

The Ecological Impact of Transportation Infrastructure*

Sam Asher[†]

Teevrat Garg[‡]

Paul Novosad[§]

March 2019

Abstract

There is a long-standing debate over whether new roads unavoidably lead to environmental damage, especially forest loss, but causal identification has been elusive. Using multiple causal identification strategies, we study the construction of new rural roads to over 100,000 villages and the upgrading of 10,000 kilometers of national highways in India. The new rural roads had precise zero effects on local deforestation. In contrast, the highway upgrades caused substantial forest loss, which appears to be driven by increased timber demand along the transportation corridors. In terms of forests, last mile connectivity had a negligible environmental cost, while expansion of major corridors had important environmental impacts.

*We are grateful for feedback on earlier versions of this paper from Prashant Bharadwaj, Jenn Burney, Jonah Busch, Eric Edmonds, Paul Ferraro, Gordon Hanson, Maulik Jagnani, Erzo Luttmer, Erin Mansur, Gordon McCord, and seminar participants at UC San Diego, UC Irvine, San Diego State University, Claremont McKenna College, George Washington University, the 2018 Pacific Development Economics Conference, the 2018 Annual Meetings of the Agricultural and Applied Economics Association and the 2017 CU Boulder Environmental and Energy Economics Workshop. Ryu Matsuura provided excellent research assistance. Garg acknowledges funding from the Center for Global Transformation at UCSD, ESRC Centre for Climate Change Economics and Policy, the Grantham Foundation for the Protection of the Environment.

[†]World Bank, sasher@worldbank.org

[‡]UC San Diego, teevrat@ucsd.edu

[§]Dartmouth College, paul.novosad@dartmouth.edu

I Introduction

Does human economic progress have an unavoidable environmental cost? This is a central question for policymakers pursuing sustainable development and has been a long-standing debate in both the conservation and the economics literature (Arrow et al., 1995; Grossman and Krueger, 1995; Stern, Common and Barbier, 1996; Andreoni and Levinson, 2001; Foster and Rosenzweig, 2003; Dasgupta, 2007; Alix-Garcia et al., 2013). A key pillar of economic development is large-scale investment in transportation infrastructure that reduces the costs of moving goods and people across space. Concern has been expressed about the potential environmental cost of such investments, and of increased trade more generally (Copeland and Taylor, 1994; Antweiler, Copeland and Taylor, 2001; Copeland and Taylor, 2004; Frankel and Rose, 2005), but researchers have struggled to identify causal estimates of the impact of transportation infrastructure on local environmental quality.

The most omnipresent of transportation investments are roads. We focus on the impact of road construction and expansion on forest loss as it is among the primary environmental concerns associated with new road construction. Forest cover loss is globally and locally important, generating global greenhouse emissions (IPCC, 2007; Jayachandran et al., 2017) and local health externalities (Bauch et al., 2015; Garg, 2017). Analysis by the IPCC suggests that restoring and protecting forests could yield almost a sixth of the emissions mitigation required to prevent runaway climate change by 2030 (IPCC, 2019).

Because of the high cost and high expected return of roads, their placement typically depends on various economic and political factors, making causal identification of their impacts difficult. For example, new roads may be targeted to regions with expanding agricultural land use; these roads may be a response to activities that are already causing forest cover reduction, making it difficult to isolate the direct impact of the roads. While many earlier studies have documented changes in forest cover following the construction of new roads, none have addressed the endogeneity of road placement beyond the inclusion of control variables and in a few cases, location fixed effects. Further, most of these studies have focused on

large highways built into the Amazon rainforest (Pfaff, 1999; Pfaff et al., 2007; Weinhold and Reis, 2008); while these highways are important in terms of potential deforestation, their impacts are of uncertain relevance for the set of potential rural roads and highways that policy-makers in developing countries are considering today. The majority of road projects in the decades ahead are likely to be last-mile roads to people not currently connected to the road network and upgrades of existing transportation corridors into modern highways.

In this paper, we take advantage of a validated satellite-based measure of forest cover (Vegetation Continuous Fields or VCF), which makes it possible to study the impacts of two large-scale transportation projects in India. The first of these was an initiative to upgrade two major transportation corridors: the 6000 km “Golden Quadrilateral” network (GQ) connecting the country’s four largest cities, and the comparably-sized “North-South and East-West” network (NS-EW) connecting the country’s four cardinal endpoints in a cross. Both corridors were already used for cross-city transportation before 2000, but over the following fifteen years they were upgraded into world class divided highways. The second project was a rural road construction program, under which over 100,000 new paved rural feeder roads were built, ten kilometers in length on average, providing new connections to over 100 million rural residents. Each project has exceeded ten billion dollars in cost to date and has caused a significant reallocation of local economic activity (Asher and Novosad, 2018; Ghani, Goswami and Kerr, 2016).

Theoretically, the effect of road investments on local forest cover can be positive or negative. New roads can increase forest cover loss by: (i) providing external markets for forest resources, especially timber and firewood; (ii) providing external markets for agricultural products, motivating extensification of agriculture into forested land; and (iii) increasing the value of land for settlement and industry, resulting in forest clearing. On the other hand, paved roads could also reduce forest cover loss by (i) improving local household and industry access to substitutes for local forest resources, especially firewood; (ii) providing access to external output and labor markets, lowering the relative returns to clearing forests

for agricultural land as well as to harvesting other forest products such as firewood. Given the substantially different nature of rural feeder roads and national highways, we can also expect the importance of any of these channels to vary by the type of road.

To evaluate the impact of rural roads, we first use a regression discontinuity approach, exploiting an implementation rule that discontinuously raised the probability of road construction in villages with population above an arbitrary threshold. Second, we use a difference-in-differences specification that exploits the exact timing of road construction. Both approaches show zero effects of new roads on forest cover. The estimates are precise; we can rule out gains larger than 0.6% and losses greater than 0.2% in forest cover up to five years after roads are completed. Further, we find zero effects for sample subgroups where we might expect losses to be greater, such as villages with greater baseline forest cover or with very poor or forest-dependent residents. We also find zero change in household firewood use in treated villages. We do identify marginal (0.5%) reductions in forest cover during the road construction period; these reductions are reversed soon after roads are completed, but there is no evidence that forest cover continues to rise. We show that ignoring these construction period effects could lead to biased impact estimates. These roads have no effect on forest cover in spite of significantly altering economic opportunities for people in villages (Asher and Novosad, 2018; Adukia, Asher and Novosad, 2019).

Causal identification for impacts of highways is much more difficult than for rural roads, because in almost all cases, new highways are small in number and are built along existing transportation corridors. We take the approach of comparing changes in forest cover in areas that are near and that are far from the new highways. While we do not have data covering the period before the construction of the Golden Quadrilateral, the North-South/East-West highway route provides a plausible counterfactual, in that it is a highway of comparable size and importance that was announced simultaneously and on a similar construction schedule to the GQ, but its construction was pushed back by approximately eight years due to bureaucratic delays. Ghani, Goswami and Kerr (2016) take a similar approach in comparing

these two networks to study the impacts of the GQ on manufacturing activity.¹

In sharp contrast to rural roads, we find that the highway upgrades have had substantial negative effects on forest cover. Following construction of the GQ, we find a 20% decline in forest cover in a 100 kilometer band around the highway, an effect that persists for at least eight years. We find no change in forest cover along the NS-EW corridor until construction accelerates in 2008, at which point we also observe local forest cover loss. The timing of relative forest loss around the construction of each corridor supports a causal interpretation of these estimates. Because forest cover in India is rising on average during the sample period, these are net effects on forest cover, combining increases in deforestation and reductions in afforestation.

These highways appear to have depleted forest cover by increasing timber demand in their vicinity, which has wide ranging effects into the hinterlands of the transport corridors. Following the construction of the GQ, we find a substantial upward trend break in employment in proximate firms that use timber and wood as primary inputs, as well as employment in logging firms. Additional tests reject the competing mechanisms; there are no increases in agricultural land use or changes in local firewood consumption along the highway corridor.

This paper makes two central contributions. First, we generate the first causal estimates of the impact of large scale transportation infrastructure investments on natural resource depletion.² In so doing, we contribute to a long literature on the trade-offs and synergies between economic development and environmental conservation.³

¹On the impacts of the Golden Quadrilateral on firms in India, see also Datta (2012) and Khanna (2016).

²Many studies describe cross-sectional relationships between roads and forest cover or forest loss (Chomitz and Gray, 1996; Angelsen and Kaimowitz, 1999; Pfaff, 1999; Cropper et al., 2001; Geist and Lambin, 2002; Deng et al., 2011; Barber et al., 2014; Li et al., 2015; Dasgupta and Wheeler, 2016). A small number of studies examine forest loss in areas with new roads but do not address the endogeneity of road placement (Pfaff et al., 2007; Weinhold and Reis, 2008). The closest study to ours is ongoing work by Kaczan (2017), who uses a difference-in-differences design similar to our first strategy (but does not look at highways), finding that India's new rural roads marginally increased forest cover. The differences may arise because Kaczan (2017) does not distinguish between construction and post-construction periods, and includes villages that never receive roads as part of the control group. We show in Section IV that both of these choices may lead to biased treatment effects.

³On the general relationship between economic development and the environment, see Den Butter and Verbruggen (1994), Arrow et al. (1995), Grossman and Krueger (1995), Stern, Common and Barbier (1996), Andreoni and Levinson (2001), Dasgupta et al. (2002), Foster and Rosenzweig (2003) and Stern (2004).

Second, this is the first paper to show that the impact of roads on deforestation is a function of which markets are being connected by those roads. Last-mile rural roads provide connectivity to small local markets, facilitating exits from agriculture but without significantly changing industry’s access to forest products (Asher and Novosad, 2018). In contrast, highways dramatically change the geographic distribution of industry (Ghani, Goswami and Kerr, 2016); in India at least, this appears to have substantial environmental consequences.

Our estimates are particularly relevant as the infrastructure agenda in Sub-Saharan Africa and South and Southeast Asia is likely to prioritize exactly the kinds of infrastructure investments that we study here — new feeder roads and expansion of existing corridors — as opposed to the large highways through virgin rainforest that have been the subject of much of the earlier work on roads and deforestation. China’s Belt and Road Initiative is the signature example, which aims to promote the construction of large scale highway corridors across Southeast and Central Asia, most of which are expansions of existing roadways (Reed and Trubetskoy, 2018). In sub-Saharan Africa, where only 30% of rural people live within two kilometers of a road, last mile access is a major policy priority (Roberts, Shyam and Rastogi, 2006).

Finally, we raise an important methodological issue in the literature on estimating impacts of infrastructure. Large-scale infrastructure often takes many years to build and involves significant land clearing and economic activity during the construction process. In both our examination of highways and of rural roads, we find that forest loss begins during the construction period; in either case, estimates based strictly on the timing of infrastructure completion would underestimate the environmental impact of roads.

The rest of the paper is organized as follows. The next section describes India’s rural road and highway construction programs. In Section III, we describe the data on forest cover and

On deforestation specifically, see Koop and Tole (1999), Burgess et al. (2012), Alix-Garcia et al. (2013), and Jayachandran et al. (2017). Assunção et al. (2017) provide causal evidence that rural *electrification* mitigated forest loss in Brazil. For an exhaustive review on drivers of deforestation, see Ferretti-Gallon and Busch (2014). For a literature review on impacts of highways and rural roads on outcomes other than the environment, see Asher and Novosad (2018).

roads, as well as other secondary datasets used in our analysis. Section [IV](#) presents empirical strategy and results describing the impact of rural roads on deforestation. Section [V](#) presents the empirical strategy and impacts of highway expansions, and Section [VI](#) concludes.

II Background: Road Construction Programs in India

In 1999 and 2000, the Government of India launched two major road construction programs — one aimed at upgrading several national highway corridors and the other at connecting the remainder of India’s population to the road network. Together, these programs marked the largest expansion of road infrastructure in Indian history and came at a joint cost exceeding \$50 billion. This section provides background information on both road construction programs.

II.A Rural Roads

In 2000, the Indian government launched the Pradhan Mantri Gram Sadak Yojana (PMGSY), or the Prime Minister’s Village Roads Scheme. The primary objective of the program was to provide new paved roads to previously unconnected villages, although in practice this also involved upgrading low quality roads in already connected villages. By 2015, over 400,000 kilometers of new roads were built, providing access to the national road network to over 100 million rural people in over 100,000 villages. Over 70% of new rural roads were routes that terminated in villages.

Rural road construction began toward the end of 2001 and was continuing steadily through the end of the sample period in 2014 (See Appendix Figure [A1](#)). Villages were selected for roads based on a set of guidelines issued by a national government body, the National Rural Roads Development Authority. Notably, the program prioritized construction of roads to larger villages; district-level implementation plans were to first target all villages with populations greater than 1000, followed by villages with population greater than 500, and finally those with population greater than 250.⁴

⁴Strictly speaking, the allocation was based on habitation population rather than village population. A habitation is a smaller unit of aggregation than the village; there are between one and three habitations in

The rules were applied on a state-by-state basis, allowing states to move from one threshold to another on their own timelines. In practice, there were several other prioritization guidelines and political patronage undoubtedly played a role, so that a village’s population relative to the threshold significantly influenced its likelihood of receiving a road but was not definitive. For instance, smaller villages could be connected if they were along the least-cost path between larger prioritized villages, and proximate villages could combine their populations to attain the eligibility thresholds. For more details, see Asher and Novosad (2018) and National Rural Roads Development Agency (2005).

II.B National Highways

In 1999, the Indian government announced a plan to modernize its major highways, the National Highways Development Project. The first component of the project was the upgrading and widening of the Golden Quadrilateral highway corridor (henceforth, GQ), so named because it connected the four major cities in India: New Delhi, Mumbai, Chennai and Kolkata. The second component was a similar upgrading of the the North-South and East-West corridor (NS-EW), which would connect the furthest corners of the country from Srinagar in the north to Kanyakumari in the south, and from Porbandar in the west to Silchar in the east. Panel A of Figure 1 shows both highway corridors along with the major cities that were connected by them.

While the GQ and NS-EW projects were commissioned around the same time, the government prioritized the implementation of the GQ and construction of the NS-EW was substantially delayed. Construction on the GQ began in 2001; 80% was completed by 2004 and 95% by 2006. In contrast, by 2006 only 10% of the NS-EW corridor was completed, almost half of which was a set of highways which were shared with the GQ (Ghani, Goswami and Kerr, 2016). By 2010, 72% of the NS-EW was completed, and 90% was completed by 2015. The delay in the construction of the NS-EW allows us to use the NS-EW corridor

each village. In practice, habitation populations were pooled to the village level in many cases (see below). We aggregate to the village level because neither additional data nor maps are available at the habitation level.

as a counterfactual for changes in forest cover in the GQ corridor during and immediately following substantial completion of the GQ.

