda \cdot \text{Topo PC1}_i + \omega_s + \epsilon_i$$

Second Stage:

$$\log(\text{GDP per capita})_i = \alpha + \beta \cdot \log(\widehat{\text{Road Density}}^{\tau})_i + \delta_s + \varepsilon_i$$

The identifying assumption is that the instruments affect GDP only through their effect on road density. In other words, both historical roads and topography must not influence income via any other channel. While this exclusion restriction cannot be formally tested, I offer three forms of justification.

First, elevation and ruggedness are physical constraints, not economic determinants, once roads and state conditions are accounted for. Second, historical road placement reflected strategic or military logic, not developmental targeting. Third, I conduct a plausibility check: I regress the fitted values of road density from the first stage on log GDP per capita and find no statistically significant relationship. This suggests that the exogenous variation induced by the instruments is not mechanically correlated with income levels, supporting—but not conclusively validating—the exclusion assumption.

To further guard against violations, Section 5 also examines robustness to alternative instrument definitions and subsample exclusions. All first-stage regressions yield F-statistics above 10, indicating strong instrument relevance.

As discussed in Section 1, the use of topography and historical networks as instruments follows a large literature on infrastructure and geography-based identification strategies. These include applications to colonial roads, elevation-based access, and natural terrain breaks, which similarly treat infrastructure placement as endogenous.

[13]The 1969 map is the earliest in the series that allows for consistent visual identification of state highways based on color differentiation. Earlier maps either lacked such differentiation or used similar shades for multiple road types, rendering classification unreliable. The 1928 map is the earliest available comprehensive record of the national road network. No earlier cartographic source offers comparable coverage.

5 Results

This section presents the core findings linking highway infrastructure to regional income. The analysis proceeds in four parts. First, I estimate the average causal effect of National Highway (NH) density on local GDP per capita using an instrumental variables (IV) strategy. Second, I examine whether this effect varies systematically with baseline income, consistent with the model's prediction of threshold-based complementarities. Third, I contrast these results with those for State Highways (SH), which do not generate comparable gains. Finally, I assess the robustness of the findings.

Baseline OLS estimates, reported in Table 1, show a positive but modest association between NH density and income: a 1% increase in NH density is linked to a 0.2–0.4% rise in GDP per capita, conditional on state fixed effects and clustering at the state level. These estimates, however, likely understate the true effect. As shown in Table 2, which serves as a falsification test, regions with higher income in 2012 were not systematically more likely to receive road upgrades over the subsequent decade. In fact, post-2011 road expansion was weakly negatively correlated with 2012 GDP, consistent with redistributive targeting of infrastructure toward poorer areas. This suggests that OLS estimates are biased downward, not upward, alleviating concerns that the observed association is driven by reverse causality. Moreover, if infrastructure delivers returns only after a threshold level of complementary inputs, pooling across heterogeneous districts may obscure nonlinearities—further motivating an instrumental variables (IV) approach.

IV Estimates of National Highway Effects. Table 6 reports IV estimates using historical road placement and topography as instruments for modern NH density. Using the 1928 colonial road map—reflecting administrative and military priorities prior to modern development patterns—I find (Column 4) that a 1% increase in NH density raises GDP per capita by 14%. These roads, laid nearly a century ago, continue to shape the economic geography of India—underscoring the persistence of early spatial advantages.

Compared to this baseline, IV estimates using 1969 road density (Column 2) are smaller (around 0.009) but remain significant. However, their credibility is limited by potential endogenous placement. Terrain-based instruments (Column 1) yield similar results but are weak in the first stage when used alone. When combined with historical instruments (Columns 3 and 5), results remain consistent, and over-identification tests generally support joint validity—though with some caveats, discussed below.

A potential concern with this identification strategy is reverse causality: economically stronger regions may have received denser highway coverage even in historical maps due to their long-standing importance. Table 8 addresses this by regressing the fitted values of NH density from the first stage on 2012 income. In Columns 2 and 3—based on terrain and 1928 colonial roads—coefficients are small and statistically insignificant, indicating no systematic relationship between historical instrument values and modern economic activity. Column 1, which uses the 1969 map, shows a large and significant coefficient, suggesting that placement in that year may already reflect economic considerations. This strengthens the case for using the 1928 instrument as the preferred source of variation for causal inference.

In specifications that use the 1928 instrument, the Wu-Hausman test fails to reject the null of exogeneity. This should not be taken to imply that NH placement is exogenous, particularly since the test's power is limited under weak instrumentation or small sample sizes. As shown in Table 2, post-2011 road expansion targeted poorer regions, suggesting that equity-based allocation characterizes infrastructure placement in India. The retention of the null in the Wu-Hausman test is more plausibly attributed to reduced power: the test's sensitivity declines when instruments are weak, when sample sizes are small (as in the 1928 subsample), or when the endogenous regressor and the instrument are only weakly correlated. Hence, the test may fail to detect endogeneity even when it exists. I therefore rely on institutional arguments and historical context—alongside instrument strength and overidentification diagnostics—to justify the IV strategy.

Taken together, these results indicate a large and robust causal effect of NH infrastructure. Having established the average return to highways, I now examine whether these gains are concentrated in richer regions, consistent with threshold-based predictions of the model.

Heterogeneous Returns by Baseline Income. Figure 7 plots IV estimates for each income quartile, using

alternative instruments. Although confidence intervals are wide—reflecting smaller within-group samples—the pattern in point estimates is clear: returns to NH density increase with baseline income. This gradient is most pronounced under the 1928 colonial instrument, albeit with lower precision.

To formally assess heterogeneity, Table 9 reports an IV regression with a binary high-income interaction term. The interaction term is large, positive, and statistically significant, while the baseline effect remains small and insignificant. This pattern reinforces the model's prediction: infrastructure raises income only where complementary factors (capital, skills) are already present.

Figure 8 presents formal estimates of the income threshold beyond which National Highway (NH) investments yield statistically significant economic returns. For each candidate threshold along the log-transformed GDP per capita distribution, the sample is split into high- and low-income groups, and separate IV regressions are estimated using 1969 road density and topography as instruments. The resulting coefficients are plotted in blue (above threshold) and red (below), with 90% confidence intervals. The dotted green line marks the lowest threshold at which the high-income coefficient becomes significantly positive. This occurs at 15.36 in transformed logs—corresponding to a raw per capita GDP of $4.68 (constant 2017 USD)—placing it above the 99.47th percentile of the national distribution. In other words, measurable gains from NH investments are confined almost entirely to the very richest regions, with no significant effect observed elsewhere.

A similar exercise using the 1928 road network and topography—the preferred, more plausibly exogenous instrument—yields a comparable threshold of 13.63 ($0.83 in constant 2017 USD), albeit with considerably noisier estimates due to weaker first-stage variation. While the 1969 instrument delivers sharper statistical precision, it is subject to potential endogeneity, as highlighted in the reverse causality analysis. Together, both exercises reinforce the same substantive point: the growth effects of highway investments are highly uneven, emerging only above a very high income threshold. This trade-off between statistical power and causal credibility is evident, but the core result—that infrastructure returns are conditional on local capacity—remains robust across instruments.

