Working Paper Series 807
(ISSN 2788-0443)

Post(-Mongol) Roads to Path Dependence

Sebastian Ottinger
Elizaveta Zelnitskaia

CERGE-EI
Prague, December 2025

Post(-Mongol) Roads to Path Dependence*

Sebastian Ottinger
(CERGE-EI and IZA)

Elizaveta Zelnitskaia
(CERGE-EI)

PRELIMINARY DRAFT

This version: Wednesday 19th November, 2025

Abstract

Why do cities emerge where they do? This paper exploits a rule-based transport network in Imperial Russia to study the origins of urban centers. The yams postal system, introduced by the Mongols in the thirteenth century and maintained by Muscovy, required relay stations in regular intervals to change horses, creating an infrastructure grid whose spacing reflected logistics rather than geography or pre-existing settlements. We digitize all stations listed in the 1777 Russian Road Guide along a sample of 15 major routes, and divide rays between consecutive stops into 0.5 km cells. In modern satellite data, cells located at the historical interval where horses were changed are about thirty percent brighter today than neighboring cells before or after that range. The effect is robust to first- and second-nature controls, ray fixed effects, and controlling of pre-1800 settlements, and is absent for the later Trans-Siberian Railway. Additional analyses show that subsequent city growth correlates little with geographic endowments, but was amplified by later infrastructure investments, suggesting that administrative accidents – not natural advantages – seeded some of Russia's urban geography. The findings illustrate how spatial inequality can arise from arbitrary historical coordination points, with lasting consequences for the distribution of economic activity.

JEL: N73, O18, R11, R12, H11

Keywords: City Location, Path Dependence, Transport Infrastructure, Natural Advantage

*We thank Štěpán Jurajda, Nikolas Mittag, David Kristian Nagy, Christian Ochsner, Giacomo Ponzetto, Jakub Steiner, and Vlada Kosenkova for valuable comments and suggestions, as well as participants of the CERGE-EI Brown Bag seminars. Any suggestions are welcome. All remaining errors are our own. Corresponding author: Sebastian Ottinger (sebastian.ottinger@gmail.com)

1 Introduction

Why do cities emerge where they do? In most countries across the world, economic activity and innovation are concentrated in a small number of cities.[1] Yet the spatial distribution of these urban centers is remarkably persistent over time: many modern cities occupy sites that were already important in medieval or even ancient periods. Understanding why some locations became focal points of settlement and others did not is a central question in urban and economic geography. Two broad views dominate this debate. One emphasizes *natural advantages* – features such as navigable rivers, fertile soils, or resource endowments that make particular sites inherently productive (Henderson, 1974; Redding and Rossi-Hansberg, 2017). The other highlights *historical accidents* – events that placed initial activity in otherwise ordinary locations, which then persisted through path dependence and agglomeration forces (Davis and Weinstein, 2002; Bleakley and Lin, 2012). Distinguishing between these explanations is empirically difficult because geography and history are deeply intertwined.

The cleanest way to disentangle nature from history would be the following ideal experiment: randomly select locations across a country, establish small settlements, and then wait for several centuries to see whether these became permanent settlement and assess which grew into major urban centers. Such an experiment would reveal whether intrinsic geographic advantages or sheer historical contingency determine long-run urban patterns. Of course, no government ever conducted such a randomized trial. Yet in rare historical circumstances, administrative or technological constraints created settlement patterns that approximate random assignment. When infrastructure had to be placed at fixed intervals determined by logistics rather than local productivity – such as the spacing of rest points for horses – officials effectively scattered potential city seeds across the landscape, largely independent of geography or pre-existing population.