Before these highways were widened and upgraded, the GQ and NS-EW routes were already significant transportation corridors, but their road quality and congestion were highly variable. The upgrading of these networks dramatically improved their quality and reliability; these were the first major long-distance divided highway networks to be developed in India. The construction of the GQ changed national supply networks and led to a substantial reallocation of manufacturing firms into the GQ corridor (Datta, 2012; Khanna, 2016; Ghani, Goswami and Kerr, 2016). The economic impact of the NS-EW corridor has so far been little studied due to its later completion date.

III Data

To estimate the effects of new roads on forest cover, we combine five different national data sources. We use a validated high resolution satellite-based measure of forest cover. Data on rural roads come from the administrative implementation data generated by the rural road construction program, and geographic data on new major highway networks come from national highway maps. While these datasets form the basis of our core specifications, we also use data from the 1991, 2001 and 2011 Population Censuses and 3rd through 6th rounds of the Economic Census to control for location characteristics and explore mechanisms of treatment effects. All of these are census datasets that describe the entire population of India and are geocoded to the village, town and subdistrict levels. This section describes the details of how we prepare and combine all of these datasets. Table 1 shows summary statistics for all variables used.

III.A Forest Cover

Detailed and reliable administrative records on forest cover and deforestation rarely exist, especially in developing countries. Instead, we obtain high resolution time series estimates of forest cover using a standardized publicly-available satellite-based dataset. Vegetation

Continuous Fields (VCF) is available at 250m resolution and provides annual tree cover from 2000–2014 in the form of the percentage of each pixel under forest cover (Townshend et al., 2011). For our primary specification, we define forest cover as the total log pixel value plus one in a given geographic area.⁵ Results are robust to using the average percentage of forest cover in each village.

The VCF measure is a prediction of the percentage of a pixel that is covered by forest, generated from a machine learning model based on a combination of images from MODIS and samples from higher resolution satellites. The measure employs not only the visible bandwidth but also other bandwidths. For example, VCF uses thermal signatures because forested areas tend to be cooler than non-forested plantation areas, allowing VCF to (partially) distinguish between forest cover and plantations. To the extent that thermal signatures and other correlates can distinguish forests from non-forest plantations, VCF substantially improves upon the Normalized Differenced Vegetation Index (NDVI) that has been widely used in understanding the causes of deforestation (for example, Foster and Rosenzweig (2003)). For all analyses, we restrict the sample of villages to those that had non-zero forest cover in 2000, a year predating the construction of all roads considered in this research. This is also the earliest year that these forest cover datasets are available.⁶

Some earlier studies have used the Global Forest Cover (GFC) dataset, which describes baseline forest cover in the year 2000, and a binary indicator for the year of deforestation for each 30mx30m pixel. In the GFC data, a pixel is considered deforested if over 90% of 2000 forest was lost by a given year, or reforested if a pixel goes from zero forest in 2000 to positive forest cover by 2012 (Hansen et al., 2013). While GFC and VCF are both based on satellite imagery, GFC is less useful for the study of forest cover in India, because forest change in India is not well summarized by a binary deforestation indicator. The VCF measures suggest that forest cover rose 15% over the sample period, an estimate consistent with official and

⁵Results are robust to using the inverse hyperbolic sine transformation instead of log plus one.

⁶Fewer than 10% of villages have zero forest cover in 2000; 95% of these villages have less than 1% forest cover in 2014; the mean of forest cover for pixels with non-zero forest is 12.76% in 2000.

international sources. Because most of these gains are in areas that had some pre-existing forest, they are not recorded by GFC. GFC also does not describe partial forest loss, while VCF does. We can replicate GFC estimates by restricting the VCF data to forest losses, but they miss a significant share of forest change in the sample period. Because 92% of villages are larger than the VCF cell size, the resolution advantage of GFC would be minimal. In the cross-section data from the year 2000, VCF and GFC have a correlation coefficient of 0.92 with each other, as compared to respective correlation coefficients of 0.71 and 0.67 with an NDVI measure based on the choices of Foster and Rosenzweig (2003). Appendix Figure A2 presents heat maps of forest cover in 2000 according to these three datasets, which convey clearly the similarity between VCF and GFC, and the difference of both of these from NDVI.

We matched forest cover data to the 2011 Population Census village, town and subdistrict boundaries using geographic boundary data purchased from ML InfoMap. In remote parts of India, we received only settlement centroids rather than village boundaries. We generated Thiessen polygons for these villages; all results are robust to excluding this set of villages. Panel B of Figure 1 shows a heat map of baseline forest cover in India. While contiguous areas of very dense forest are geographically concentrated, areas with 20-40% of their land covered by forest are found throughout the country.

III.B Rural Roads

We scraped village-level administrative data describing the construction of rural roads from the program’s online management portal.⁷ For each road, the data provide the names of connected villages, the date when the contract for road construction was awarded, and the date of road completion. While data were reported at the sub-village (habitation) level, we aggregated the data to the village level to match our other data sources. We define a village as treated if any habitation in the village was provided with a new road. The data construction and scraping approach is described in detail in Asher and Novosad (2018). The dataset describes over 100,000 new roads built between 2001 and 2014; we limit our

⁷The data is publicly available at <http://omms.nic.in>.

sample to areas with non-zero forest cover and no paved road in the baseline year, leaving approximately 65,000 new roads in the analysis sample.⁸

III.C Highways

Construction dates and geocoordinates for the Golden Quadrilateral and North-South and East-West corridors were generously shared with us by Ghani, Goswami and Kerr (2016). We linked these to the village, town and subdistrict polygons described above by calculating straight line distances from polygon centroids to the nearest point on each highway.

III.D Population and Economic Censuses

We matched all villages and towns from the 1991, 2001 and 2011 population censuses using a combination of incomplete keys provided by the Registrar General and a set of fuzzy matching algorithms based on village and town names. The population censuses describe village and town public goods, village amenities (such as schools and medical centers) and household characteristics, including the primary source of cooking fuel. Fuel use is reported as the share of households in a location using firewood (68% of households at baseline), imported fuels (chiefly propane, 8%) or local nonwood fuels (crop residue and dung, 22%) as a primary source of energy. Fuel use is reported at the subdistrict level in 2001 and at the village level in 2011.

The Economic Censuses are complete enumerations of all nonfarm establishments undertaken in 1990, 1998, 2005 and 2013, including informal and non-manufacturing firms. We matched these on village names to the three population censuses using a fuzzy matching algorithm. The Economic Census reports total employment and industry for all firms. We create variables describing total employment in (i) firms engaged in logging and (ii) firms whose primary input is raw lumber, which include sawmilling and planing of wood, manufacture of wooden products such as furniture and wooden containers, manufacture of cork,

⁸Results are robust to including upgrades and/or villages with no forest cover at baseline. These would be expected to attenuate non-zero treatment effects, thus their exclusion if anything biases us against finding zero effects.

and manufacture of pulp and paper products. The industry categorization for the 2005 Economic Census places logging firms in the same industry category as firms engaged in the conservation of forest plantations, management of forest tree nurseries and other afforestation categories. We therefore exclude 2005 from analysis of employment in logging firms.

IV Impacts of Rural Feeder Roads on Forest Cover

This section describes the impact of new feeder roads on local deforestation. The main challenge to causal identification of the impacts of rural roads is endogeneity. Because roads are costly to build, their placement is typically correlated with other factors that could also be predictors of deforestation. For example, roads could be targeted to places that are expected to grow or to places that are lagging economically. Road placement may also depend on geographic (e.g. slope, terrain, soil quality) or political factors. Any of these scenarios would bias OLS estimates of the effect of new roads on deforestation.⁹ Causal identification of the impact of new roads therefore relies on some kind of variation in road placement or timing that is plausibly exogenous. To study the impact of rural roads, we rely on (i) an implementation rule that led to a discontinuity in the probability of a village getting a new road based on arbitrary population cutoffs; and (ii) variation in the specific year that a targeted village was treated. We focus our analysis on forest cover in the vicinity of connected villages. Because newly connected rural villages are mostly small and isolated, and because most of the new roads terminate in villages rather than providing new long distance corridors, these roads are unlikely to have had important general equilibrium effects on more distant areas.

⁹Appendix Table A1 shows estimates from cross-sectional OLS regressions of village-level log forest cover in 2001 on an indicator variable that takes the value one if a village has a paved road in 2001. While the bivariate relationship is strongly negative and highly statistically significant, the estimate gets progressively closer to zero as we add village-level controls and fixed effects, implying substantial selection on observables in the presence of roads. Selection on unobservables is plausibly also important, making the OLS estimates unreliable for causal inference.

IV.A Rural Roads: Regression Discontinuity Specification

We begin by exploiting the eligibility rule that prioritized villages for new roads based on arbitrary population thresholds. Given the imperfect compliance with these eligibility rules (described in Section II), we employ a fuzzy regression discontinuity (RD) design. We limit the RD analysis to states in which administrators adhered closely to population threshold rules.¹⁰

We use an optimal bandwidth local linear regression discontinuity specification (Imbens and Lemieux, 2008; Imbens and Kalyanaraman, 2012; Gelman and Imbens, 2014) to identify the change in forest cover caused by a new road at the treatment threshold. We use the following two stage least squares specification:

$$\begin{aligned} Treatment_{vds} = & \gamma_0 + \gamma_1 \cdot (pop_{vds} \geq T_s) + \gamma_2(pop_{vds} - T_s) + \gamma_3(pop_{vds} - T_s) \cdot (pop_{vds} \geq T_s) \\ & + \nu_d + \boldsymbol{\theta} \mathbf{X}_{vds} + \epsilon_{vds} \end{aligned} \quad (1)$$

$$\begin{aligned} Forest_{vds} = & \beta_0 + \beta_1 \cdot Treatment_{vds} + \beta_2(pop_{vds} - T_s) + \beta_3(pop_{vds} - T_s) \cdot (pop_{vds} \geq T_s) \\ & + \mu_d + \boldsymbol{\kappa} \mathbf{X}_{vds} + \eta_{vds} \end{aligned} \quad (2)$$

$Forest_{vds}$ is forest cover in village v , district d and state s , and $Treatment_{vds}$ is an indicator equal to one if a new road was built in village v . pop_{vds} is the population of village v and T_s is the treatment threshold used in state s .¹¹ μ_d and ν_d are district fixed effects; we find virtually identical results with fixed effects at higher or lower geographic scales. \mathbf{X}_{vds} is a control for

¹⁰We identified these states with the help of officials at NRRDA. They include Chhattisgarh, Gujarat, Madhya Pradesh, Maharashtra, Orissa and Rajasthan. The difference-in-differences analysis below uses all states that built any roads in the sample period.

¹¹The treatment threshold varies with state because some states used a threshold of 500 and others were using a threshold of 1000. States used the lower treatment threshold when they had few villages with population over 1000 that did not already have roads. Officials at the National Rural Roads Development Agency provided us with information on which states were using which cutoffs, which we then verified in the data. Madhya Pradesh used both the 500 and 1000 treatment thresholds for roads built in the same period; we include separate fixed effects for the set of villages in the neighborhood of each threshold. Because the optimal regression discontinuity bandwidth is close to 100, there is no overlapping between these two groups. Few villages around the lowest population threshold of 250 received roads so we do not use this threshold for analysis.

baseline forest cover; like the fixed effects, the control is unnecessary for identification but improves precision. This is a cross-sectional regression where β_1 identifies the effect of new roads on forest cover in a given year. Outcomes are measured in the final year in the sample data, which is 2013.¹²

Appendix Figure A3 shows regression discontinuity balance tests for a set of variables measured in the baseline period; Appendix Table A2 presents the regression estimates on these tests using Equation 1. None of the regression discontinuity estimates are significantly different from zero at baseline. Appendix Figure A4 shows that the density of the running variable is continuous around the treatment threshold (McCrary, 2008).

IV.B Rural Roads: Regression Discontinuity Results

Figure 2 shows a graphical representation of the regression discontinuity estimates of the impact of rural roads on forest cover. Panel A shows the first stage; the Y axis shows the share of sample villages that received new roads by 2013 under PMGSY as a function of their population relative to the treatment threshold. Villages above the threshold are about 16 percentage points more likely to receive new roads and the discontinuity is evident. Panel B shows the first stage estimate separately for each outcome year; each point in the figure represents the γ_1 coefficient from Equation 1, where the dependent variable takes the value one if a village received a new road by the year indicated on the X axis. We can see that roads built before 2007 were not prioritized according to the population threshold rule; the first stage of the RD becomes noticeable after 2008 and continues to rise until 2014.

Panel C of Figure 2 plots village-level log forest cover in 2013 against the population relative to the treatment threshold, in population bins. If roads significantly affected local forest cover, we would expect to see a discontinuity at the treatment threshold analogous to that in Panel A; no such treatment effect is evident. Panel D shows the reduced form treatment effect of above-threshold population on forest cover (β_1) in each year separately;

¹²We find similar results if we pool outcome years from 2010 through 2013 and cluster standard errors at the village level (not shown). Standard errors are slightly smaller with this alternative approach, at the cost of putting more weight on roads which have been built for shorter periods of time.

as in Panel B, each point is an estimate from a separate regression, where the dependent variable is the log of forest cover for the year on the X axis. If the new rural roads significantly affected forest cover, we would expect to see a change in the coefficient following 2008 when administrators began to adhere to the population implementation rule. Instead, the effect is very close to zero both before and after 2008, indicating that new rural roads had negligible effects on forest cover.

Table 2 shows analogous regression estimates, where the dependent variable is forest cover as measured in 2013. Column 1 shows the first stage estimate of a 16 percentage point increase in the probability of road treatment for villages just above the eligibility threshold. Columns 2 and 3 confirm there is no reduced form effect on either log or average forest cover. Columns 4 through 6 test for treatment effects in villages that might be expected to respond more to new roads. These are: villages with above-median baseline forest cover (Column 4); villages with above-median population shares of constitutionally described “backward” communities (Scheduled Tribes) who often derive livelihoods from forests (Column 5); and villages with below median assets, who might depend more on forests for fuelwood (Column 6). There is no evidence of impacts of roads in any of these groups.¹³ Columns 7 and 8 show IV estimates on log and average forest cover. The IV estimates respectively rule out a 0.14 gain and a 0.11 loss in log forest cover with 95% confidence, or approximately a one percentage point change in average forest cover. The average treated village in the sample received a new road in 2008, so these estimates reflect cumulative forest change five years after a village is connected. Results are robust to different controls or fixed effects and different bandwidth choices.¹⁴ Appendix Table A5 uses the RD specification to show further that there are no changes in household fuel use following completion of a new road.