State Highways as a Contrast. The results for SHs offer a useful empirical contrast. Table 5 shows that SH density is only weakly predicted by terrain, and historical maps yield modest instrument strength. Even when first-stage relevance is restored (Columns 2–3), second-stage estimates in Table 7 are near zero and precisely estimated.

This asymmetry reinforces the interpretation that infrastructure amplifies rather than equalizes economic capacity. While national highways appear to complement existing productive potential, state highways do not exhibit systematic effects on income. There are several likely explanations. First, state highways are often designed to connect remote or underdeveloped regions that are typically bypassed by National Highways. Such areas usually lack the complementary inputs necessary to activate meaningful productivity gains, consistent with the threshold mechanism outlined in the model. Second, placement decisions for state highways may reflect political or geographic targeting rather than economic considerations. Third, since state highways are financed and maintained by state governments, they are typically subject to more stringent fiscal and administrative constraints.

Taken together, these findings support the model's central implication: infrastructure delivers high returns only in settings where sufficient complementary capacity is already present. National highways appear to magnify spatial advantages rather than level the playing field.

Robustness. I conduct two robustness checks. First, I implement a leave-one-state-out (LOSO) analysis to confirm that results are not driven by any single region. Second, I replicate the full IV pipeline at the $0.25° \times 0.25°$ grid resolution to ensure that findings are not artifacts of district-level aggregation.

Figure 9 shows LOSO IV estimates under two specifications: Panel (a) uses the 1969 + topography instrument, and Panel (b) uses the preferred 1928 + topography design. In both cases, point estimates remain stable across state exclusions. Confidence intervals widen for smaller states, particularly under the 1928 specification, but the core result holds; no single geography drives the observed effect.

At the finer grid resolution, I re-estimate both average and heterogeneous effects. Table 11 reports over-

identification tests. While most pass comfortably, Column 2 (1928 + topography) shows a significant Sargan statistic ($p < 0.05$), indicating potential inconsistency between the instruments.

Accordingly, while I continue to use the 1928 + topography specification for average effects, I rely primarily on the topography-only instrument for interpreting heterogeneity. State Highways, meanwhile, continue to show null effects despite improved instrument strength.

Figure 10 presents quartile-wise IV estimates at the 0.25° × 0.25° resolution. While average effects are often statistically indistinct from zero—a consequence of saturated fixed effects and compressed variation—the upward gradient across income remains consistent. Replicating this analysis at the state level restores significance and sharpens the slope. This lends further support to the threshold-based interpretation.

Taken together, these findings provide strong and consistent evidence for the main hypothesis. National Highways raise income not universally, but selectively—where complementary inputs are already in place. This pattern is robust to sample exclusions, spatial resolution, and alternative identification strategies. Infrastructure, in this setting, acts not as a redistributive tool, but as an amplifier of regional advantage.

6 Conclusion

This paper studies the causal impact of highway infrastructure on local economic development in India, combining newly digitized historical road maps with high-resolution grid-level GDP data. Three main findings emerge. First, national highway (NH) density has a large and statistically significant positive effect on local income. Second, these gains are highly uneven: the benefits are concentrated in wealthier regions, while poorer areas see little measurable improvement. Third, state highways (SHs) show no robust effect once endogeneity is addressed. These patterns are consistent across specifications and supported by extensive robustness checks.

To help explain these results, I develop a simple dynamic framework in which infrastructure raises output only after a region surpasses a threshold level of complementary capacity. This reflects the reality that roads alone are insufficient without skilled labor, institutional quality, or market access. The empirical findings closely track this logic: NH effects are strong in high-income regions and negligible elsewhere. Threshold estimates imply that only the top 0.5–2% of Indian regions benefit substantially—underlining that infrastructure returns are conditional and nonlinear.

The policy implications are twofold. First, there is a fundamental equity-efficiency tradeoff: building highways in lagging regions may promote spatial inclusion but does not automatically translate into growth. Second, absorptive capacity is critical—physical infrastructure must be complemented by targeted investments that enable regions to leverage connectivity.

Historical legacies further reinforce these dynamics. Regions with stronger colonial road exposure continue to exhibit persistent economic advantages, highlighting the compounding effects of early infrastructure. It matters not just where roads are built, but when—and what surrounds them.

While the empirical setting is India, the underlying mechanisms are likely more general. Similar threshold effects are observed in other low-income contexts—for instance, in Ethiopia, where road expansion between 1997 and 2016 generated strong gains in richer regions but little or even negative growth in poorer ones (Alder et al. (2022)). This convergence in patterns suggests that infrastructure-led development may be systematically shaped by underlying complementarities. Future work could test this hypothesis using comparable spatial GDP and road data in other developing countries as they become available to comment more on the external validity of the empirical design.

Various limitations warrant mention. First, the analysis remains partial equilibrium: it does not model migration, spatial spillovers, or general equilibrium reallocation of firms and workers. While the 0.25° grid resolution and district fixed effects mitigate some spatial correlation, future work could formally incorporate spillovers using a network-based or structural approach. Second, the framework is deliberately agnostic about microfoundations—sectoral dynamics, household behavior, and welfare impacts are not explicitly modeled. These choices limit the scope for formal policy simulations or normative evaluation. Finally, while the identification strategy

leverages historical variation, additional validation is warranted. Falsification tests are planned once updated grid-level data—such as those currently under development by Rossi-Hansberg and Zhang—become available. These tests will provide further evidence on whether infrastructure placement truly precedes growth or simply reflects latent trends.

In sum, this study underscores both the promise and the limits of infrastructure as a tool for development. Highways can drive growth—but primarily in regions with the institutions, skills, and market linkages to absorb them. Absent those conditions, new roads may not level the playing field; they may merely pave over old inequalities.

References

Acemoglu, D., Johnson, S., and Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5):1369–1401.

Alder, S., Croke, K., Duhaut, A., Marty, R. A., and Vaisey, A. B. (2022). The Impact of Ethiopia's Road Investment Program on Economic Development and Land Use: Evidence from Satellite Data. Technical report, The World Bank.

Allen, T. and Atkin, D. (2022). Volatility and the gains from trade. *Econometrica*, 90(5):2053–2092.

Asher, S. and Novosad, P. (2020). Rural roads and local economic development. *American Economic Review*, 110(3):797–823.

Asturias, J., García-Santana, M., and Ramos, R. (2019). Competition and the welfare gains from transportation infrastructure: Evidence from the Golden Quadrilateral of India. *Journal of the European Economic Association*, 17(6):1881–1940.

Dell, M. (2010). The persistent effects of Peru's mining mita. *Econometrica*, 78(6):1863–1903.