This paper turns to a particular – but also particularly suited – setting to approximate such an experiment: the yams postal network of Imperial Russia. The system, introduced

[1] In 2023, the world's 1,000 largest cities generated 60% of global GDP, yet they were home to only about 30% of the global population. Urbanization is accelerating worldwide: since 2007, more people have lived in cities than in rural areas, and in high-income countries, roughly four out of five people now reside in urban areas. (Oxford Economics, 2024; Ritchie, Samborska, and Roser, 2024). Similar patterns are observed in developing countries, where urban primacy – the share of a nation's urban population living in its largest city – is often over 30% and closely linked to the concentration of economic activity, with the largest city contributing an even larger share of GDP and innovation output (Frick, 2018).

by the Mongols in the thirteenth century and maintained by Muscovy thereafter, required couriers to change horses at regular distances. These relay stations, spaced roughly thirty kilometers apart, formed a rule-based transport grid designed solely for communication efficiency. Crucially, while similar postal systems existed elsewhere – in Persia, China, and later across Europe – those were typically built over densely settled land where stops coincided with existing towns, or their detailed records have not survived. In contrast, the Russian system was implemented across vast, sparsely populated regions, and its 1777 Russian Road Guide provides a comprehensive list of stations and routes. This combination of low prior settlement density and precise administrative documentation makes the Russian post roads an unusually clean setting to study how accidental, rule-driven infrastructure placement could seed enduring patterns of urban activity.

To study the long-run imprint of this network, we digitize all stations listed for 15 major postal routes in the Foreign and Russian Road Guide published by Ruban in 1777 – the first comprehensive enumeration of postal routes in the Russian Empire. The guide details the sequence and spacing of stops on each road. For each pair of consecutive stops, we draw a straight line connecting them and divide it into 0.5-kilometer cells, yielding roughly forty thousand observations. These cells form the empirical canvas on which we observe modern economic activity using satellite night-light intensity, a widely used and highly granular proxy for local population and income. We further merge each cell with measures of natural geography – elevation, soil and caloric suitability, and presence of rivers – and with indicators for modern infrastructure such as railways and highways.

Our empirical strategy exploits the mechanical regularity of the postal system's design. Because the next station was expected to occur at a fixed interval from the previous one, the distance alone predicts where a station would likely have been built, regardless of local geography. This feature allows us to compare modern brightness along each historical route as a function of distance from the preceding station, controlling for route fixed effects and all observable characteristics. The identifying idea is straightforward: if station placement was exogenous with respect to geography, any discontinuous rise in modern activity should occur precisely at the historical horse-changing interval. The results reveal a remarkably persistent spatial imprint. Along every route, night-light intensity today rises sharply at the distance where horses once had to be changed: cells located within roughly thirty kilometers (25-35 kilometers) of the previous station are about thirty percent brighter than neighboring cells before or after that range. The effect is statistically strong and robust to controlling for both first-nature endowments and second-nature infrastructure. We document extensive robustness to shifting the precise kilometer range, expanding or shrinking the ray length, and aggregating the analysis to larger segments.

To sharpen inference, we also use the 25–35-kilometer range as an instrument for the actual presence of a historical station, capturing the causal effect of rule-based station placement on long-term urban development. Instrumental-variable estimates yield even larger magnitudes, reinforcing the interpretation that administrative design – rather than natural advantage – determined which places became enduring centers of population. What began as logistical waypoints on an imperial communication grid thus evolved into a durable backbone of Russia's urban geography.

We examine several alternative explanations for the observed pattern and show that none can account for the persistence we document. The first concern is that post stations may not have been placed randomly at all, but simply located in or near pre-existing population centers. If so, the apparent effect of the postal network would merely reflect older settlement patterns rather than new ones created by administrative design. Several pieces of evidence argue against this interpretation. Eighteenth-century Russia was exceptionally sparsely populated, and large stretches of the postal network crossed entirely unsettled territory. To verify that pre-existing habitation does not drive the result, we control for monasteries founded before 1800. The brightness peak at the horse-changing interval remains almost unchanged, suggesting that the urban imprint we observe originated with the postal system itself.