¹³Appendix Table A3 shows further that roads do not significantly affect forest cover in villages defined by high or low town distance, market access, nor in villages in subdistricts with above median employment in the logging sector or in industries that are heavy consumers of wood.

¹⁴Results at many different bandwidths are shown in Appendix Table A4.

IV.C Rural Roads: Difference-in-Differences Specification

The regression discontinuity design estimates causal impacts of roads under minimal assumptions, but is limited to estimating a LATE in the neighborhood of the treatment threshold in states that closely followed implementation rules on population thresholds. We can make greater use of our data and obtain tighter treatment estimates using a difference-in-differences specification that exploits the differential timing of road treatment in each village. For this empirical test, we limit the sample of villages to those that received a road at *some* point during the road construction program, and use outcomes in later-treated villages as a control group for villages that were treated earlier. We specifically estimate the following equation:

$$Forest_{vdt} = \beta_1 \cdot Award_{vdt} + \beta_2 \cdot Complete_{vdt} + \alpha_v + \gamma_{dt} + \mathbf{X}_v \cdot \boldsymbol{\nu}_t + \eta_{vdt} \quad (3)$$

$Forest_{vdt}$ is a measure of forest cover in village v and district d in year t . $Award_{vdt}$ is an indicator that takes the value one for the years where a contract has been awarded for the construction of a road to village v but the road construction is not yet complete. $Complete_{vdt}$ is an indicator that takes the value one for all years following the completion of a new road to village v . We separate these two periods because the road construction process may have effects on forest cover (such as clearing of forested area to make room for the physical placement of roads) that are theoretically distinct from the economic effects of a village having a new road. Village fixed effects (α_v) control for all village-level time-invariant unobservables, while district-year fixed effects (γ_{dt}) control for any pattern of regional shocks.¹⁵ We also interact a vector of baseline village controls \mathbf{X}_v (baseline forest cover, village population and distance from the village to the nearest towns) with year fixed effects. These control for any differential time path of forest cover that is correlated with baseline village characteristics. These controls are particularly important because larger villages are more likely to be treated earlier due to program implementation rules. Standard

¹⁵Results are unchanged by replacing these with state-year or subdistrict-year fixed effects.

errors are clustered at the village level to account for serial correlation.

We can interpret β_1 and β_2 as the effects of road construction activities and the effects of new roads, respectively; both coefficients describe outcomes relative to the period before any construction began. We restrict our sample from the universe of villages in India to those that had no road in 2000 and had a road completed during the study period. We do this so as not to compare villages that received new roads with those that did not; the endogeneity problem in such a comparison is severe.¹⁶ Identification rests on the assumption that, among the set of villages that received roads in the sample period, there are no other systematic changes specific to villages in the years that roads were awarded and completed that are not caused by the roads themselves.

IV.D Rural Roads: Difference-in-Differences Results

The difference-in-difference estimates of the impact of rural roads on village-level forest cover are summarized by Figure 3. These graphs show the residual of log forest cover — after taking out fixed effects and controls described above — as a function of the number of years elapsed since a road was completed in a given village. Panel A shows all previously-unconnected villages that received new roads between 2001 and 2014. Panel B restricts the set of villages to those with above median forest cover in 2000. We show only four years before and after road construction because wider windows have more variable sample composition across estimates; this occurs because we observe different length of pre- and post-periods for different villages depending on their date of treatment.¹⁷ Two patterns are evident in the figure. First, there is a statistically significant reduction in forest cover approximately two years before road construction is complete. Second, forest cover marginally increases in the four years after road completion recovering some or all of the pre-treatment drop.

Given that these rural roads took one to two years to build, this pattern is consistent

¹⁶As we show above, a minority of roads were allocated strictly due to the village population thresholds. There are enough of these to estimate a regression discontinuity test on local compliers, but not enough to assume that all treated villages are selected as good as randomly.

¹⁷Appendix Figure A5 shows a wider time window around treatment; the pattern is the same.

with a small degree of forest loss (approximately 0.5%) during the road construction period, with partial or complete recovery afterward. We test this directly in Table 3, which shows estimates from Equation 3. Our main estimate in Column 1 shows that villages lose 0.5% of their forest cover during the period between the awarding of a road construction contract and the completion of a road. However, that forest loss is fully restored in the period after the road has been completed; the estimate of 0.002 log points on the completion indicator can be interpreted as the difference in forest cover between the post-road and the pre-award periods. Relative to the pre-award period, we can rule out gains larger than 0.6% and declines larger than 0.2% in forest cover. In Column 2, we show that failing to account for the award period would lead to the estimation of a marginal forest cover gain of 0.5% because it would incorrectly attribute the construction period loss to the pretrend. This result highlights the importance of accounting for the construction period when studying the environmental impacts of new infrastructure. Columns 3 and 4 present estimates where forest cover is measured as the average share of each pixel that is covered by forest; results are similar. These estimates are based on different lengths of post-construction periods in different villages, but on average they show effects for four years after treatment.¹⁸

Table 4 shows these estimates along the same dimensions of heterogeneity described above. Effects are broadly similar whether we cut the sample on baseline forest cover, population share of Scheduled Tribes, or asset poverty. There is thus no evidence that our zero results are hiding differential positive and negative effects in different places.¹⁹ It is also

¹⁸Appendix Table A6 shows that these estimates are robust to a range of specifications including the use of village time trends, subdistrict-year fixed effects (instead of district-year fixed effects) and using a limited sample of roads for which we have at least 4 (or 5) years of both pre-treatment and post-treatment data. Appendix Table A7 shows additional specifications. Column 1 adds villages that did not receive roads in the sample period, the specification used in Kaczan (2017). Like Kaczan (2017), we find a positive treatment coefficient; however, Column 2 shows that this is not robust to the inclusion of village-specific time trends, indicating that never-treated villages are on different forest cover trends from treated villages. Columns 3 and 4 show that our main estimate is robust to village-specific time trends. Column 5 and 6 define the treated area as a circle around the village with a radius of 5km and 50km, respectively; as in the main specification, we find no treatment effects at these radii.

¹⁹Equally, we find no effects of roads on forest cover when splitting the sample on distance to the nearest town or on market access, nor in subdistricts with above median employment in logging or in industries with high consumption of wood (Appendix Table A8).

unlikely that outmigration of individuals following road construction is significantly biasing our findings; these rural roads are not associated with significant population change at the village level (Asher and Novosad, 2018). Rural-to-urban migration has also been much slower in India over the sample period than in other countries at comparable levels of income.

The panel estimates confirm the finding in the regression discontinuity analysis, using a different set of villages with a different local average treatment effect; the evidence is clear that new rural roads have had a negligible effect on local forest cover.

V Impacts of Major Highways on Forest Cover

In this section, we aim to identify the causal impact of highways on local forest cover. The identification challenge is that highways are typically built to connect cities with current or anticipated economic growth; if economic growth is correlated with forest cover changes for any reason other than the direct effect of highways, then we cannot interpret the correlation between highways and forests as a causal effect.

We therefore focus on a set of places that happen to be in between the targeted endpoints of India’s new highways, as in Ghani, Goswami and Kerr (2016). Both the Golden Quadrilateral (GQ) and the North-South and East-West (NS-EW) corridors were upgraded with the objective of improving connections between India’s major cities and regions; the connection of secondary cities and intermediate places on the route was a secondary priority. Because these intermediate regions were targeted incidentally rather than directly, the placement of the highways is less likely to be driven by existing or anticipated economic growth.

We can further generate a plausible counterfactual that describes how forest cover would have changed in the absence of the highway upgrades. Like the GQ, the NS-EW route was an important transportation corridor in 2000 and was to be upgraded before 2005 as part of NHDP, but the project did not begin in earnest until several years after the GQ was completed. Our main estimates examine forest changes along the GQ corridor during and after the construction years, as compared to regions further from the GQ. We then test for effects along the NS-EW route using a similar specification, showing there are no effects

along the second corridor until after 2008, as would be expected given the construction delay.

As a starting point, Figure 4 plots kernel-smoothed local regression estimates of mean forest cover and forest cover change as a function of distance from each highway. Initial forest cover (Panel A) is broadly similar across the two highways. Panel B shows forest cover change from 2000-2008, also by distance to each highway. Relative to the NS-EW (dashed line), forest cover within 100 km of the GQ (solid line) falls substantially between 2000 and 2008. At further distances the effects are similar across the two highways, though there may be smaller relative gains for the GQ. We present this as suggestive evidence of relative forest loss along the GQ corridor during and after its construction. The rest of this section generates formal tests for change, controlling for fixed effects and other factors that may have simultaneously influenced forest change.

V.A Highways: Empirical Specification

The simplest form of the difference-in-differences specification is described by the following equation:

$$Forest_{ist} = \beta_0 + \beta_1 CLOSE_{is} + \beta_2 POST_t + \beta_3 CLOSE_{is} * POST_t + \epsilon_{ist} \quad (4)$$

In this specification, i indexes a subdistrict in state s and time t , $CLOSE_{is}$ is an indicator for subdistricts close to the highway, and $POST$ indicates years following the completion of the highway. $Forest_{ist}$ is a measure of forest cover in subdistrict i and state s at time t , usually log total forest cover. β_3 describes the differential change in forest between locations that are near and far from the highway network after the highway is built, controlling for the same geographic difference before the highway was built. If new highways cause deforestation, we expect β_3 to be less than zero. We conduct our analysis at the subdistrict level, because subdistricts are contiguous regions that cover the whole of India for which we can calculate a range of demographic and socioeconomic controls. We weight results by subdistrict area.²⁰

²⁰Results from a town- and village-level analysis with subdistrict clusters deliver nearly identical results. We could in principle conduct analysis at the grid cell level, but this would require imputation for control

There are approximately 4000 subdistricts in India.

We extend this simple specification in three ways. First, because we do not have strong priors on which distances are near and which are far, we use a flexible set of distance indicators to nonparametrically identify highway effects at a range of distances. Estimates can still be interpreted as the difference from a given band to the omitted (most remote) distance band. This ensures that our result is not dependent upon a particular definition of closeness. Second, because the construction of India’s national highways were multiyear projects, we separate the $POST_t$ indicator into multiple periods to capture construction and post-construction effects. Third, we add a wide set of fixed effects and controls to improve precision and reduce bias from omitted variables. The most flexible estimating equation is:

$$Forest_{ist} = \sum_{d=1}^D \sum_{t=2001}^{2014} \beta_{d,t}(DIST_i \in (d^-, d^+), YEAR = t) + \gamma_{st} + \mathbf{X}_i \cdot \boldsymbol{\nu}_t + \psi_d + \eta_{ist} \quad (5)$$

The distance to the highway is divided into D bands, the boundaries of which are indexed by d . We include a distance band fixed effect ψ_d , state-year fixed effect γ_{st} and a vector of subdistrict controls (\mathbf{X}_i) interacted with year fixed effects ($\boldsymbol{\nu}_t$). The latter control for any differential time path of forest cover that is correlated with baseline subdistrict characteristics. Controls are the same as in Equation 3. We include locations up to a distance $D + E$ from the Golden Quadrilateral; the outer boundary (D, E) is the omitted distance category against which the other estimates can be compared. Unless otherwise specified, we define (D, E) as the 200-300km distance band.²¹ $\beta_{d,t}$ identifies the change in forest cover from 2000 to year t , at distance range d from the highway, relative to the omitted distance range (D, E) . The $\beta_{d,t}$ coefficients can thus be directly interpreted as the effect of highway construction on forest cover after t years. If new highways cause proximate forest cover loss, we expect $\beta_{d,t}$ to take on negative values for low values of d in the periods t after highway construction has begun. For graphs, we include a set of indicator variables $\beta_{d,2000}$ which

variables not available at the grid cell level.

²¹Alternate choices of the range of the omitted group, including using the remainder of the country does not appreciably affect our estimates.

describe baseline forest cover as a function of distance from the highway.²² Standard errors are clustered at the subdistrict level to account for serial correlation. Because the regression above may have hundreds of coefficients, we pool years or distances in different specifications to improve interpretability.

We exclude areas within 200km of the nodal towns on the highway routes, as we wish to identify effects on intermediate regions rather than at the highway end points, as in Goswami Ghani and Kerr (Ghani, Goswami and Kerr, 2016). Estimates of NS-EW treatment effects omit areas that are within 200km of the GQ as they are plausibly being treated by the other highway network. We do not omit NS-EW regions from the GQ regressions because NS-EW construction has barely begun during the periods of interest for the GQ analysis; however, regression results are not changed by omitting places within 200km of NS-EW.

V.B Highways: Estimates on Forest Cover

Panel A of Figure 5 plots coefficient estimates from a single estimation of Equation 5, with distances from the GQ highway divided into 10km bands, and years divided into a single pre-construction year (2000), the construction period (2001-2004), and two post construction periods (2005-2008 and 2009-2012). All estimates describe the difference between a given 10km distance band from the GQ and the omitted category of 250-300km.²³ The solid black line describes baseline forest cover as a function of distance from the GQ corridor. The remaining lines show that forest cover within 100 km of the GQ declines rapidly during the GQ construction period and then continues to fall in the years following construction.

²²We do not include subdistrict fixed effects because we want to generate coefficients on the distance band indicators for the omitted year 2000 — these coefficients describe the baseline differences in forest cover between places that were near and far from the highway. However, inclusion of subdistrict fixed effects does not meaningfully change the results. We use state-year fixed effects rather than district-year fixed effects because we wish to test for meaningful effects of distance from highways that may extend beyond the radius of districts. District-year fixed effects would absorb true effects of the GQ that span distances larger than districts. The analysis of Goswami Ghani and Kerr (Ghani, Goswami and Kerr, 2016) is entirely district level, giving us reason to expect meaningful cross-district effects. As expected, the inclusion of district-year fixed effects attenuates our results slightly but does not change the direction of effects nor eliminate statistical significance.

²³We include coefficients for the 200-250 km bands in order to plot treatment effects at these ranges. Effects in closer bands are very similar if we restrict the distance indicators to 200km and use 200-300km as the omitted group, because there are few differences across years in the 200-250km range.