Donaldson, D. (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5):899–934.

Donaldson, D. and Hornbeck, R. (2016). Railroads and American economic growth: A "market access" approach. *The Quarterly Journal of Economics*, 131(2):799–858.

Ghatak, M. (2015). Theories of poverty traps and anti-poverty policies. *The World Bank Economic Review*, 29:S77–S105.

Ghatak, M. (2017). Comment on Chapters 9 and 10. In *The Economics of Poverty Traps*, pages 383–393. University of Chicago Press.

Nunn, N. and Puga, D. (2012). Ruggedness: The blessing of bad geography in Africa. *Review of Economics and Statistics*, 94(1):20–36.

Rossi-Hansberg, E. and Zhang, J. (2025). Local GDP Estimates Around the World. Technical report, National Bureau of Economic Research.

Sachs, J., McArthur, J. W., Schmidt-Traub, G., Kruk, M., Bahadur, C., Faye, M., and McCord, G. (2004). Ending Africa's poverty trap. *Brookings papers on economic activity*, 2004(1):117–240.

A Proofs of Propositions

Proof of Proposition 1 (Existence and Uniqueness of High-Productivity Steady State)

Existence. Suppose there exists $\varepsilon > 0$ such that for all $K_i \in (\hat{K}, \hat{K} + \varepsilon)$,

$$AK_i^{\alpha} R^{\beta+\gamma} - F > \frac{\delta}{s} K_i.$$

Define the function

$$g(K_i) = s\left(AK_i^{\alpha} R^{\beta+\gamma} - F\right) - \delta K_i.$$

Step 1: g is continuous on $(0, \infty)$, as it's comprised of continuous functions.

Step 2: By assumption, $g(K_i) > 0$ for $K_i \in (\hat{K}, \hat{K} + \varepsilon)$; choose one such point K^+.

Step 3: As $K_i \to \infty$, $g(K_i) \to -\infty$ since $\alpha < 1 \Rightarrow K_i^{\alpha} = o(K_i)$.

Step 4: By the Intermediate Value Theorem, g has at least one root $\overline{K_i} > \hat{K}$ with $g(\overline{K_i}) = 0$. □

Uniqueness. To establish uniqueness, we first establish and then invoke the following lemma.

Lemma 1 (Uniqueness of Root for Strictly Concave Functions). *Let $f \in C^2(a, b)$ be strictly concave. If $f(x_0) > 0$ for some $x_0 \in (a, b)$ and $\lim_{x \to b^-} f(x) < 0$, then $f(x) = 0$ has a unique solution in (x_0, b).*

Proof. Since $f \in C^2(a, b)$, it is also continuously differentiable and continuous on (a, b). Given $f(x_0) > 0$ and $\lim_{x \to b^-} f(x) < 0$, the Intermediate Value Theorem ensures that $f(x) = 0$ has at least one solution in (x_0, b).

Suppose, for contradiction, that there exist two distinct roots $x_1 < x_2$ in (x_0, b) with $f(x_1) = f(x_2) = 0$. By strict concavity, for all $\lambda \in (0, 1)$,

$$f(\lambda x_1 + (1 - \lambda)x_2) > \lambda f(x_1) + (1 - \lambda) f(x_2) = 0,$$

so $f(x) > 0$ for all $x \in (x_1, x_2)$.

Since $\lim_{x \to b^-} f(x) < 0$, there exists $x_3 \in (x_2, b)$ such that $f(x_3) < 0$.

Now, fix $\lambda \in (0, 1)$ and define $x_\lambda = \lambda x_1 + (1 - \lambda)x_2 \in (x_1, x_2)$. Then $f(x_\lambda) > 0$, and by IVT applied to f over $[x_\lambda, x_3]$, there exists $\tilde{x}_\lambda \in (x_\lambda, x_3)$ such that $f(\tilde{x}_\lambda) = 0$.

Clearly as $\tilde{x}_\lambda > x_\lambda$ and $x_\lambda > x_1$, we have $\tilde{x}_\lambda > x_1$. We now consider the relative position of $\tilde{x}_\lambda$ and x_2. There are three cases:

Case 1. $\tilde{x}_\lambda < x_2$

Since $x_1 < \tilde{x}_\lambda < x_2$, we can express $\tilde{x}_\lambda = \mu x_1 + (1 - \mu)x_2$ for some $\mu \in (0, 1)$. Then by strict concavity:

$$f(\tilde{x}_\lambda) > \mu f(x_1) + (1 - \mu) f(x_2) = 0,$$

contradicting the fact that $f(\tilde{x}_\lambda) = 0$.

Case 2. $\tilde{x}_\lambda = x_2$

We now have exactly two roots, $x_1 < x_2$, with $f(x_1) = f(x_2) = 0$, and $f(x) > 0$ on (x_1, x_2). Consider the behavior of f' near x_1. Since $f(x_0) > 0$ and $f(x_1) = 0$, with x_1 being the first root, the function is decreasing as it approaches x_1, so:

$$\lim_{h \downarrow 0} \frac{f(x_1) - f(x_1 - h)}{h} = f'_-(x_1) < 0.$$

On the other hand, since $f(x) > 0$ on (x_1, x_2) and $f(x_2) = 0$, f must initially rise to the right of x_1, implying:

$$\lim_{h \downarrow 0} \frac{f(x_1 + h) - f(x_1)}{h} = f'_+(x_1) > 0.$$

But since $f \in C^2 \implies f \in C^1$, the derivative must be continuous, so $f'_-(x_1) = f'_+(x_1)$, which contradicts the strict sign change at x_1. Hence, this case is not possible.

Case 3. $\tilde{x}_\lambda > x_2$

Now there are three distinct roots: $x_1 < x_2 < \tilde{x}_\lambda$. Since $f \in C^2$, we get that f'' is well defined. We apply Rolle's Theorem on $[x_1, x_2]$ and $[x_2, \tilde{x}_\lambda]$ to find:

$$\exists\, \eta_1 \in (x_1, x_2),\ \eta_2 \in (x_2, \tilde{x}_\lambda) \text{ such that } f'(\eta_1) = f'(\eta_2) = 0.$$

But, the strict concavity of a twice continuously differentiable function implies $f''(x) < 0$, so f' is strictly decreasing, and cannot attain the same value (0) at two distinct points. This gives us our contradiction.