A second explanation is that the results could simply reflect the influence of proximity to major urban centers such as Moscow or St. Petersburg. If post stations near these hubs grew faster because of agglomeration spillovers or market access rather than the spacing rule, the pattern could arise mechanically from centrality. To rule this out, we exclude all routes leading to Moscow and St. Petersburg, and in other specifications drop any stations that later developed into large cities. The characteristic brightness spike persists in these restricted samples, appearing just as clearly on peripheral routes as on core ones. The effect therefore reflects a general feature of the postal system, not a by-product of metropolitan growth.

A third possibility is that the 25–35 km spacing merely captures a universal regularity in how settlements form – a generic "system of cities" pattern driven by transport costs or the spatial range of economic interactions, rather than any specific institutional rule. Classic theories of urban hierarchies (Henderson, 1974; Fujita, Krugman, and Venables, 1999) predict that equilibrium distances between towns depend on the technology of transport: when movement is costly, many small centers arise at short intervals, while cheaper or faster transport allows fewer, more widely spaced hubs. To examine whether our pattern simply reflects such technology-driven spacing, we replicate the analysis for the Trans-Siberian Railway, constructed under entirely different technological constraints in the late nineteenth century. Along that network, the peak in modern brightness occurs around ten kilometers – the typical distance between water stops for steam locomotives – and no effect appears at thirty kilometers. This comparison indicates that the spacing of economic activity mirrors the underlying transport technology: what was thirty kilometers for horses became ten kilometers for steam. The regularity we document thus reflects a specific logistical legacy of the horse-based *yams* system rather than a universal equilibrium distance.

Having established that the initial placement of post stations generated quasi-random "city seeds," we next examine why some of these settlements prospered while others remained small. To do so, we focus on the relevant 25–35 km range along each route – the interval where a post station was typically located – and aggregate all variables to this *city-seed* level. This coarser perspective allows us to compare segments across rays and relate their modern development to geographic and infrastructural characteristics. Our analysis reveals two main insights. First, natural geography matters only weakly: elevation and overall caloric suitability are positively related to modern brightness, but most crop-specific suitabilities are insignificant or even negative. The only clear positive association is with potato suitability – a crop introduced long after the network was designed – confirming that planners could not have selected locations based on latent agronomic potential. Second, the presence of roads or railways is strongly correlated with modern settlement size, suggesting that subsequent transport infrastructure investment "locked in" city seeds . Together, these findings suggest that while favorable geography modestly influenced which city seeds thrived, the dominant driver of persistence was administrative design itself – later reinforced, rather than determined, by second-nature infrastructure.

Taken together, the evidence paints a consistent picture of how an administrative system designed for speed of communication inadvertently shaped the spatial foundations of the

Russian economy. By tracing the locations of eighteenth-century postal relay stations and comparing economic activity along precise geographic rays, the paper isolates an episode in which potential city sites were assigned according to a simple logistical rule rather than natural geography. The resulting settlements – initially established for horse exchange rather than production – became enduring urban centers whose presence remains visible in satellite data centuries later. By combining detailed historical sources with modern spatial data, and by exploiting the mechanical spacing of a pre-industrial transport network, the analysis provides the first quasi-experimental evidence that accidents of administrative design, rather than inherent advantages, played a decisive role in determining where cities in Russia first emerged, and that subsequent, second nature transport infrastructure investments shaped how their pattern persisted over time.

Related Literature. Our work connects to several strands of research on the origins and persistence of urban systems. A first body of work emphasizes *first-nature geography* as the foundation of city formation. Natural advantages such as access to waterways, defensible terrain, or fertile soils are seen as the 'seeds" from which urban centers grew and later persisted through cumulative agglomeration forces. For Europe, Bosker and Buringh (2017) show that medieval city locations closely align with advantageous geographic features, while Düben and Krause (2023) document similar mechanisms for imperial China. More recent work has pushed the frontier of measurement: Nguyen, Brakman, Garretsen, and Kohl (2023) infer initial geographical conditions from medieval German toponyms and show that references to rivers, hills, and clearings predict later population size. These studies collectively argue that modern urban sizes and hierarchies mirror deep natural fundamentals.[2] Yet geography and history are often inseparable: most cities arose where both natural endowments and early institutions interacted, making it difficult to isolate purely geographic causation.