Effects are slightly smaller in the 100-150km bandwidth, and statistically indistinguishable from zero at a distances greater than 150km from the highway.

To alleviate the concern that these forest losses are explained by existing trends in forest loss along existing highway corridors, we run the same estimation for subdistricts along the NS-EW corridor and show results in Panel B. As predicted, there are no differential changes in forest cover close to the NS-EW route before 2008. Net forest loss along NS-EW begins in the 2009-2012 period, and the distance effects then look similar to those of the GQ.²⁴ This distance pattern of forest cover loss is similar to that found in the Amazon (Pfaff et al., 2007). Effects along the NS-EW corridor may be slightly smaller than along the GQ both because construction took place slowly and was still incomplete at the end of the sample period in 2014; the network structure of highways mean that the value of any particular segment depends on the completion status of other segments. Concentration of industry along the GQ corridor may have also reduced the importance of NS-EW as an intercity transportation corridor by the time the NS-EW upgrades took place.

Table 5 presents regression estimates from Equation 5, with distances in 50km bands for legibility. Each estimate describes the difference in forest cover between a given distance band and the omitted category of 200-300km. Columns (1) and (2) describe the impact of the GQ on forest cover. The top four rows of the table show estimates of construction period impacts on forest cover at various distance bands. Places within 50km of the new highway network lose 27 log points of forest cover (Column 1) or 1.3 percentage points of forest cover (Column 2, on a base of 7.5%), and the effects shrink at greater distances. The next four rows show similar effects (still relative to year 2000) in the post-construction period of 2005-2008. The final four rows show estimates of baseline differences between the GQ and the regions further away; the level differences are small relative to the treatment effects.

Columns 3 and 4 of Table 5 show comparable estimates along the NS-EW corridor for

²⁴Standard errors are omitted in the figure for visual clarity. For the GQ, differences between the 2000 estimates and the 2001-2004 estimates are statistically significant at the 1% level for all estimates up to 150km. For the NS-EW, differences between the 2000 estimates and the 2009-2012 treatment estimates are statistically distinguishable at the 1% level until the 180km estimate.

the same time periods. There are no detectable changes in forest cover close to the NS-EW in the time period when significant forest loss took place near the GQ. Along with Figure 5, this should alleviate any concern that the GQ treatment effects are driven by generalized deforestation along existing highway corridors from 2001-2008. These results are robust to instrumenting for highway location using straight line instruments connecting the nodal cities of the highway network, as employed by Ghani, Goswami and Kerr (2016); we present analogous reduced form estimates to the above in Appendix Table A9. These estimates alleviate the concern that the particular routing of the GQ (but not the NS-EW) was specifically targeted to places that may have already been losing forest cover.

An alternate estimation approach would be to exploit the timing of construction of each segment of the highway and to study forest cover in the years before and after the nearest segment to a given location was built. This would be analogous to the difference-in-differences estimation used to study the impact of rural roads in Section IV. The results, which are consistent with the findings above, are presented in Appendix Table A10.²⁵ One useful feature of this analysis is that it is directly analogous to our difference-in-differences estimates of the impacts of rural roads (Equation 3 and Table 3), making clear that the differential effects of highways and rural roads are not an artifact of the different empirical strategies used to evaluate them.

A final concern with these estimates of forest change is that they could be describing displacement of forest loss from the hinterlands to the highway corridors, or even net afforestation in the hinterlands. This concern arises frequently in studies of transportation projects with national scale, and is typically only resolved by assumption through a structural modeling approach, which is beyond the scope of this paper. This said, large displacement effects are made less plausible by the low quality of the broader road network, and by

²⁵We did not use this as a primary specification because there are substantial network effects in highway construction. A region without an upgraded highway segment may experience greater transportation access if the regions around it experience upgrades. Equally, a nearby upgrade may not be of much value if it is an isolated upgrade to a not-yet-upgraded part of the corridor. We therefore would expect estimates from a specification exploiting timing to underestimate the full impacts of the highway construction, which is what we find in Appendix Table A10.

the high transportation costs during the sample period, weakening market connections with the hinterlands of these highways. While we cannot entirely rule out that there may be some displacement effects, Panel B of Figure 4 suggests that effects are driven by the highway corridor regions rather than the hinterlands; forest cover diverges between the places close to the GQ and NS-EW, but not in those far away.

V.B.1 Mechanisms for Highway Effects

We consider four possible mechanisms for the forest cover loss caused by India’s major highway networks: (i) increased demand for timber products by firms due to local growth; (ii) increased demand for firewood due to shifts in household fuel consumption; (iii) expansion of agriculture into previously forested lands; and (iv) clearing of trees for settlements and industry. This section presents suggestive evidence that the deforestation along India’s major highways is predominantly caused by increased logging driven by local timber demand.

To identify potential mechanisms, we use the regression specification used to identify effects of highways on forest cover (Equation 5), with data from the economic and population censuses which were undertaken in various periods between 1990 and 2013. Because the results above suggest forest cover changes within 100km of the highway network, we use 100km distance bands and define 200-300km as the omitted distance band.²⁶ The years in the sample are determined by census availability.

We first look at changes in employment in the list of industries that are directly downstream from timber harvesting (described in Section III). Panel A of Figure 6 shows results from a regression of log employment in major wood-consuming sectors on the usual set of year-distance-band fixed effects. We graph the point estimates on the 0-100km coefficient in each year that the Economic Census is available (*i.e.* the coefficients $\beta_{0-100km,1990}$, $\beta_{0-100km,1998}$, $\beta_{0-100km,2005}$, and $\beta_{0-100km,2013}$ from Equation 5). These estimates can be interpreted as the difference in residual log employment between the 0-100km distance band and the 200-300km distance band, after controlling for state-year fixed effects and the con-

²⁶We find virtually identical results when we use 50km distance bins.

trols described above. In 1990 and 1998 (before the GQ was begun), there is no significant difference between areas close to the highway corridor and areas that are far from it, nor is there a significant trend. By 2005, we see a 5% increase in employment in wood-consuming firms in the GQ corridor relative to the hinterland, which continues to rise through 2013. Panel B shows similar results for employment in logging firms; we omit 2005 because logging firms were not distinguished from firms engaged in afforestation in the 2005 Economic Census. These two graphs suggest that demand for wood from downstream firms is a plausible explanation for local deforestation after construction of the GQ.²⁷ Logging firms were more common along the transport corridor even before the GQ was built, but there is no suggestion of a pretrend that could explain what we see after highway construction. Note that employment in other sectors of the economy, which also consume wood, exhibit similar treatment effects close to the new highways (Ghani, Goswami and Kerr, 2016).

Panels C through E of Figure 6 show the effects of the GQ upgrades on household fuel consumption, for which data are available in 2001 and 2011. We observe marginal increases in firewood and imported wood use, and comparable reductions in use of local non-wood fuels. These effects are not statistically significantly different from zero, nor are they large enough to explain a 20% reduction in forest cover in the neighborhood of the GQ corridor. Panel F shows that land use shifts slightly *away* from agricultural uses following the construction of the GQ, breaking a previous upward trend, making agricultural extensification an unlikely explanation for the treatment effects.

It is difficult to directly test the last hypothesis that net forest loss has come from the expansion of land dedicated to settlement and industry, because data on land dedicated to settlement and industry only becomes available in the 2011 Population Census. However, it is implausible that settlement and industrial expansion could explain a 20% reduction in forest cover in a distance band as wide as 100km around a 6000km long highway corridor.

²⁷Appendix Tables A11 and A12 show complete regression results for all mechanism tests. As would be expected, employment in logging firms is more geographically diffuse, as those firms reach further into the GQ's hinterland.

In 2011, only 6.7% of rural land was used for settlement and industry.

The large relative forest losses along the GQ corridor are unlikely to be direct effects of the road construction process. Although we found evidence of these effects in the construction of rural roads in Section IV, these were very local, temporary, and an order of magnitude lower in size. It is also implausible that direct construction effects would extend more than a few kilometers from the roadway.

In conclusion, we find evidence that expansion of industry demand for timber can explain forest loss in the GQ corridor, and we can rule out agricultural expansion, changes in household fuel consumption and settlement expansion as mechanisms.

VI Conclusion

The development, maintenance and expansion of transportation infrastructure is an important driver and correlate of economic development around the world. In this paper, we provide causal estimates of the ecological impact of two transportation investments with global significance: India’s massive expansion of rural roads and its upgrading of national highways. Using identification strategies established in the literature, we find that: (i) the new rural roads had negligible effects on forest cover and; (ii) the highway expansions had a large negative effect on forest cover, which may have been driven by the expansion of wood-using industries. Methodologically, we demonstrate the critical importance of accounting for endogeneity and separately estimating the effects of the construction period from the post-completion period.

Because of the different implementation and spatial structure of rural roads and highways, it was necessary to use different empirical tests to evaluate the environmental impacts of each. However, our differential results are unlikely to be the result of having used these different strategies. Our analysis of rural roads ruled out even small effects on forest cover extending six years after the onset of construction and four years after road completion—a time horizon in which the highway expansions already had substantial impacts on local forest cover. We also find no effects of rural roads on forests in places with high local demand for wood from

either individuals or firms—which are the most likely drivers of forest change along highway corridors. Finally, while most of our analysis of rural roads focuses on local effects, we also found no effects at the wider distance horizons at which forest loss responds to highway expansion.

Globally, road expansion is expected to dramatically increase through the course of the 21st century. Some additional 25 million kilometers of road infrastructure is projected to be built by 2050, a 60% increase over 2010 levels. Nine out of ten of these roads will be built in developing countries (Laurance et al., 2014). At the same time, tropical forests in developing countries are increasingly under threat. These forests not only provide global carbon benefits but also provide important local ecosystems which support biodiversity as well as the generally poor populations that rely on them (Barrett, Garg and McBride, 2016). Against the background of this tension between economic development and environmental conservation, understanding the relationship between roads and forests is fundamental to a successful strategy for sustainable development.

Crucially, we show that the impact of road construction depends on what those roads connect. The fiscal costs of the two large scale transportation investments that we study were similar, but they had vastly different environmental consequences. Expansion of existing highway corridors caused changes in the spatial distribution of industry, which had dramatic effects on forest use in India. In contrast, building roads to connect smallholder farmers to new markets had virtually no impact on local forests, even for those farmers most likely to draw some part of their livelihoods from those forests.

References

- Adukia, Anjali, Sam Asher, and Paul Novosad.** 2019. “Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction.” American Economic Journal: Applied Economics (forthcoming).
- Alix-Garcia, Jennifer, Craig McIntosh, Katharine R.E. Sims, and Jarrod R. Welch.** 2013. “The Ecological Footprint of Poverty Alleviation: Evidence from Mexico’s Oportunidades Program.” Review of Economics and Statistics, 95(2): 417–435.
- Andreoni, James, and Arik Levinson.** 2001. “The Simple Analytics of the Environmental Kuznets Curve.” Journal of Public Economics, 80(2): 269–286.
- Angelsen, Arild, and David Kaimowitz.** 1999. “Rethinking the Causes of Deforestation: Lessons from Economic Models.” The World Bank Research Observer, 14(1): 73–98.
- Antweiler, Werner, Brian R. Copeland, and Scott M. Taylor.** 2001. “Is Free Trade Good for the Environment?” American Economic Review, 91(4).
- Arrow, Kenneth, Bert Bolin, Robert Costanza, Partha Dasgupta, Carl Folke, Crawford S. Holling, Bengt-Owe Jansson, Simon Levin, Karl-Göran Mäler, Charles Perrings, et al.** 1995. “Economic Growth, Carrying Capacity, and the Environment.” Ecological Economics, 15(2): 91–95.
- Asher, Sam, and Paul Novosad.** 2018. “Rural Roads and Local Economic Development.” Mimeo.
- Assunção, Juliano, Molly Lipscomb, Ahmed Mushfiq Mobarak, and Dimitri Szerman.** 2017. “Agricultural Productivity and Deforestation in Brazil.” Mimeo.
- Barber, Christopher P., Mark A. Cochrane, Carlos M. Souza, and William F. Laurance.** 2014. “Roads, Deforestation, and the Mitigating Effect of Protected Areas in the Amazon.” Biological Conservation, 177: 203–209.
- Barrett, Christopher B., Teevrat Garg, and Linden McBride.** 2016. “Well-being dynamics and poverty traps.” Annual Review of Resource Economics, 8: 303–327.
- Bauch, Simone C., Anna M. Birkenbach, Subhrendu K. Pattanayak, and Erin O. Sills.** 2015. “Public Health Impacts of Ecosystem Change in the Brazilian Amazon.” Proceedings of the National Academy of Sciences, 112(24): 7414–7419.
- Burgess, Robin, Matthew Hansen, Benjamin A. Olken, Peter Potapov, and Stefanie Sieber.** 2012. “The Political Economy of Deforestation in the Tropics.” The Quarterly Journal of Economics, 127(4): 1707–1754.
- Chomitz, Kenneth M., and David A. Gray.** 1996. “Roads, Land Use, and Deforestation: a Spatial Model Applied to Belize.” The World Bank Economic Review, 10(3): 487–512.
- Copeland, Brian R., and M. Scott Taylor.** 1994. “North-South Trade and the Environment.” The Quarterly Journal of Economics, 109(3).