Having ruled out all possible positions of $\tilde{x}_\lambda$, we conclude the assumption of two distinct roots is false. Hence, the root is unique. □

Apply the lemma 1 to $g(K_i)$, which is C^2 and strictly concave for $\alpha \in (0, 1)$:

$$g''(K_i) = sA\alpha(\alpha - 1)K_i^{\alpha-2}R^{\beta+\gamma} < 0.$$

By assumption, $g(K^+) > 0$, and $\lim_{K_i \to \infty} g(K_i) < 0$. Hence, the steady state $\overline{K_i} \in (\hat{K}, \infty)$ is unique. □

Discussion: Multiple Steady States under Weaker Feasibility conditions

If accumulation is only feasible at some $K > \hat{K}$ (and not necessarily immediately after the threshold), then $g(K)$ may admit two distinct roots. This yields three steady states: a low stable $\underline{K_i}$, an unstable middle $\tilde{K}_i \in (\hat{K}, \overline{K_i})$, and a high stable $\overline{K_i}$. Convergence to $\tilde{K}_i$ is not generic and is a zero-measure event in continuous space (as it requires starting at $\tilde{K}_i$), hence not empirically relevant. In this case, however, simply crossing the threshold $\hat{K}$ is not enough to converge to $\overline{K_i}$, instead, region i must have initial capital endowment above $\tilde{K}_i$ to ensure convergence to the high steady state. If $K_i < \tilde{K}_i$, region i will converge to $\underline{K_i}$.

This motivates the stronger condition of immediate feasibility used in Proposition 1, which ensures convergence to the high steady state as long as $K_i > \hat{K}$.

Figure 5 illustrates this possibility.

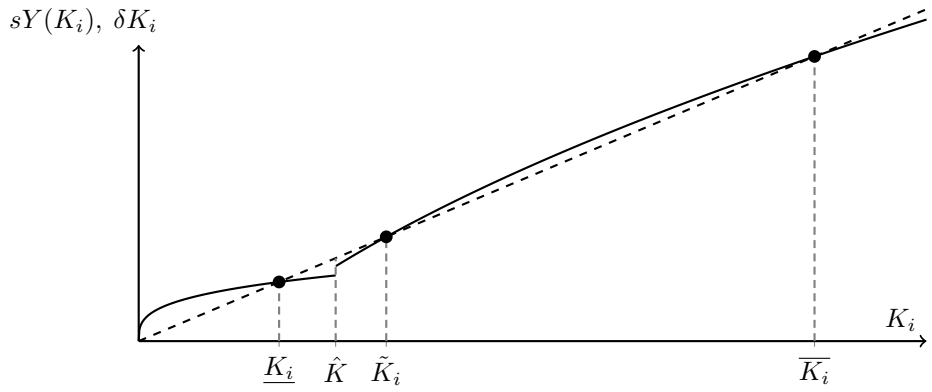

Figure 5: Multiple Steady States under Weaker Feasibility Condition

Proof of Proposition 2 (Heterogeneous Infrastructure Returns)

We compute:

$$\left.\frac{\partial Y}{\partial R}\right|_{K=\underline{K_i}} = A\beta \underline{K_i}^{\alpha} R^{\beta-1}, \quad \left.\frac{\partial Y}{\partial R}\right|_{K=\overline{K_i}} = A(\beta+\gamma)\overline{K_i}^{\alpha} R^{\beta+\gamma-1}.$$

Taking their ratio:

$$\frac{\partial Y/\partial R|_{\overline{K_i}}}{\partial Y/\partial R|_{\underline{K_i}}} = \frac{\beta+\gamma}{\beta} \cdot \left(\frac{\overline{K_i}}{\underline{K_i}}\right)^{\alpha} \cdot R^{\gamma} > 1,$$

since $\overline{K_i} > \underline{K_i}$, $\gamma > 0$, and $R \geq 1$. Hence, infrastructure returns are strictly higher in high-capital regions. □

B Main Tables and Figures

B.1 OLS Results

Table 1: OLS Estimates: Road Density and Log GDP per Capita

	(1)	(2)	(3)
log(AWM–NH) density	0.002***		
	(0.000)		
log(GQ–NS) density	0.004***		
	(0.001)		
log(GQ–EW) density	0.000		
	(0.001)		
log(SH) density			0.000
			(0.001)
log(NH) density		0.002***	
		(0.001)	
Num. Obs	1134	1134	1134
R^2	0.737	0.730	0.728

Notes: Dep. variable is log(GDP per Capita, 2012)
All models include state fixed effects and standard errors clustered at the state level.
AWM–NH denotes all-weather motorable National Highways.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

B.2 Endogenous Placement

Table 2: OLS Estimates: NH Growth (2011–21) on 2012 Log GDP per Capita

	NH Growth 2011–21
log(GDP per capita, 2012)	-2.700^{+}
	(1.438)
Num. Obs	1114
R^2	0.085

Notes: Dep. variable is Δlog(NH density) between 2011 and 2021.
State fixed effects included and standard errors clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

B.3 PCA Results for Topography Instrument

Table 3: Topographic PCA Summary Statistics

(a) Variance Explained by Each Principal Component

Metric	PC1	PC2	PC3
Standard Deviation	1.547	0.678	0.385
Proportion of Variance	0.798	0.153	0.049
Cumulative Proportion	0.798	0.951	1.000

(b) Variable Loadings on Each Principal Component

Variable	PC1	PC2	PC3
Mean Elevation	−0.614	−0.136	−0.778
Variance of Elevation	−0.574	−0.599	0.558
Mean TRI	−0.541	0.789	0.290

B.4 Evidence of Persistence

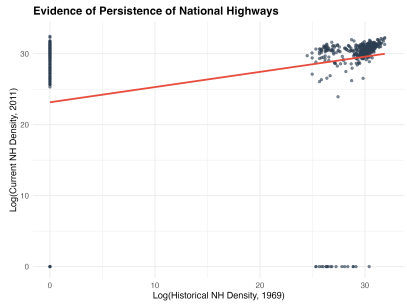

(a) NH Persistence, 1969–2011

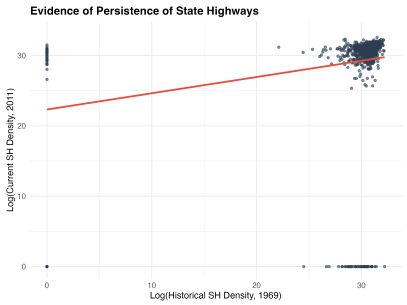

(b) SH Persistence, 1969–2011

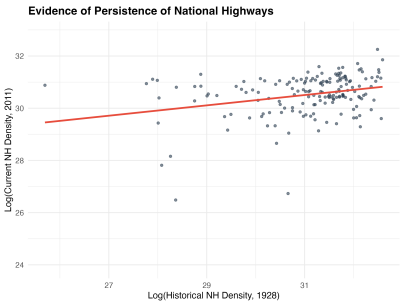

(c) NH Persistence, 1928–2011

Figure 6: Persistence in Road Infrastructure

B.5 First Stage IV Results

Table 4: First-Stage Regressions: National Highways (log density)

	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Topo PC1	1.516*		0.225		1.093*
	(0.634)		(0.518)		(0.377)
log(NH density, 1969)		0.225***	0.222***		
		(0.030)	(0.029)		
log(NH density, 1928)				0.223*	0.182*
				(0.084)	(0.075)
Num. Obs	1113	1017	1015	169	169
R^2	0.068	0.166	0.162	0.0939	0.203