A second literature stresses the role of *institutional or technological "seeds"* – instances where human design, rather than nature, determined where settlements first appeared. Infrastructure and policy choices can create durable spatial hierarchies even when their original purposes fade. For example, political centralization often leaves a lasting imprint on the urban system: capital cities tend to be disproportionately large and economically dominant

[2] Michaels and Rauch (2018) documents how important city location can be for national economic growth.

relative to other cities of similar geographic potential (e.g. Ades and Glaeser, 1995; Henderson and Wang, 2007). Beyond institutional factors, a range of historical shocks and other state interventions — including the introduction of roads or railroads, the drawing of new, even short-lived, borders, or direct policies for city formation — left long-lasting marks on spatial development (Donaldson and Hornbeck, 2016; Jedwab and Moradi, 2016; Heblich, Redding, and Sturm, 2020; Dalgaard, Kaarsen, Olsson, and Selaya, 2022; Nagy, 2022; Ochsner, 2025; Schweiger, Stepanov, and Zacchia, 2022; Glaeser and Ponzetto, 2018). These studies illustrate how exogenous administrative or technological shocks can reshape the map of economic activity. Our paper complements this perspective by examining a much earlier and arguably purer case of mechanical design: the Mongol and Muscovite postal network, whose horse-relay stations were placed at nearly fixed intervals dictated by travel logistics. Unlike deliberate town founding or route optimization, the spacing rule was invariant to local productivity or geography, creating an unusually clean source of quasi-random settlement assignment.

A third strand of research concerns *spatial path dependence*. Once a location gains an early advantage – whether geographic, institutional, or accidental – subsequent agglomeration and learning can perpetuate its prominence (Bleakley and Lin, 2012; Dalgaard et al., 2022; Allen and Donaldson, 2020). Our contribution is to demonstrate a new mechanism of such persistence: a rule-based administrative system that created settlement sites independent of geography.

Fourthly, our setting contributes to the growing literature on the *economic history of the Russian Empire* and the spatial foundations of its development. Recent research shows how Russia's regional development has been shaped by geography and institutional legacies – notably, the long-term impact of serfdom and regional frictions – resulting in persistent disparities in wealth, productivity, and economic structure (e.g. Markevich and Zhuravskaya, 2018; Zhuravskaya, Guriev, and Markevich, 2024; Markevich, 2019; Cheremukhin, Golosov, Guriev, and Tsyvinski, 2017). Within this context, the Mongol-origin *yams* system provides a rare quasi-experimental environment: it imposed a standardized spacing of relay stations across largely unsettled terrain, long before industrialization or large-scale colonization (Gurlyand, 1900). By linking pre-industrial logistical rules to modern economic outcomes, we offer one of the clearest empirical demonstrations that bureaucratic design – rather than natural advantage – can durably shape a country's urban geography.

Finally, a growing literature examines how postal systems shaped economic development by lowering the cost of communication. Hanlon, Heblich, Monte, and Schmitz (2022) show that Britain's 1840 Uniform Penny Post, which introduced a flat nationwide postage rate, sharply increased knowledge flows and patenting by reducing distance-dependent communication costs. Related studies find that the presence or reform of post offices often signaled stronger state capacity and spurred local innovation: nineteenth-century U.S. counties with denser postal coverage exhibited higher patenting rates (Acemoglu, Moscona, and Robinson, 2016), and bureaucratic modernization of the American Post Office further boosted inventive activity (Aneja and Xu, 2022). Historical reconstructions of continental postal routes likewise highlight the administrative and technological challenges of overland communication: in pre-industrial France, the state-sponsored post expanded slowly and often misaligned with economic geography, yielding little effect on city growth (Sasaki, 2025). To the best of our knowledge, we are the first to leverage one particular feature of postal systems – the regular distance to the next post station – for causal identification.