- Copeland, Brian R., and Scott M. Taylor.** 2004. "Trade, Growth, and the Environment." Journal of Economic Literature, 42(1).
- Cropper, Maureen, Jyotsna Puri, Charles Griffiths, Edward B. Barbier, and Joanne C. Burgess.** 2001. "Predicting the Location of Deforestation: The Role of Roads and Protected Areas in North Thailand." Land Economics, 77(2): 172–186.
- Dasgupta, Partha.** 2007. "The Idea of Sustainable Development." Sustainability Science, 2(1): 5–11.
- Dasgupta, Susmita, and David Wheeler.** 2016. "Minimizing Ecological Damage from Road Improvement in Tropical Forests." World Bank Policy Research Working Paper No. 7826.
- Dasgupta, Susmita, Benoit Laplante, Hua Wang, and David Wheeler.** 2002. "Confronting the Environmental Kuznets Curve." Journal of Economic Perspectives, 16(1): 147–168.
- Datta, Saugato.** 2012. "The Impact of Improved Highways on Indian Firms." Journal of Development Economics, 99(1): 46–57.
- Den Butter, F.A.G., and Harmen Verbruggen.** 1994. "Measuring the Trade-Off Between Economic Growth and a Clean Environment." Environmental and Resource Economics, 4(2): 187–208.
- Deng, Xiangzheng, Jikun Huang, Emi Uchida, Scott Rozelle, and John Gibson.** 2011. "Pressure Cookers or Pressure Valves: Do Roads Lead to Deforestation in China?" Journal of Environmental Economics and Management, 61(1): 79–94.
- Donaldson, Dave, and Richard Hornbeck.** 2016. "Railroads and American Economic Growth: A "Market Access" Approach." Quarterly Journal of Economics, 131(2): 799–858.
- Ferretti-Gallon, Kalifi, and Jonah Busch.** 2014. "What Drives Deforestation and What Stops it? A Meta-Analysis of Spatially Explicit Econometric Studies." Mimeo.
- Foster, Andrew D., and Mark R. Rosenzweig.** 2003. "Economic Growth and the Rise of Forests." The Quarterly Journal of Economics, 118(2): 601–637.
- Frankel, Jeffrey A., and Andrew K. Rose.** 2005. "Is Trade Good or Bad for the Environment? Sorting Out the Causality." Review of Economics and Statistics, 87(1).
- Garg, Teevrat.** 2017. "Ecosystems and Human Health: The Local Benefits of Forest Cover in Indonesia." Mimeo.
- Geist, Helmut J., and Eric F. Lambin.** 2002. "Proximate Causes and Underlying Driving Forces of Tropical Deforestation: Tropical Forests are Disappearing as the Result of Many Pressures, Both Local and Regional, Acting in Various Combinations in Different Geographical Locations." BioScience, 52(2): 143–150.

- Gelman, Andrew, and Guido Imbens.** 2014. “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs.” NBER Working Paper No. 20405.
- Ghani, Ejaz, Arti Grover Goswami, and William R. Kerr.** 2016. “Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing.” Economic Journal, 126(591): 317–357.
- Grossman, Gene M., and Alan B. Krueger.** 1995. “Economic Growth and the Environment.” The Quarterly Journal of Economics, 110(2): 353–377.
- Hansen, M. C., P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Komareddy, A. Egorov, L. Chini, C. O. Justice, and J. R. G. Townshend.** 2013. “High-Resolution Global Maps of 21st-Century Forest Cover Change.” Science, 342(6160): 850–853.
- Imbens, Guido, and Karthik Kalyanaraman.** 2012. “Optimal Bandwidth Choice for the Regression Discontinuity Estimator.” Review of Economic Studies, 79(3): 933–959.
- Imbens, Guido W., and Thomas Lemieux.** 2008. “Regression Discontinuity Designs: A Guide to Practice.” Journal of Econometrics, 142(2): 615–635.
- IPCC.** 2007. The Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Geneva, Switzerland.
- IPCC.** 2019. Global Warming of 1.5 Degree C: Special Report on the Impacts of Global Warming. Geneva, Switzerland.
- Jayachandran, Seema, Joost de Laat, Eric F. Lambin, Charlotte Y. Stanton, Robin Audy, and Nancy E. Thomas.** 2017. “Cash for Carbon: A Randomized Trial of Payments for Ecosystem Services to Reduce Deforestation.” Science, 357(6348): 267–273.
- Kaczan, David.** 2017. “Can roads contribute to forest transitions?” Mimeo.
- Khanna, Gaurav.** 2016. “Road Oft Taken: The Route to Spatial Development.” Mimeo.
- Koop, Gary, and Lise Tole.** 1999. “Is There an Environmental Kuznets Curve for Deforestation?” Journal of Development Economics, 58(1): 231–244.
- Laurance, William F., Gopalasamy Reuben Clements, Sean Sloan, Christine S. O’connell, Nathan D. Mueller, Miriam Goosem, Oscar Venter, David P. Edwards, Ben Phalan, Andrew Balmford, et al.** 2014. “A Global Strategy for Road Building.” Nature, 513(7517): 229.
- Li, Man, Alessandro De Pinto, John M. Ulimwengu, Liangzhi You, and Richard D. Robertson.** 2015. “Impacts of Road Expansion on Deforestation and Biological Carbon Loss in the Democratic Republic of Congo.” Environmental and Resource Economics, 60(3): 433–469.

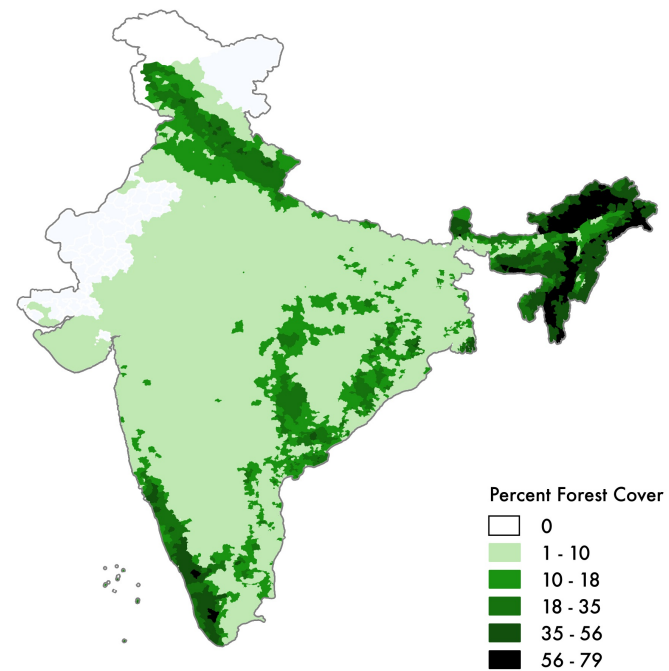
- McCrary, Justin.** 2008. "Manipulation of the running variable in the regression discontinuity design: A density test." Journal of Econometrics, 142(2): 698–714.
- National Rural Roads Development Agency.** 2005. "Pradhan Mantri Gram Sadak Yojana - Operations Manual." Ministry of Rural Development, Government of India.
- Pfaff, Alexander, Juan Robalino, Robert Walker, Steven Aldrich, Marcellus Caldas, Eustaquio Reis, Stephen Perz, Claudio Bohrer, Eugenio Arima, William Laurance, et al.** 2007. "Road Investments, Spatial Spillovers, and Deforestation in the Brazilian Amazon." Journal of Regional Science, 47(1): 109–123.
- Pfaff, Alexander S.P.** 1999. "What Drives Deforestation in the Brazilian Amazon?: Evidence From Satellite and Socioeconomic Data." Journal of Environmental Economics and Management, 37(1): 26–43.
- Reed, Tristan, and Alexandr Trubetskoy.** 2018. "The Belt and Road Initiative and the Value of Urban Land." Mimeo.
- Roberts, Peter, KC Shyam, and Cordula Rastogi.** 2006. "Rural access index: a key development indicator." The World Bank.
- Stern, David I.** 2004. "The Rise and Fall of the Environmental Kuznets Curve." World Development, 32(8): 1419–1439.
- Stern, David I., Michael S. Common, and Edward B. Barbier.** 1996. "Economic Growth and Environmental Degradation: the Environmental Kuznets Curve and Sustainable Development." World Development, 24(7): 1151–1160.
- Townshend, J., M. Hansen, M. Carroll, C. DiMiceli, R. Sohlberg, and C. Huang.** 2011. "User Guide for the MODIS Vegetation Continuous Fields product Collection 5 version 1."
- Weinhold, Diana, and Eustaquio Reis.** 2008. "Transportation Costs and the Spatial Distribution of Land Use in the Brazilian Amazon." Global Environmental Change, 18(1): 54–68.

Figure 1

Panel A: Highway Construction
(2001-2014)



Panel B: Forest Cover
(2001)




Panel A shows a map of the Golden Quadrilateral and North-South/East-West corridor highways. Panel B shows a heat map of forest cover in 2001. Areas are shaded according to average share of each pixel that is covered by forest.


Figure 2

Regression Discontinuity Estimates of Impact of
Rural Roads on Forest Cover

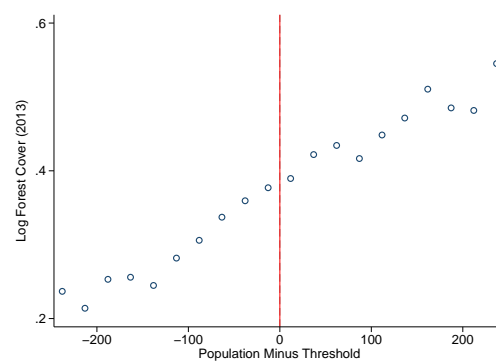
Panel A: First Stage (2013)




Panel B: First Stage (by year)



Panel C: Reduced Form (2013)




Panel D: Reduced Form (by year)




The figure shows regression discontinuity estimates of the impact of new rural roads on local deforestation. Panel A shows the first stage probability of a village receiving a new road before 2013 as a function of its population relative to the population threshold. Each point shows the mean of the Y variable in a given population bin. Panel B shows the first stage RD estimate of a village receiving a new road by the year indicated on the X axis. Each point is an estimate from an RD first stage regression. Panel C is analogous to Panel B; the dependent variable is the log of forest cover in 2013. The points show the mean of this variable in each population bin; population is shown relative to the population treatment threshold. Panel D shows reduced form RD estimates of the impact of being above the population threshold on forest cover in each year on the X axis. All estimates in Panels B and D use the same specification as Table 2, and include district-population threshold fixed effects and a control for baseline forest cover.

Figure 3
Difference-in-Differences Estimates of Impact of
Rural Roads on Forest Cover

Panel A: Full Sample




Panel B: High Baseline Forest



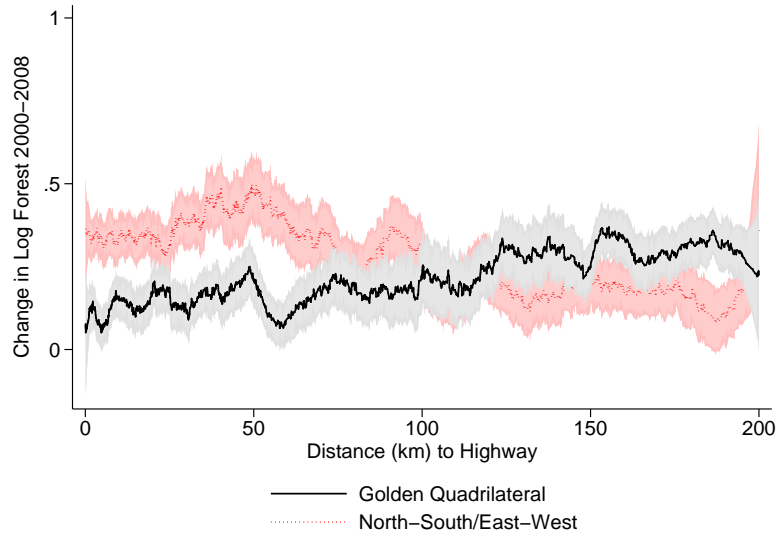
The figure shows year-by-year estimates of log forest cover in villages that received new roads between 2001 and 2013. Villages are grouped on the X axis according to the year relative to road completion. Each point thus shows the average value of log forest cover in villages in a given year relative to the treatment year, controlling for village fixed effects, district*year fixed effects, baseline population * year and baseline log forest cover * year interactions. Standard errors are clustered at the village level. The year before road completion is omitted ($t=-1$); forest cover is thus shown relative to this period.

Figure 4
Forest Cover and Forest Cover Change
Along Highway Corridors (2000-2008)

Panel A: Pre-Expansion Forest Cover (2000)

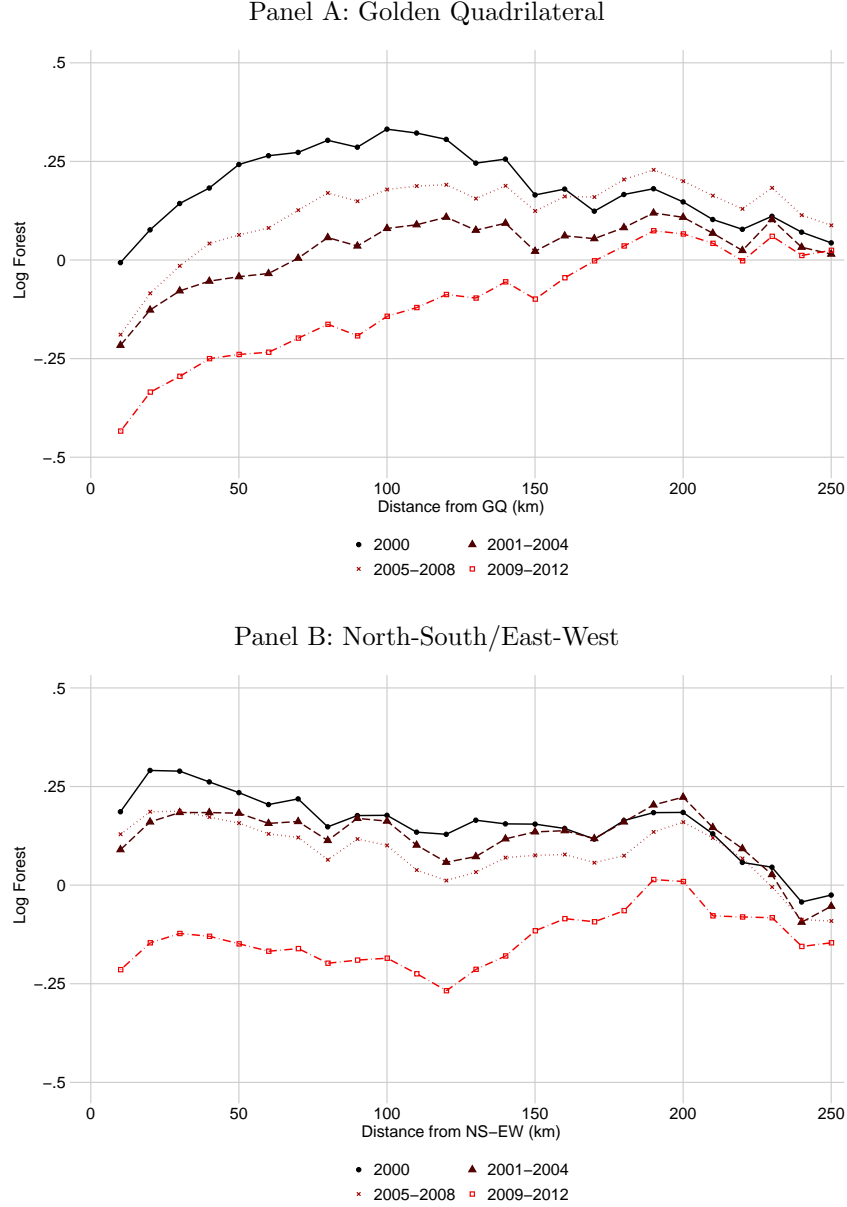


Panel B: Forest Cover Change (2000-2008)



Panel A shows a kernel-smoothed regression of log subdistrict forest cover in 2000 on distance to the corridors where the Golden Quadrilateral and North-South/East-West highways will be expanded. Panel B plots kernel-smoothed regression estimates of change in log subdistrict forest cover from 2000 to 2008 against distance to each highway network. By 2008, there was very little construction on the NS-EW corridor, so we treat it here as a control group. The plots display means that are unadjusted for any fixed effects or controls. 95% confidence intervals are displayed in the shaded areas.