Notes: Dep. variable is log(NH density, 2011).
All specifications include state fixed effects and standard errors clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Table 5: First-Stage Regressions: State Highways (log density)

	(1)	**(2)**	**(3)**
Topo PC1	0.046		-0.737^{+}
	(0.650)		(0.400)
log(SH density, 1969)		0.235*	0.236*
		(0.098)	(0.098)
Num. Obs	1113	1017	1015
R^2	0.114	0.166	0.169

Notes: Dep. variable is log(SH density, 2011).
All specifications include state fixed effects.
Standard errors clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

B.6 Second Stage IV Results

Table 6: Second-Stage IV Regressions: National Highways (log density)

	(1) Topo Only	(2) 1969 Only	(3) 1969 + Topo	(4) 1928 Only	(5) 1928 + Topo
NH (log density)	0.018*	0.009***	0.009***	0.144+	0.114+
	(0.008)	(0.002)	(0.002)	(0.085)	(0.061)
Num. Obs	1113	1017	1015	169	169
R^2	0.607	0.733	0.732	0.687	0.697
First Stage F stat	8.184**	119.712***	58.734***	19.9***	23.3***
Wu-Hausman Stat	4.047*	12.324***	12.062***	0.816	0.573
Sargan Stat	—	—	0.002	—	0.314

Notes: Dep. variable is log(GDP per capita) in 2012. Instruments indicated by column headers.
All specifications include state fixed effects and standard errors clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Table 7: Second-Stage IV Regressions: State Highways (log density)

	(1) Topography Only	(2) 1969 Only	(3) 1969 and Topography
SH (log density)	0.602	0.008	0.006
	(8.702)	(0.007)	(0.007)
Num. Obs	1113	1017	1015
R^2	-89.444	0.746	0.749
First Stage F stat	0.015	42.612***	23.182***
Wu–Hausman Stat	5.056*	1.774	1.288
Sargan Stat	—	—	0.769

Notes: Dep. variable is log(GDP per capita) in 2012. Instruments indicated by column headers.
All specifications include state fixed effects and standard errors clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

B.7 Reverse Causality Checks on the IV Results

Table 8: Reverse Causality Checks: Predicting Fitted NH Density from Income

	(1) Joint IV (1969)	(2) Topography IV	(3) Joint IV (1928)
log(GDP per capita, 2012)	1.923***	0.043	0.219
	(0.401)	(0.061)	(0.132)
Num. Obs	1023	1217	169
R^2	0.419	0.870	0.195

Notes: Dep. variable is the fitted value of NH density from the first stage equation.
No significant relationship suggests the absence of reverse causality.
All models include state fixed effects and standard errors clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

B.8 Heterogeneous Returns to National Highways by Income

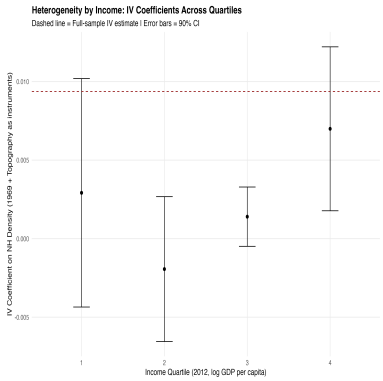

(a) Heterogeneous Returns, using Topography and 1969 Roads

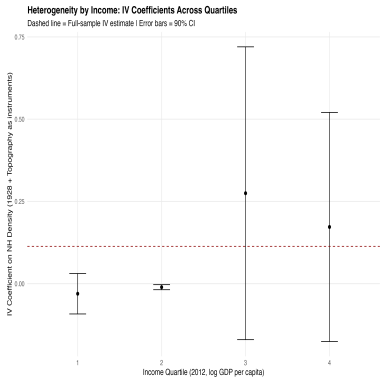

(b) Heterogeneous Returns, using Topography and 1928 Roads

Figure 7: Heterogeneous Returns at the $0.5 \times 0.5°$ resolution

Table 9: Interaction IV Estimates Using 1928 Historical Roads

	(1) **Q4 only**	**(2)** **Q3 + Q4**
NH Density (2011)	0.053	0.043
	(0.045)	(0.041)
NH Density (2011) × Rich Group	0.016***	0.017***
	(0.004)	(0.005)
Num. Obs	169	169
R^2	0.825	0.797

Notes: Dependent variable is log GDP per capita (2012).
All models include state fixed effects. Standard errors are clustered at the state level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