The remainder of the paper proceeds as follows. Section 2 provides historical context on the *yams* system and the evolution of the Russian postal network. Section 3 describes the digitization of historical sources, construction of the empirical grid, and data on modern outcomes and controls. Section 4 develops the identification strategy and presents the core results, while Section 5 rules out alternative explanations. Section 6 investigates why some post-station towns grew more than others, distinguishing the roles of first and second nature. Section 7 concludes and lists our next steps in this ongoing work.

2 Historical Background

This section provides the necessary historical background for the subsequent empirical analysis. The postal relay network of Imperial Russia provides the institutional and empirical foundation for our analysis. Its defining feature – the nearly uniform spacing of relay stations – creates the quasi-experimental variation central to our identification strategy. We trace the evolution of the *yam* system from its thirteenth-century Mongol origins to its codification under Muscovy and expansion across the empire, showing that station placement followed a simple logistical rule based on horse endurance rather than geography or trade. This mechanical design effectively scattered potential settlement sites across a sparsely populated landscape. The subsections below outline the system's institutional development, document its spatial layout, and illustrate how its rule-based spacing continues to align with modern economic activity.

2.1 Mongol Origins: The *Yam* System

The origins of Russia's postal relay system trace back to the administrative practices of the Mongol Empire. The Mongols introduced the *yam*, a system of organized horse relays and waystations, during their conquest of Rus' in the early 1240s (Belyakov, 2022; Sinan and Derya, 2023; Gurlyand, 1900). The system was primarily designed to ensure rapid communication and transportation of tribute, correspondence, and goods across vast territories. *Yams* were typically supported by local communities, which provided horses, fodder, and facilities for messengers; these contributions were originally a form of tribute (Gurlyand, 1900).

The Mongol *yam* model influenced the structure of settlements along key roads, particularly along routes connecting major centers to the steppe, and facilitated movement across sparsely populated areas (Belyakov, 2022). Although some scholars debate the direct continuity between the Mongol *yam* and the later Muscovite system, primary sources indicate that the Mongol model provided both a template and terminology adopted by Muscovy (Gurlyand, 1900; Sinan and Derya, 2023).

2.2 Adoption and Adaptation in Muscovy

In Muscovy, the *yam* evolved into a formal state-run relay system during the fifteenth century, with evidence suggesting that by Ivan III's reign (1462–1505), organized stations were already in operation along main roads (Semenov, 2016). These stations were strategically placed to allow messengers to travel efficiently, with an average spacing of 30–40 versts (32–43 km)[3], depending on terrain and road quality (Gurlyand, 1900). These intervals reflected the endurance limits of the typical Russian post horse rather than the distribution of population or resources.

As in the Mongol system, the fixed distance between stops was adjusted only minimally

[3]A verst equals 1.067 kilometers. Our empirical spacing in later sections is slightly smaller (around 30 kilometers) because we measure the *direct linear distance* between consecutive stations, whereas historical regulations referred to the actual travel distance along the winding postal road, which was typically 10–20% longer.

for local conditions: builders occasionally shifted a station a short distance – for instance, from a flood-prone riverbank to higher ground or to a site with available fodder – but the prescribed interval itself was rarely altered.[4] Whereas postal systems in Western Europe were superimposed upon dense settlement networks, Muscovy's *yam* routes often traversed vast stretches of uninhabited land. As a result, many post stations were established in entirely new locations, creating the first permanent settlements in their vicinity.

Each station, or *yam*, was managed by a small team of *yamschiki*, who oversaw horses, supplies, and lodging for government messengers. Early regulations required settlements to provide the necessary resources through *podvody*, a pre-existing system whereby locals furnished horses and food, which later coexisted alongside the formalized *yamskaya gon'ba* (Gurlyand, 1900; Belyakov, 2022). By the mid-sixteenth century, special documents called *podorozhnye* regulated the use of horses and the order of relay changes, and *zagonnye knigy* recorded distances and horse usage (Gurlyand, 1900).