Figure 5
Difference-in-Differences Estimates of Impact of
Highways on Forest Cover, by Distance Bands




The figure shows point estimates from Equation 5, with distance from the Golden Quadrilateral highway network (Panel A) and distance from the North-South/East-West highway network (Panel B) divided into 10km bands. Each point on the graph shows, for a given set of years (shown in the legend), the average value of log forest cover at a given distance band from the given highway network, relative to the omitted distance band of 290 to 300 km from the highway. All estimates control for state*year fixed effects, baseline population * year and baseline log forest cover * year interactions.

Figure 6
Mechanism Tests for Impact of Highways on Deforestation


Panel A: Employment in Wood-Using Firms




Panel B: Employment in Logging Firms




Panel C: Share of Energy from Firewood




Panel D: Share of Energy from Imported Fuels



Panel E: Share of Energy from Local Non-Wood Sources



Panel F: Agricultural Share of Village Land



The figure shows point estimates from Equation 5, with distances from the Golden Quadrilateral highway network specified in 100km bands. Each figure shows the point estimate on the 0-100km distance indicator, interacted with the year shown in the X axis. The omitted category is the set of places that are 200-300 kilometers from the Golden Quadrilateral network. The dependent variables in Panels A and B are log employment in respectively wood-consuming firms and in logging firms. Logging was not specified in the 2005 Economic Census so this point is omitted. In Panels C through E the dependent variable is the share of households' cooking fuel that takes the form of (C) firewood; (D) imported fuels, primarily propane; and (E) crop residue and animal waste. In Panel F, the dependent variable is the share of village land dedicated to agriculture. All estimates are from regressions with state-year fixed effects and standard errors are clustered at the subdistrict level. Appendix Table A12 shows the full set of estimates from the regressions that produced these graphs.

Table 1
Summary Statistics

Village-level Statistics

	Mean	Standard Deviation	Observations
New road before 2011	0.17	0.38	256885
Road completion year	2007	2	45338
Population share with no assets (2002)	0.69	0.31	171249
Population share Scheduled Tribes (2001)	0.22	0.39	256885
Agricultural share of village land (2001)	0.64	0.28	372246
Share energy from firewood (2001)	0.67	0.26	409298
Share energy from imports (2001)	0.07	0.09	409298
Share energy from local nonwood (2001)	0.26	0.26	409298

Subdistrict-level Statistics

	Mean	Standard Deviation	Observations
Average forest cover (2000)	12.76	14.66	4019
Average forest cover (2014)	14.69	14.49	4019
Distance to Golden Quadrilateral	218.50	212.33	4019
Distance to North-South East-West	191.48	155.68	4019
Employment in wood-using firms	141.41	299.77	4019
Employment in logging firms	9.78	92.66	4019

The table shows summary statistics for the samples used for village- and subdistrict-level analyses. Road completion year is shown only for villages that received new roads between 2001 and 2011. The sample for the first four village-level variables consists of the set of villages that did not have a road at baseline. The sample for agricultural land and energy shares consists of all villages with non-zero forest cover at baseline.

Table 2
Regression Discontinuity Estimates of Impact of
Rural Roads on Forest Cover

	First Stage			Reduced Form			IV	
	Any Road	Log Forest	Avg Forest	High Baseline	High ST	Low Assets	Log Forest	Avg Forest
Above Population Threshold	0.185*** (0.011)	-0.002 (0.015)	-0.007 (0.106)	-0.011 (0.018)	-0.008 (0.020)	0.007 (0.023)		
New Road							-0.008 (0.079)	-0.373 (0.703)
N	22365	22365	22365	11214	11174	8875	22368	22368
r2	0.25	0.81	0.62	0.72	0.83	0.80	0.81	0.42

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows regression discontinuity treatment estimates of the effect of new village roads on local forest cover, estimated with Equation 1. In Column 1, the dependent variable is an indicator that takes the value one if a village received a new road in the sample period. Above Population Threshold is an indicator for a village population being above the treatment threshold. Columns 2 through 6 show reduced form estimates of the effect of being above the treatment population threshold. The dependent variables in Columns 2 and 3 respectively are log village forest cover and average covered share of each village pixel; the data source is Vegetation Continuous Fields. Columns 4 through 6 run the log forest cover specification on subgroups defined respectively by (i) above-median forest cover villages; (ii) above median share of Scheduled Tribes in a village; and (iii) below median baseline village assets. Columns 7 and 8 show IV estimates of the treatment effects of new roads, using respectively log and average forest cover as dependent variables. The outcome variable in Columns 2 through 8 is measured in 2013. All estimates include district-population threshold fixed effects and a control for baseline forest cover.

Table 3
Difference-in-Differences Estimates of Impact of
Rural Roads on Forest Cover

	Log Forest		Average Forest	
	(1)	(2)	(3)	(4)
Award Period	-0.005*** (0.002)		-0.033*** (0.013)	
Completion Period	0.002 (0.002)	0.005*** (0.002)	0.009 (0.015)	0.013 (0.012)
District-Year F.E.	Yes	Yes	Yes	Yes
Village F.E.	Yes	Yes	Yes	Yes
N	688275	688275	688275	688275
r2	0.94	0.94	0.92	0.92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows difference-in-differences estimates of the impact of new village roads on local forest cover. We define forest cover as log village forest cover (Columns 1 and 2) and average covered share of each village pixel (Columns 3 and 4); the data source is Vegetation Continuous Fields. The sample consists strictly of villages that received new roads between 2001 and 2013, and were not accessible by paved road in 2001. Award Period is an indicator variable that takes the value one for years after a road contract was awarded and before the road was completed. Completion period is an indicator variable that marks the years after a village's new road was built. All regressions include district*year fixed effects, village fixed effects, baseline population * year fixed effects, and baseline forest * year fixed effects. Standard errors are clustered at the village level to correct for serial correlation.

Table 4
Rural Roads and Deforestation:
Heterogeneity of Difference-in-Differences Estimates

	Baseline Forest		ST Share		Asset Poverty	
	High	Low	High	Low	Poor	Not Poor
Award Period	-0.005*	-0.005***	-0.003	-0.006***	-0.004	-0.006***
	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)	(0.002)
Completion Period	-0.002	0.002	0.000	0.003	0.001	0.001
	(0.004)	(0.002)	(0.003)	(0.003)	(0.004)	(0.003)
N	341280	346455	344010	343860	265470	422430
r2	0.86	0.92	0.93	0.95	0.94	0.95

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows difference-in-differences estimates of the impact of new village roads on local forest cover, along three dimensions of heterogeneity. Forest cover is defined as log village forest cover; the data source is Vegetation Continuous Fields. Columns 1 and 2 respectively show estimates for villages with above and below median baseline forest cover. Columns 3 and 4 respectively show estimates for villages and above and below median population share of members of Scheduled Tribes. Columns 5 and 6 respectively show estimates for below- and above-median shares of households who report no assets in the 2002 Below Poverty Line survey. The sample consists strictly of villages that received new roads between 2001 and 2013, and were not accessible by paved road in 2001. Award Period is an indicator variable that takes the value one for years after a road contract was awarded and before the road was completed. Completion period is an indicator variable that marks the years after a village's new road was built. All regressions include district*year fixed effects, village fixed effects, baseline population * year fixed effects, and baseline forest * year fixed effects. Standard errors are clustered at the village level to correct for serial correlation.

Table 5

Difference-in-Differences Estimates of Impact of
Highways on Forest Cover


	GQ (Treatment)		NSEW (Placebo)	
	Log Forest	Avg Forest	Log Forest	Average Forest
GQ Construction Period * (0-50km)	-0.265*** (0.055)	-1.306*** (0.248)	-0.038 (0.072)	0.183 (0.305)
GQ Construction Period * (50-100km)	-0.278*** (0.056)	-1.215*** (0.247)	-0.001 (0.065)	0.077 (0.286)
GQ Construction Period * (100-150km)	-0.221*** (0.051)	-1.085*** (0.235)	-0.001 (0.060)	-0.176 (0.267)
GQ Construction Period * (150-200km)	-0.102** (0.044)	-0.444** (0.198)	0.023 (0.057)	0.004 (0.264)
GQ Post Period * (0-50km)	-0.210*** (0.061)	-1.161*** (0.253)	-0.013 (0.062)	0.413 (0.319)
GQ Post Period * (50-100km)	-0.185*** (0.060)	-1.023*** (0.249)	0.022 (0.061)	0.106 (0.308)
GQ Post Period * (100-150km)	-0.131** (0.058)	-0.855*** (0.226)	0.019 (0.060)	-0.200 (0.294)
GQ Post Period * (150-200km)	-0.008 (0.051)	-0.198 (0.197)	0.028 (0.060)	-0.012 (0.301)
Distance 0-50km	0.022 (0.021)	0.145 (0.100)	-0.014 (0.020)	0.265 (0.722)
Distance 50-100km	0.021 (0.020)	0.151 (0.099)	-0.021 (0.018)	0.394 (0.717)
Distance 100-150km	0.026 (0.019)	0.133 (0.083)	-0.006 (0.021)	-0.582 (0.784)
Distance 150-200km	0.015 (0.014)	0.075 (0.060)	-0.000 (0.016)	-0.560 (0.697)
N	26766	26766	19062	19062
r2	0.89	0.91	0.92	0.85

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows treatment estimates for the impact of the construction of the GQ highway network on forest cover in the proximity of the highway, according to Equation 5. We define forest cover as log subdistrict forest cover (Columns 1 and 3) and average covered share of each subdistrict pixel (Columns 2 and 4); the data source is Vegetation Continuous Fields. The distance variables are indicators that identify places within a given distance band from the GQ (Columns 1 and 2) or the NS-EW highway network (Columns 3 and 4). The omitted category is the band of places at a distance of 200-300km from the highway network. These distance band indicators are then interacted with time period indicators. The construction period (rows 1 through 4) is 2001 to 2004. The post period (rows 5 through 8) is 2005 to 2008. Columns 3 and 4 estimate a placebo specification with distances to the NS-EW highway network, where construction had barely begun by 2008. The sample includes data from 2000 to 2008; 2000 is the omitted period. We omit years after 2008 as the placebo group is treated in those years. In Columns 3 and 4, we exclude places within 150km of the GQ network to prevent sample contamination. All estimates include state-year fixed effects and standard errors are clustered at the subdistrict level to account for serial correlation.