B.9 Estimating the Threshold Level of Income

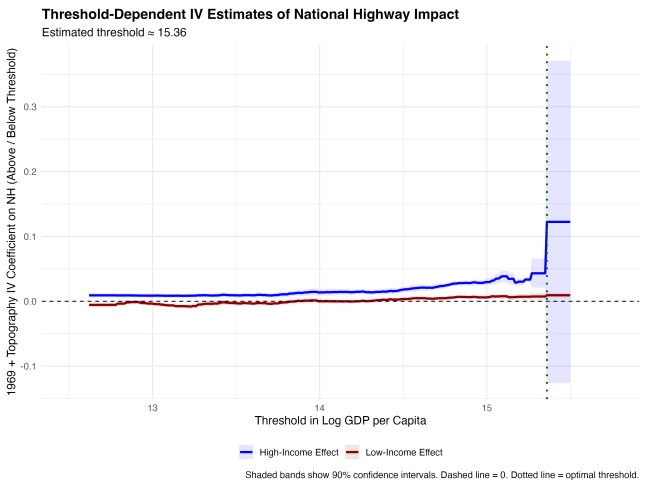

Figure 8: Threshold IV Estimates Using 1969 Roads and Topography

B.10 Robustness Checks

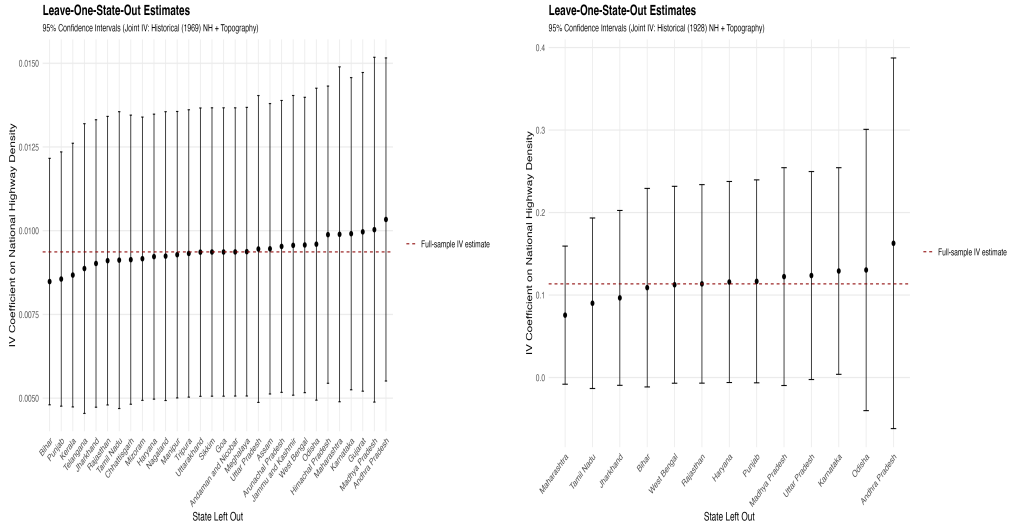

(a) Leave-One-State-Out Results using 1969 + Topography Instrument

(b) Leave-One-State-Out Results using 1928 + Topography Instrument

Figure 9: Robustness checks - Leave One State Out estimates

IV Results at the 0.25° × 0.25° resolution

Table 10: First-Stage Regressions (0.25°): Joint Instrumentation

	(1) NH (2011)	(2) NH (2011)	(3) NH (2011)	(4) SH (2011)
Topographic Suitability (PC1)	1.330^{+}	1.892***	1.708**	0.647
	(0.710)	(0.509)	(0.571)	(0.450)
Log NH Density (1969)	0.400***			
	(0.026)			
Log SH Density (1969)				0.286***
				(0.034)
Log NH Density (1928)		0.129^{+}		
		(0.076)		
Num.Obs.	3024	342	3803	3024
R^2	0.460	0.367	0.331	0.455

All specifications include district fixed effects and standard errors clustered at the district level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Table 11: Second-Stage IV Regressions (0.25°): Joint Instrumentation

	(1) 1969 and Topo	(2) 1928 and Topo	(3) Topo	(4) 1969 and Topo
Log NH Density (2SLS/IV)	0.005***	0.044	0.005	
	(0.001)	(0.034)	(0.005)	
Log SH Density (2SLS)				0.001
				(0.002)
Num.Obs.	3024	342	3803	3024
R^2	0.832	0.873	0.821	0.837
First Stage F stat	262.401***	25.121***	15.821***	90.903***
Wu Hausman Stat	17.504***	0.582	0.566	0.221
Sargan Stat	0.209	5.4555*	—	0.387

Notes: Dep. variable is log(GDP per capita) in 2012. Instruments used are indicated by column headers.
All specifications include district fixed effects and standard errors clustered at the district level.
$^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

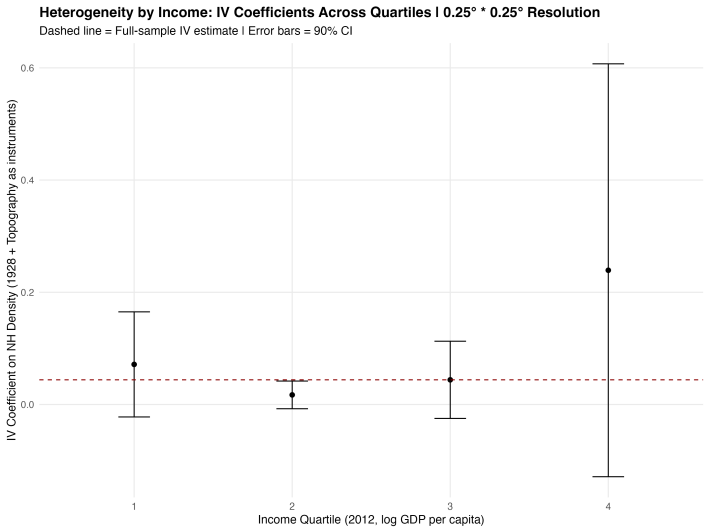

(a) IV estimates using 1928 + Topography

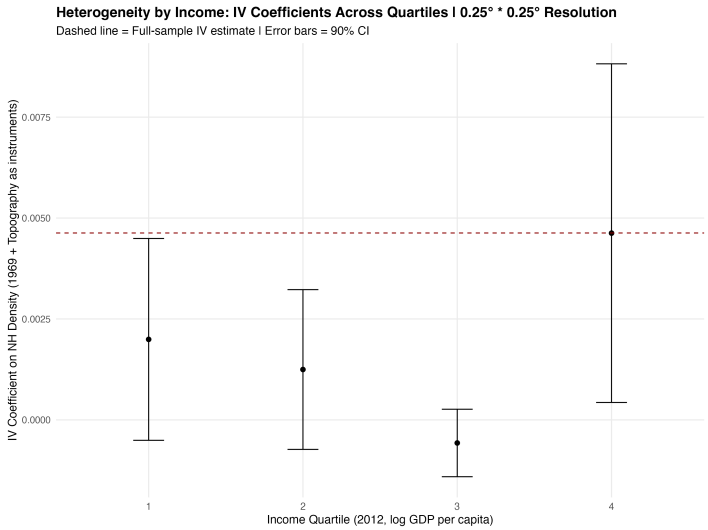

(b) IV estimates using 1969 + Topography

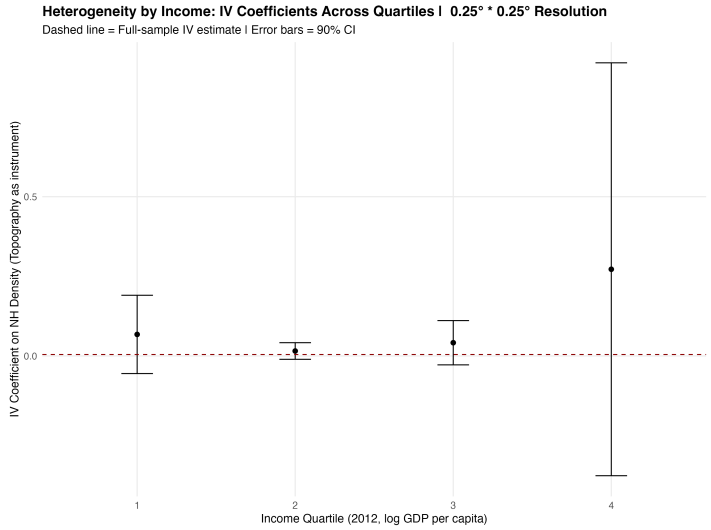

(c) IV estimates using Topography

Figure 10: Heterogeneity at the 0.25°×0.25° Resolution
Heterogeneous IV estimates using different instruments. Dashed lines denote full sample IV estimates.
All specifications include district fixed effects and standard errors clustered at the district level.

C Historical Road Network and Topographic Instrument Construction

C.1 Image Classification and Road Digitization (MATLAB)

The classification of historical road maps was carried out using a supervised image classification approach implemented in MATLAB. Starting with high-resolution scans of Survey of India road maps (e.g., 2011), the raw raster images were read into MATLAB, and road classes were manually labeled using polygon-based selections. Each road class (e.g., National Highways, Golden Quadrilateral segments, and State Highways) was assigned a distinct numerical code. Around 20–30 polygons were drawn per class to generate a sufficiently rich training set of RGB pixel values. These labeled samples were used to train a Random Forest classifier (`TreeBagger`), which was then applied to the entire image in chunks of 25 million pixels using parallel processing to avoid memory overflow.