From the mid-sixteenth century, a reform introduced the *yamskie okhotniki*, a dedicated class of postal workers responsible for managing the stations and resources, replacing the prior dependence on local population contributions (Gurlyand, 1900). This system gradually expanded to newly conquered territories, including Siberia at the end of the sixteenth century (Semenov, 2012, 2019).

2.3 Expansion and Bureaucratization under the Empire

By the seventeenth and early eighteenth centuries, the *yam* system had become a structured state institution. Oversight was centralized under the *Yamskoy Prikaz*, established in the early seventeenth century to manage station placement, standards for *podvody*, staffing of *yamskie okhotniki*, and accounting of taxes (Gurlyand, 1900). In Siberia and the Trans-Ural regions, river routes played a significant role, with stations providing resources for both land and water transport during seasonal navigation (Semenov, 2012; Vinogradov, 2022; Konoplev, 2021).

The system continued to evolve: settlements either provided *podvody* or, later, monetary payments (*mirskie otpuski*) to fund postal support (Gurlyand, 1900). By the late seven-

[4] Such micro-adjustments introduce minor endogeneity in precise locations along a route: local officials might place a station one or two kilometers before or after the nominal distance when encountering rivers, marshes, or unsuitable terrain. Because these deviations were idiosyncratic and symmetric around the rule-based interval, they form the basis for our instrumental-variable strategy later, where the *expected* spacing serves as an exogenous predictor of actual station placement. See the discussion in Section 4.

teenth century, the administration formalized wages for *yamskie okhotniki*, relieving local populations of direct provisioning duties (Gurlyand, 1900). The spatial logic of the system remained consistent, with stations spaced according to the endurance of horses rather than population density, and the network served as the backbone of government communication across both European and Asian parts of Russia.

By the mid-eighteenth century, the network linked provincial capitals, military forts, and frontier posts in a continuous grid spanning the empire's European and Asiatic territories. Two major sources document the system at its administrative height. The first is I. F. Truskott's *Postal Map of the Russian Empire* (1760), the earliest comprehensive cartographic representation of the network. The second is Vasilii Ruban's *Foreign and Russian Road Guide* (*Inostrannye i Rossiiskie putevye knigi*, 1777), which provides a complete enumeration of all stations then in operation, listing the sequential order of stops and the distances between them (Truskott, 1760; Ruban, 1777). Because these sources predate large-scale industrialization, they capture a pre-modern transport system whose spatial logic was dictated by horse-based communication rather than trade or population density.

2.4 Illustrative Example: The Moscow–St. Petersburg Route

To illustrate the structure of the network and the empirical logic of our analysis, Figure 1 combines four complementary views of the eighteenth-century postal route between Moscow and St. Petersburg, and anticipates our main finding.

Panel (a) shows the full route as documented in Ruban's *Foreign and Russian Road Guide* (1777), with post relays marked in red. Panel (b) zooms in on a short segment of this route, overlaying modern satellite night-light intensity on the historical rays connecting consecutive stations. Panels (c) and (d) summarize the quantitative regularities we exploit: the near-regular spacing of stations, and the corresponding peaks in modern light intensity at those intervals. Together, the panels show how a simple logistical rule – placing horse relays at nearly fixed distances – generated a durable spatial imprint still visible in the distribution of modern economic activity.