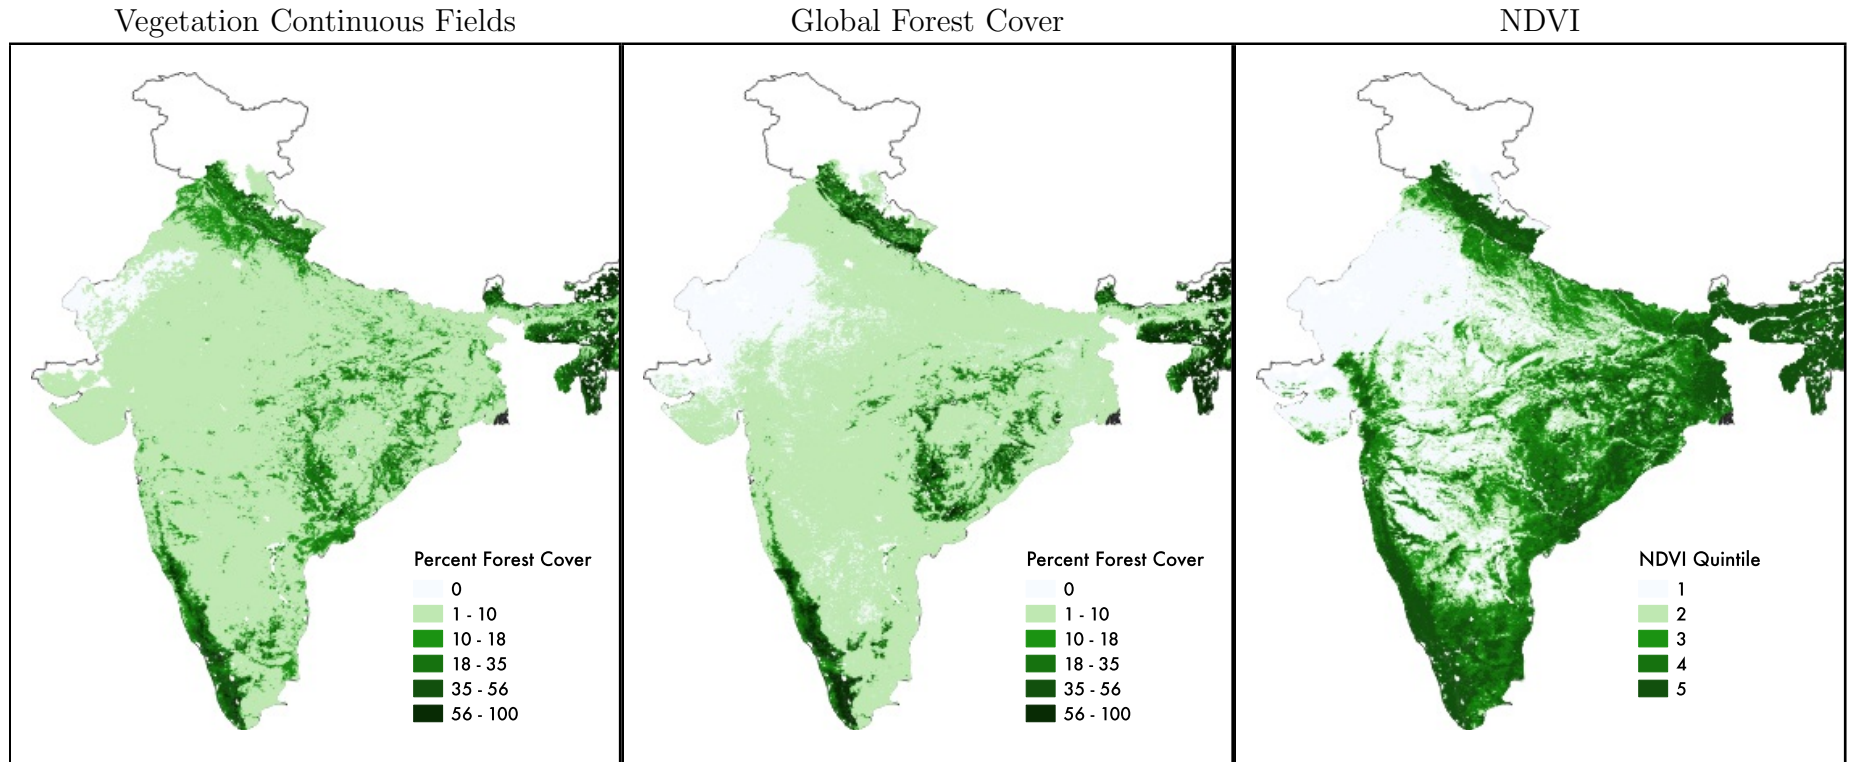
A Online Appendix

Figure A1
Road Program Completion Dates



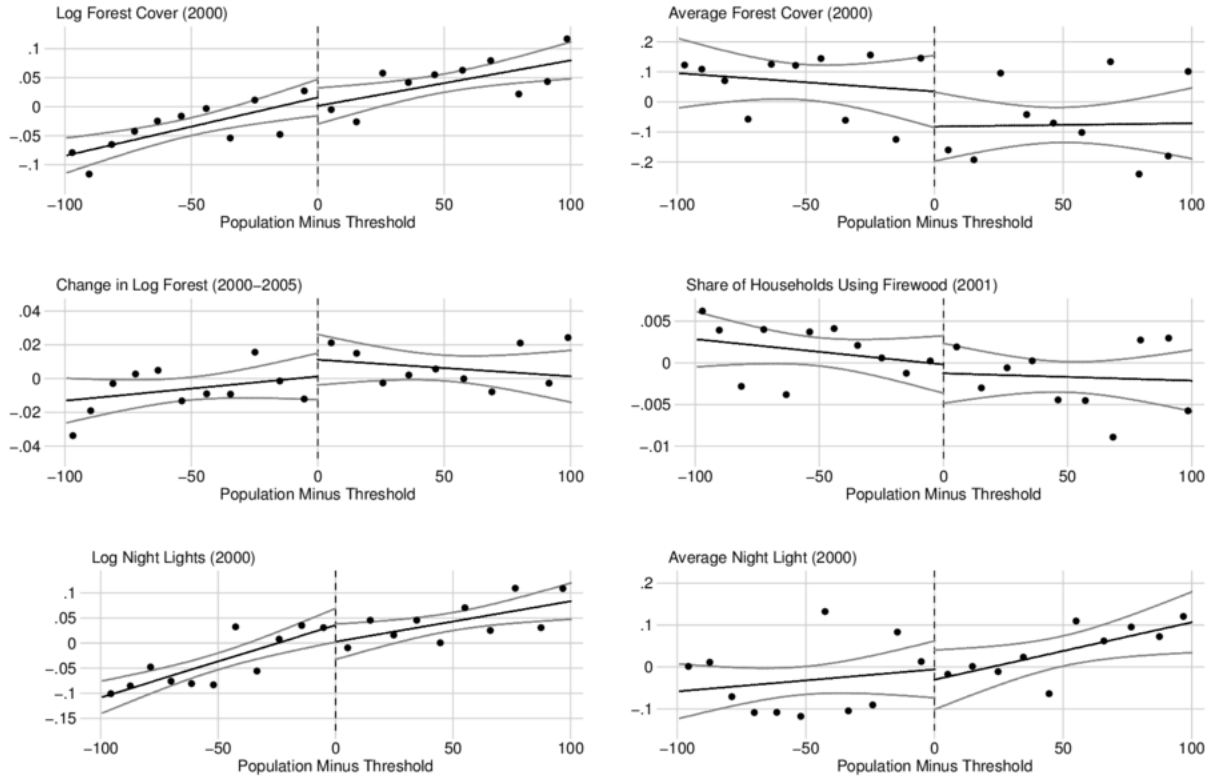
The figure shows the number of roads completed under the PMGSY road construction program, by year.

Figure A2
Forest Cover in 2000 According to Different Remote Sensing Sources




The figure compares three different remote sensing measures of forest cover in India for the year 2000. Gridded forest cover data are aggregated to village and town polygons according to the 2011 population census. Vegetation Continuous Fields and Global Forest Cover can be directly interpreted as the average share of land covered by forest in the year 2000. For NDVI, we report quintiles of the NDVI index measured in February of 2000, the time of year when forest is most distinguishable from crop cover (Foster and Rosenzweig, 2003).

Figure A3
Regression Discontinuity Balance Tests



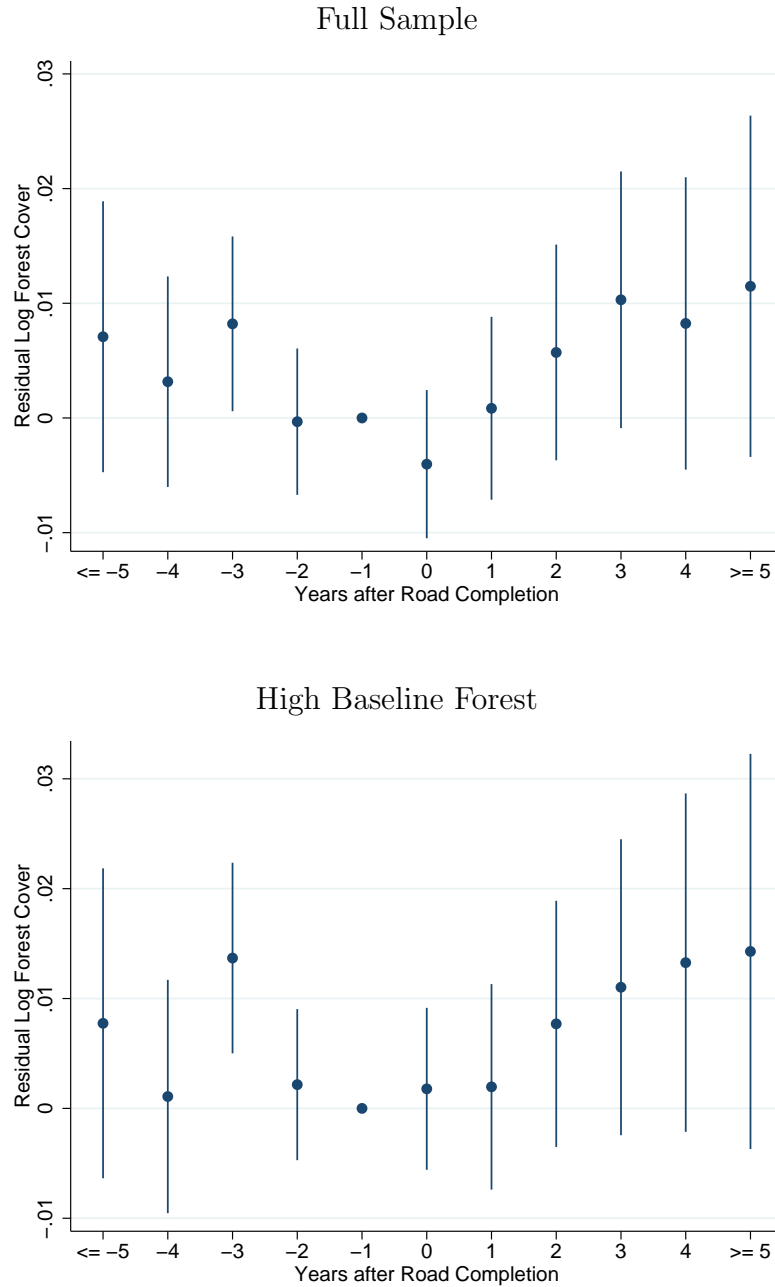
The figure displays a graphical form of the regression discontinuity balance test. Each graph shows the means of a variable measured at baseline in bins defined by population relative to the rural road program treatment threshold. The linear fits and standard errors are estimated from Equation 1. The vertical line shows the treatment threshold; the jump in the fit at this line is the regression discontinuity treatment estimate. The dependent variable in each panel (left-to-right, top-to-bottom) is (A) log forest cover in 2000; (B) average forest cover in 2000; (C) change in log forest cover from 2000 to 2005; (D) the share of households whose primary cooking fuel is firewood; (E) the log of night light luminosity in 2000; and (F) average night light luminosity in 2000. In Panel C, we omit villages with roads built before 2006, to ensure that balance estimates are not contaminated by the small number of treated villages before this date. All estimates include district-population threshold fixed effects.

Figure A4
Regression Discontinuity Density Test



The figure displays a graph from a regression discontinuity density test (McCrary, 2008). The X axis shows the population relative to the road program treatment eligibility threshold. The Y axis shows a kernel estimate of the density of villages in a given normalized population band. The lines display non-parametric fits to the density function along with 95% confidence intervals.

Figure A5
 Difference-in-Differences Estimates of Impact of
 Rural Roads on Forest Cover (Long Panel)



The figure shows year-by-year estimates of log forest cover in villages that received new roads between 2001 and 2013. The figure is identical to Figure 3, but with an additional estimate for the 5th year before and after treatment. Villages are grouped on the X axis according to the year relative to road completion. Each point thus shows the average value of log forest cover in villages in a given year relative to the treatment year, controlling for village fixed effects, district*year fixed effects and baseline population * year and baseline log forest cover * year interactions. Standard errors are clustered at the village level.

Table A1
 OLS Regressions of Forest Cover on Rural Road Indicators

	(1)	(2)	(3)	(4)	(5)
Paved Road in 2001	-0.162*** (0.010)	-0.270*** (0.010)	-0.041*** (0.009)	-0.025*** (0.009)	-0.018** (0.009)
Population		0.476*** (0.013)	0.862*** (0.012)	0.896*** (0.012)	0.910*** (0.012)
Population ²		-0.037*** (0.004)	-0.085*** (0.003)	-0.092*** (0.003)	-0.095*** (0.003)
Distance in km to town of 10,000					0.007*** (0.000)
Distance in km to town of 100,000					-0.000 (0.000)
Constant	3.405*** (0.004)	3.083*** (0.008)	2.786*** (0.007)	2.764*** (0.007)	2.533*** (0.013)
Fixed Effects	None	None	State	State	District
N	270871	270871	270871	270871	270871
r ²	0.00	0.02	0.19	0.28	0.28

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows estimates from OLS regressions of village-level log forest cover in 2001 on an indicator variable that takes the value one if a village has a paved road in 2001. Column 1 presents the bivariate estimates, and Columns 2 through 5 present estimates with progressively greater numbers of controls and fixed effects. Forest cover is calculated from Vegetation Continuous Fields. Population is measured in millions of people.

Table A2
Regression Discontinuity Balance Tests

Variable	RD Estimate	N
Log Forest (2000)	-0.012 (0.027)	22365
Average Forest (2000)	-0.093 (0.101)	22365
Share Cooking with Firewood	-0.001 (0.003)	22317
Log Forest Change (2000-2005)	0.012 (0.012)	22365
Mean Night Light (2000)	-0.019 (0.060)	21542
Log Night Light (2000)	-0.032 (0.029)	21542

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows estimates from a regression discontinuity balance test. We run the regression discontinuity specification defined by Equation 1 on variables measured before any rural road construction took place, and report the reduced form treatment estimates. Row 4 (Log Forest Change 2000-2005) tests for pretrends in forest cover; we exclude villages with roads built before 2006 for this sample. All estimates include district-population threshold fixed effects.

Table A3
Regression Discontinuity Estimates of Impact of
Rural Roads on Forest Cover (Heterogeneity Tests)

	<u>Dist. to Town</u>		<u>Market Access</u>		<u>Logging</u>	<u>Ind. Wood Use</u>
	Low	High	Low	High	High	High
Above Population Threshold	-0.008 (0.019)	0.005 (0.022)	-0.001 (0.023)	-0.001 (0.021)	-0.002 (0.015)	0.021 (0.020)
N	11180	11166	10076	10074	22365	11283
r ²	0.82	0.81	0.81	0.81	0.81	0.81

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows reduced form regression discontinuity estimates of the impact of rural roads on forest cover. The specifications are comparable to those in Columns 2 and 3 of Table 2 but for subsamples related to treatment heterogeneity of interest. Columns 1 and 2 respectively show estimates for villages with above and below median distance to the nearest town of 100,000 or higher. Columns 3 and 4 respectively show estimates for villages with above and below median market access. Market access is measured with a trade elasticity of 8, the value suggested in Donaldson and Hornbeck (2016); results are similar under other elasticities. Columns 5 and 6 restrict the sample respectively to villages with above median employment in logging and employment in wood-using industries. The outcome variable is log forest cover in 2013. All estimates include district-population threshold fixed effects and a control for baseline forest cover.

Table A4
Regression Discontinuity Estimates of Impact of
Rural Roads on Forest Cover (Alternate Bandwidths)

	Log Forest (2013)				Average Forest (2013)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Above Population Threshold	0.013 (0.021)	-0.002 (0.015)	0.000 (0.012)	-0.004 (0.010)	0.194 (0.151)	-0.007 (0.106)	-0.003 (0.088)	-0.015 (0.077)
Bandwidth	50	100	150	200	50	100	150	200
N	11275	22365	33474	44571	11275	22365	33474	44571
r ²	0.81	0.81	0.81	0.81	0.61	0.62	0.61	0.61

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows reduced form regression discontinuity estimates of the impact of rural roads on forest cover. The specifications are comparable to those in Columns 2 and 3 of Table 2 but with alternate bandwidths defined in the bandwidth row of the table. The outcome variable is log forest cover in 2013. All estimates include district-population threshold fixed effects and a control for baseline forest cover.

Table A5
Regression Discontinuity Estimates of Impact of
Rural Roads on Household Fuel Use

	Imports	Local Non-Wood	Firewood
Above Population Threshold	-0.002 (0.002)	0.002 (0.006)	0.000 (0.006)
N	22317	22317	22317
r2	0.28	0.42	0.42

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows reduced form regression discontinuity treatment estimates of the effect of new village roads on village-level household fuel use, estimated with Equation 1. The dependent variable is the share of households in a village that use imported fuel sources (primarily propane, Column 1); dung and crop residue (Column 2); and firewood (Column 3) as primary fuel sources for cooking. The dependent variables are measured in 2011. In addition to district-population threshold fixed effects, controls include the baseline fuel share reported in 2001 (at the subdistrict level) and forest cover in 2000.