To clean the classified image, a two-step filter was implemented. First, a 3×3 neighborhood consensus filter corrected small misclassifications. Second, isolated specks under 8 pixels were removed using morphological operations. This produced a cleaned classification raster with consistent labels for each road class.

To further eliminate map artifacts—such as digits, labels, or hollow place names—a second shape-based classifier was trained. Using polygon selections, features like area, perimeter, eccentricity, solidity, and aspect ratio were extracted for both roads and textual elements. These were fed into another Random Forest, which successfully separated textual artifacts from road segments. The cleaned output was then skeletonized (1-pixel thinning) and vectorized to generate shapefiles for each road class.

C.2 Spatial Harmonization and Density Estimation (R)

The vectorized road shapefiles were imported into R and intersected with grid-level GDP data. The GDP dataset was first filtered for India (2012) and merged with a cleaned grid shapefile that included `grid_id` and cell identifiers. For road density estimation, the road layers were first re-projected to EPSG:3857 and intersected with the India grid shapefile using `qgisprocess` in R. Road lengths by class were computed per grid using the `native:fieldcalculator` algorithm, the `qgis:statisticsbycategories` algorithm was used to aggregate grid-wise totals.

These road lengths were normalized by grid area (in km²) to compute class-specific road densities. The final variables include densities for National Highways (NH), two Golden Quadrilateral segments (NS and EW), and State Highways (SH). All densities were monotonically transformed by multiplying by 10^{12} and adding 1 before being logged for use in regressions, following the same transformation applied to the GDP variable.

C.3 Constructing the Topographic Instrument

To instrument for road placement, a terrain-based instrument was developed by merging grid-level elevation data (mean and variance) with terrain ruggedness index (TRI), sourced from the Shuttle Radar Topography Mission (SRTM). The elevation and TRI rasters were clipped to India's boundaries and spatially averaged within each grid. To summarize these correlated terrain features, a principal component analysis (PCA) was applied to the standardized values of `elev_mean`, `elev_variance`, and `tri_mean`.

The first principal component (PC1), which captures over 80% of the total variation, was used as the topographic instrument (`topo_PC1`). This variable was merged into the main grid-level dataset and used in the IV regressions. While PC1 is a reasonably strong instrument for national highway density, it is a weak predictor of state highway coverage once state fixed effects are introduced—reinforcing the need for a historical instrument (e.g., 1969 road map) in those regressions.