Table A6
Difference-in-Differences Estimates of Impact of
Rural Roads on Forest Cover (Robustness Tests)

	(1)	(2)	(3)	(4)
Award Period	-0.004** (0.002)	-0.006*** (0.002)	-0.010*** (0.003)	-0.007*** (0.002)
Completion Period	0.001 (0.002)	-0.000 (0.002)	-0.002 (0.004)	-0.002 (0.003)
District-Year F.E.	Yes	No	Yes	Yes
Subdistrict-Year F.E.	No	Yes	No	No
Village F.E.	Yes	Yes	Yes	Yes
Village Time Trends	Yes	No	No	No
Panel Sample	Full	Full	+/- 5 Years	+/- 4 Years
N	688275	681555	374010	481665
r ²	0.95	0.96	0.94	0.94

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows difference-in-differences estimates of the impact of new village roads on local forest cover, under alternate sample definitions. We define forest cover as log village forest cover; the data source is Vegetation Continuous Fields. Specifications are identical to those in Table 3 with the following changes. Column 1 includes village-specific time trends. Column 2 uses subdistrict-year fixed effects instead of district-year fixed effects. Column 3 restricts the sample to villages with roads for which we can observe at least 5 years of data before road completion and 5 years after. Column 4 does the same, with 4 years. The sample consists strictly of villages that received new roads between 2001 and 2013, and were not accessible by paved road in 2001. Award Period is an indicator variable that takes the value one for years after a road contract was awarded and before the road was completed. Completion period is an indicator variable that marks the years after a village's new road was built. All regressions include district*year fixed effects, village fixed effects, baseline population * year fixed effects, and baseline forest * year fixed effects. Standard errors are clustered at the village level to correct for serial correlation.

Table A7
Difference-in-Differences Estimates of Impact of
Rural Roads on Forest Cover (Alternate Specifications)

	(1)	(2)	(3)	(4)	(5)	(6)
Award Period	0.004*** (0.001)	-0.002 (0.001)	-0.005*** (0.002)	-0.004** (0.002)	-0.005** (0.002)	-0.001 (0.001)
Completion Period	0.017*** (0.001)	0.000 (0.002)	0.002 (0.002)	0.001 (0.002)	0.003 (0.003)	-0.000 (0.002)
District-Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Village F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Village Time Trends.	No	Yes	No	Yes	No	No
Village Definition	Boundary	Boundary	Boundary	Boundary	5 km radius	50 km radius
N	3359370	3359370	689745	689745	688275	688275
r2	0.95	0.96	0.94	0.95	0.94	0.97

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows difference-in-differences estimates of the impact of new village roads on local forest cover, under alternate sample definitions. We define forest cover as log village forest cover; the data source is Vegetation Continuous Fields. Specifications are identical to those in Table 3 with the following changes. Column 1 includes villages that did not receive PMGSY roads at any time as part of the control group. Column 2 adds village-specific time trends to show that the positive treatment estimate in Column 1 is driven by differential trends in never-treated villages. Columns 3 and 4 repeat these two specifications with the standard set of villages to show that village-specific time trends do not affect our main estimates. Column 5 estimates the standard specification from Table 3 with treated villages only, but the dependent variable includes forest cover in a 5km radius from the village centroid. Column 6 uses a 50km centroid. Award Period is an indicator variable that takes the value one for years after a road contract was awarded and before the road was completed. Completion period is an indicator variable that marks the years after a village's new road was built. All regressions include village fixed effects, baseline population * year fixed effects, and baseline forest * year fixed effects. Standard errors are clustered at the village level to correct for serial correlation.

Table A8
Rural Roads and Deforestation:
Heterogeneity of Difference-in-Differences Estimates

	<u>Dist. to Town</u>		<u>Market Access</u>		<u>Logging</u>	<u>Ind. Wood Use</u>
	Low	High	Low	High	High	High
Award Period	-0.004*	-0.006**	-0.010***	-0.002	-0.010***	-0.007***
	(0.002)	(0.003)	(0.003)	(0.002)	(0.003)	(0.002)
Completion Period	-0.002	0.005	0.006	-0.004	-0.002	-0.001
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)
N	343950	343770	298665	298575	205020	343215
r2	0.95	0.94	0.94	0.94	0.95	0.94

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows difference-in-differences estimates of the impact of new village roads on local forest cover, along five additional dimensions of heterogeneity. Forest cover is defined as log village forest cover; the data source is Vegetation Continuous Fields. Columns 1 and 2 respectively show estimates for villages with above and below median distance to the nearest town of 100,000 or higher. Columns 3 and 4 respectively show estimates for villages with above and below median market access. Market access is measured with a trade elasticity of 8, the value suggested in Donaldson and Hornbeck (2016); results are similar under other elasticities. Columns 5 and 6 restrict the sample respectively to villages with above median employment in logging and employment in wood-using industries. The sample consists strictly of villages that received new roads between 2001 and 2013, and were not accessible by paved road in 2001. Award Period is an indicator variable that takes the value one for years after a road contract was awarded and before the road was completed. Completion period is an indicator variable that marks the years after a village's new road was built. All regressions include district*year fixed effects, village fixed effects, baseline population * year fixed effects, and baseline forest * year fixed effects. Standard errors are clustered at the village level to correct for serial correlation.

Table A9

Difference-in-Differences Estimates of Impact of
Highways on Forest Cover:
Straight Line Instrumental Variables

	GQ (Straight Line)		NSEW (Straight Line)	
	Log Forest	Avg Forest	Log Forest	Average Forest
GQ Construction Period * (0-50km)	-0.144*** (0.026)	-0.442*** (0.142)	-0.013 (0.055)	0.379 (0.314)
GQ Construction Period * (50-100km)	-0.181*** (0.027)	-0.749*** (0.142)	-0.010 (0.052)	0.199 (0.297)
GQ Construction Period * (100-150km)	-0.129*** (0.029)	-0.495*** (0.164)	0.014 (0.048)	-0.336 (0.272)
GQ Construction Period * (150-200km)	-0.086*** (0.031)	-0.362** (0.178)	-0.030 (0.044)	-0.157 (0.268)
GQ Post Period * (0-50km)	-0.113*** (0.026)	-0.422*** (0.133)	-0.015 (0.053)	0.528 (0.332)
GQ Post Period * (50-100km)	-0.110*** (0.027)	-0.518*** (0.137)	-0.061 (0.054)	-0.051 (0.329)
GQ Post Period * (100-150km)	-0.057** (0.028)	-0.138 (0.154)	-0.027 (0.051)	-0.374 (0.316)
GQ Post Period * (150-200km)	-0.050* (0.029)	-0.186 (0.162)	-0.051 (0.042)	-0.361 (0.320)
Distance 0-50km	0.019*** (0.007)	0.179*** (0.051)	-0.056*** (0.017)	-0.166 (0.801)
Distance 50-100km	0.033*** (0.008)	0.282*** (0.055)	-0.057*** (0.016)	1.903** (0.886)
Distance 100-150km	0.033*** (0.007)	0.273*** (0.051)	-0.045*** (0.016)	0.480 (0.819)
Distance 150-200km	0.024*** (0.005)	0.195*** (0.039)	-0.019* (0.011)	-0.342 (0.866)
N	26397	26397	14958	14958
r ²	0.90	0.90	0.94	0.86

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows reduced form estimates from regressions of distance band * time period interactions on forest cover. Distance bands are calculated to straight line approximations of the Golden Quadrilateral (Columns 1 and 2) and North-South/East-West (Columns 3 and 4) highway corridors. The estimating equation is Equation 5. We define forest cover as log subdistrict forest cover (Columns 1 and 3) and average covered share of each subdistrict pixel (Columns 2 and 4); the data source is Vegetation Continuous Fields. The omitted distance category is the set of subdistricts at a distance of 200-300km from each set of straight line approximations. The construction period (rows 1 through 4) is 2001 to 2004. The post period (rows 5 through 8) is 2005 to 2008. Columns 3 and 4 estimate a placebo specification with distances to the NS-EW highway network, where construction had barely begun by 2008. The sample includes data from 2000 to 2008; 2000 is the omitted period. We omit years after 2008 as the placebo group is treated in those years. In Columns 3 and 4, we exclude places within 150km of the GQ network to prevent sample contamination. All estimates include state-year fixed effects and standard errors are clustered at the subdistrict level to account for serial correlation.

Table A10
 Difference-in-Differences Estimates of Impact of
 Highways on Forest Cover: Exploiting Construction Timing

	(1)	(2)
Distance _{GQ} ≤ 100 km	-0.085*** (0.013)	-0.114*** (0.017)
Construction Year	0.056*** (0.014)	0.128*** (0.019)
Post-Construction Year	0.152*** (0.021)	0.215*** (0.026)
Construction Year * Distance _{GQ} ≤ 100 km	-0.048*** (0.014)	-0.126*** (0.019)
Post-Construction Year * Distance _{GQ} ≤ 100 km	-0.044*** (0.014)	-0.116*** (0.021)
Distance _{GQ} ∈ (100, 200) km		-0.038** (0.016)
Construction Year * Distance _{GQ} ∈ (100, 200) km		-0.114*** (0.019)
Post-Construction Year * Distance _{GQ} ∈ (100, 200) km		-0.107*** (0.022)
N	41730	41730
r ²	0.88	0.88

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows alternate treatment estimates for the impact of the construction of the GQ highway network on forest cover in the proximity of the highway, which exploit the timing of construction of different segments. The estimating equation takes the form: $Forest_{ist} = \beta_0 + \beta_1 CLOSE_{is} + \beta_2 CONSTR_t + \beta_3 POST_t + \beta_4 CLOSE_{is} * CONSTR_t + \beta_5 CLOSE_{is} * POST_t + \epsilon_{ist}$. $CLOSE_{is}$ takes the value one if a subdistrict is within 100 km of a segment of the GQ, and $POST_t$ refers to a time period after construction. The interaction between $CLOSE_{is}$ and $CONSTR_t$ describes the impact of GQ construction on a subdistrict's forests, and the interaction with $POST_t$ describes the impact after construction of the segment is complete. In the table, Construction Year and Post-Construction Year refer to the construction and completion period for the nearest GQ segment to a subdistrict centroid. The Distance variables are indicators that take the value 1 if the nearest segment is within the given distance band. The sample is the set of subdistricts with centroids within 300km of the GQ. The omitted category is thus the band of places at a distance of 200-300km from the highway network. We define forest cover as log subdistrict forest cover; the data source is Vegetation Continuous Fields. All estimates include state-year fixed effects and baseline subdistrict population, forest cover and town distance interacted with forest observation year. Standard errors are clustered at the subdistrict level to account for serial correlation.

Table A11
Mechanism Tests for Impact of Highways on Deforestation:
Employment in Wood-Using Firms

	Wood Use	Logging
(0-100km from GQ) * 1(Year == 1990)	-0.012 (0.088)	0.171*** (0.056)
(100-200km from GQ) * 1(Year == 1990)	-0.224*** (0.080)	0.178*** (0.055)
(0-100km from GQ) * 1(Year == 1998)	0.003 (0.090)	0.159** (0.062)
(100-200km from GQ) * 1(Year == 1998)	-0.270*** (0.088)	0.151*** (0.057)
(0-100km from GQ) * 1(Year == 2005)	0.271*** (0.082)	
(100-200km from GQ) * 1(Year == 2005)	-0.088 (0.077)	
(0-100km from GQ) * 1(Year == 2013)	0.417*** (0.089)	0.306*** (0.080)
(100-200km from GQ) * 1(Year == 2013)	0.015 (0.082)	0.298*** (0.075)
N	11724	8793
r ²	0.38	0.28

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows estimates of the impact of the Golden Quadrilateral highway on log employment in timber-related firms. The table shows the full specifications used to generate Figure 6, panels A and B. The dependent variable is log employment in firms for which timber is the primary input (sawmilling, pulp and paper, manufacture of wooden containers, wooden furniture and cork boards, Column 1) and log employment in logging firms (Column 2). Each row shows the interaction of an indicator for a given distance band from the Golden Quadrilateral, interacted with an indicator for a given Economic Census year. The omitted distance category is 200-300km. The estimates thus show the difference between log employment in each sector/year/distance band with log employment in the same sector/year at distance 200-300km from the highway. Logging is not specifically identified in the 2005 Economic Census, so this estimate is omitted. All regressions include state-year fixed effects and cluster standard errors at the subdistrict level.

Table A12
Mechanism Tests for Impact of Highways on Deforestation:
Land and Fuel Use

	Ag Land Share	Fuel: Firewood	Fuel: Imported	Fuel: Local Non-wood
(0-100km from GQ) * 1(Year == 1991)	8.218*** (1.746)			
(100-200km from GQ) * 1(Year == 1991)	1.217 (1.479)			
(0-100km from GQ) * 1(Year == 2001)	6.989*** (2.575)	-0.015 (0.010)	0.010** (0.004)	0.005 (0.010)
(100-200km from GQ) * 1(Year == 2001)	-3.972* (2.281)	0.051*** (0.009)	-0.006* (0.004)	-0.045*** (0.008)
(0-100km from GQ) * 1(Year == 2011)	-2.103 (2.427)	-0.021* (0.011)	0.022*** (0.005)	-0.001 (0.010)
(100-200km from GQ) * 1(Year == 2011)	-5.070** (2.225)	0.044*** (0.009)	0.004 (0.004)	-0.047*** (0.009)
N	8793	5862	5862	5862
r2	0.69	0.59	0.43	0.67

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows estimates of the impact of the Golden Quadrilateral highway on land and fuel use. The table shows the full specifications used to generate Figure 6, panels C through F. The dependent variable is the share of village land dedicated to agriculture (Column 1); the share of households in a village that use firewood (Column 2); imported fuel sources (primarily propane, Column 3); and dung and crop residue (Column 4) as primary fuel sources for cooking. Each row shows the interaction of an indicator for a given distance band from the Golden Quadrilateral, interacted with an indicator for a given Population Census year. The omitted distance category is 200-300km. The estimates thus show the difference between the outcome variable in each sector/year/distance band with the outcome variable in the same sector/year at distance 200-300km from the highway. All regressions includes state-year fixed effects and cluster standard errors at the subdistrict level.