D Elevation and Road Maps

D.1 Elevation Map

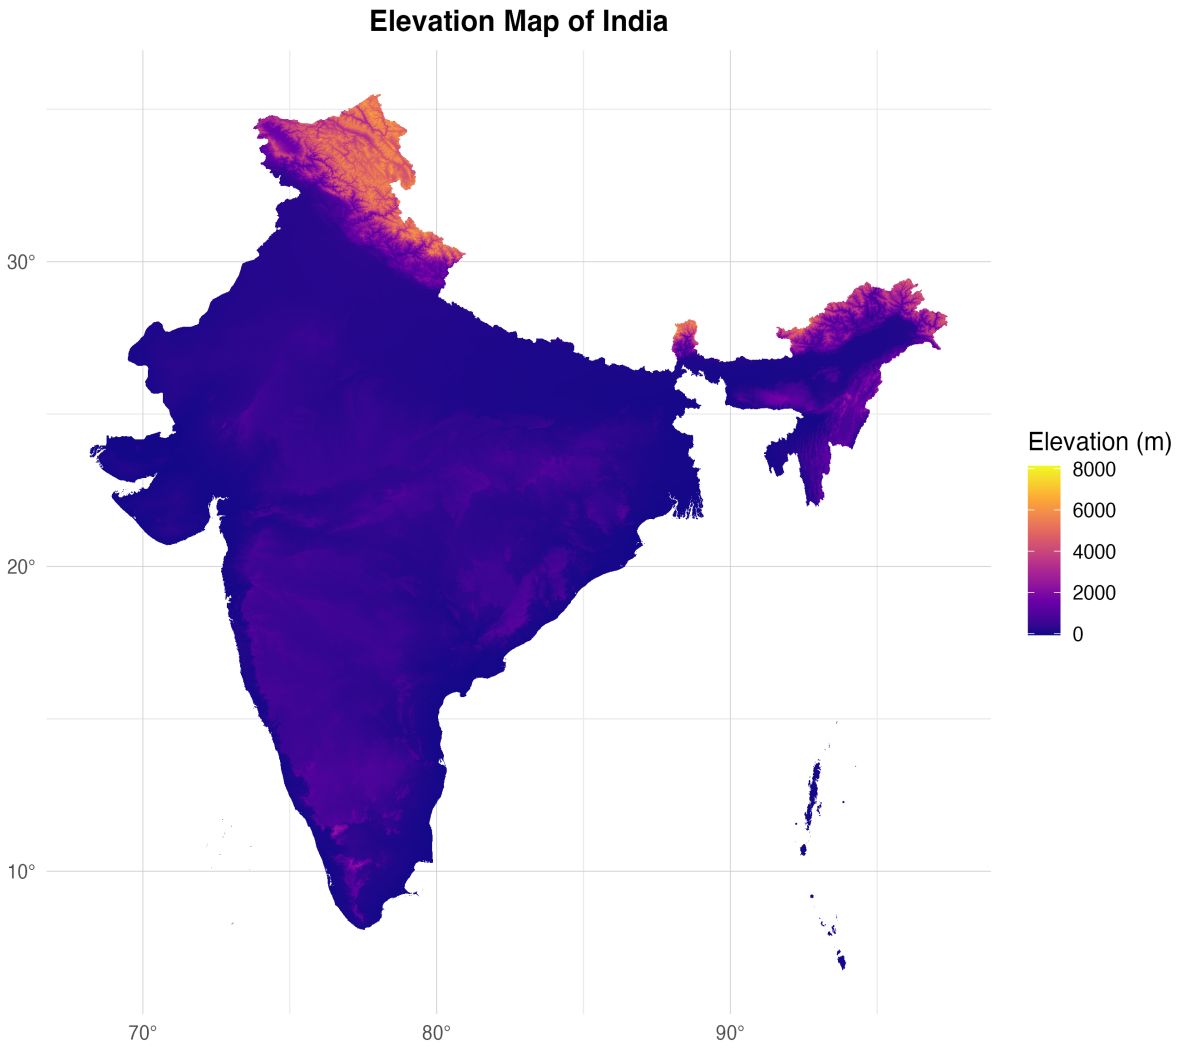

Figure 11: Elevation Map of India

D.2 Historical & Contemporary Road Maps

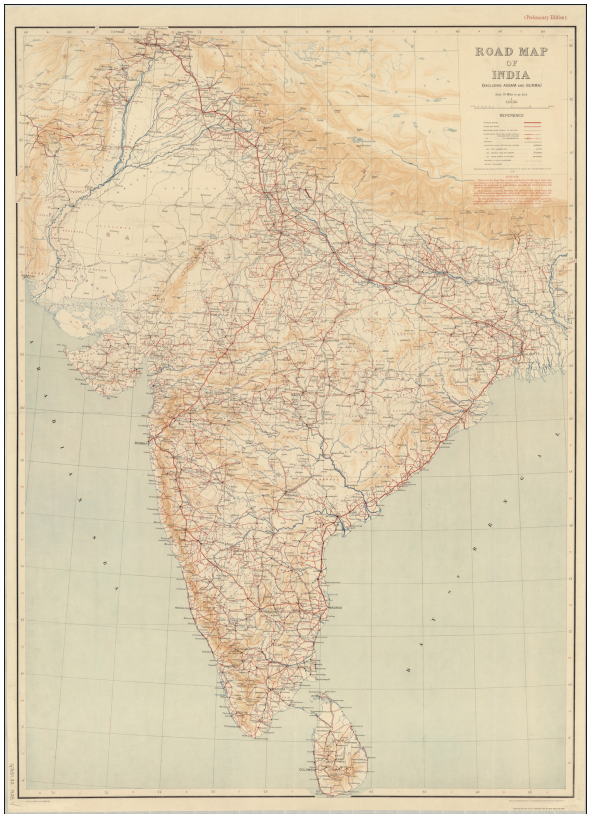

(a) 1928 Map

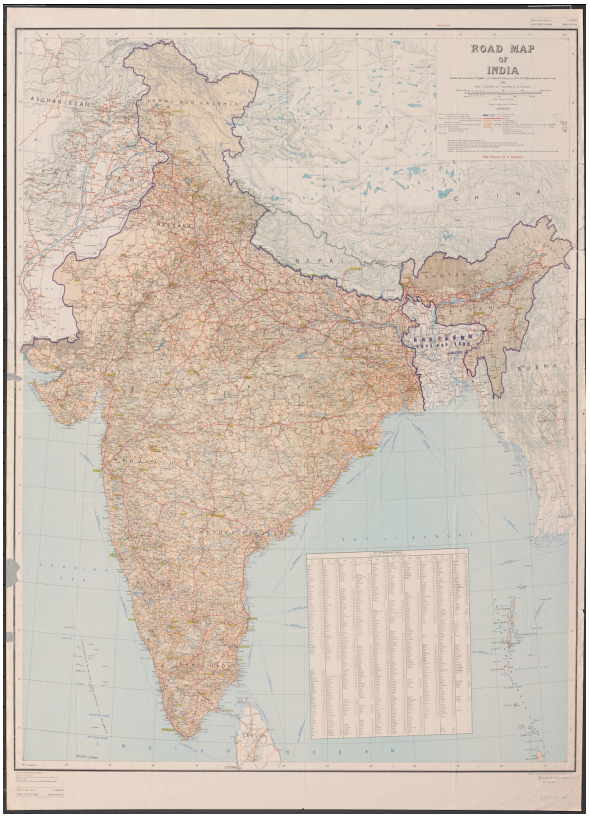

(b) 1969 Map

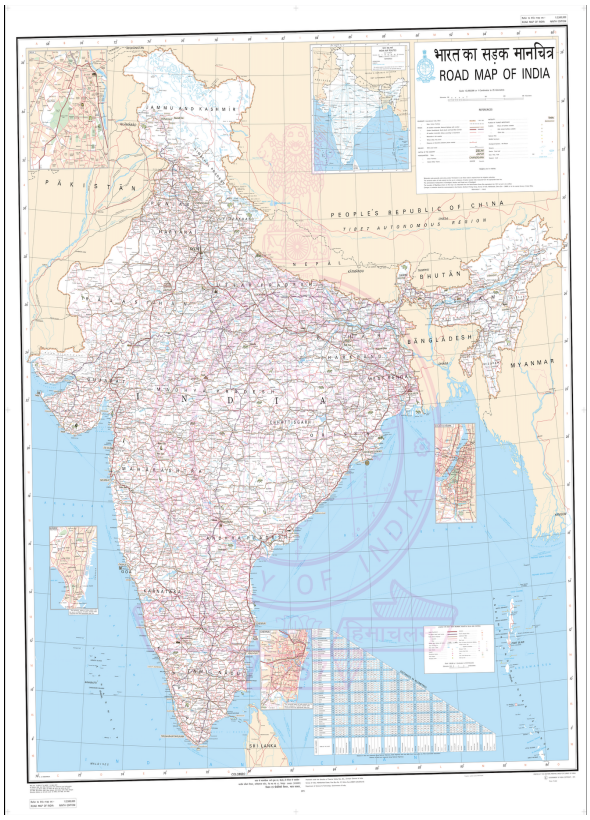

(c) 2011 Map

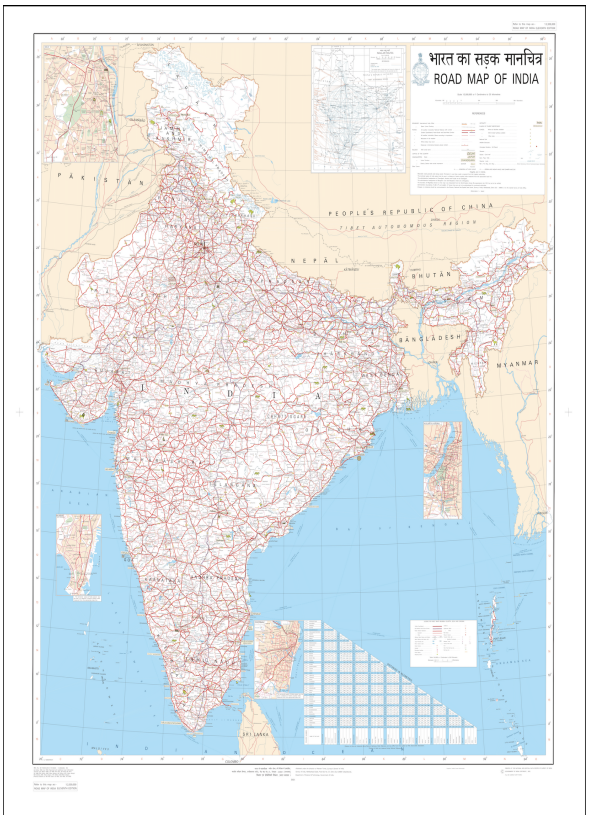

(d) 2021 Map

Figure 12: Raw Historical and Contemporary Road Maps of India (1928–2021)