3 Data and Measurement

Our empirical analysis builds directly on the historical reconstruction of the Imperial Russian postal network presented in the previous section. This part of the paper details how we

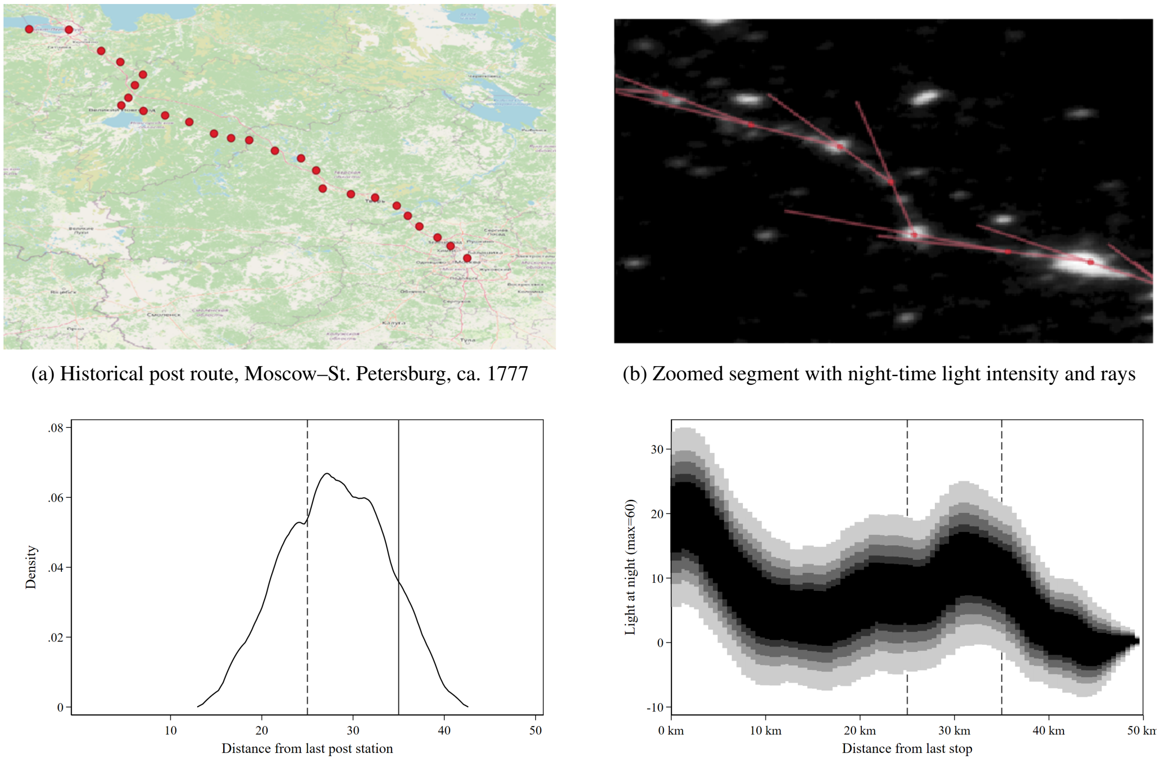

(a) Historical post route, Moscow–St. Petersburg, ca. 1777

(b) Zoomed segment with night-time light intensity and rays

(c) Distribution of distances between consecutive stations

(d) Modern light intensity along the same rays

Figure 1: Example: The Moscow–St. Petersburg postal route and its modern imprint

Notes: Panel (a) maps the eighteenth-century postal route from Moscow to St. Petersburg with relay stations marked in red. Panel (b) zooms in on a single segment, showing modern VIIRS night-light intensity (2020) along straight-line rays connecting stations. Panel (c) displays the histogram of distances between consecutive stations, with a pronounced mode at 25–35 km. Panel (d) plots modern brightness along the same route, peaking around the historical horse-changing interval. Together, these panels visualize the mechanical spacing of the Imperial postal grid and its enduring economic footprint.

translate that archival and cartographic evidence into a modern, spatially explicit dataset suited for econometric analysis. We begin by describing the historical sources from which the historical sources from which the names and locations of eighteenth-century post stations and routes were digitized and geographically identified. We then explain how these historical connections are transformed into a fine-grained empirical grid that defines our unit of observation – half-kilometer cells arrayed along straight-line rays between consecutive stations. Next, we discuss our measure of contemporary population density and economic activity, proxied by satellite night-time lights, and the geographic and infrastructural covariates used as controls. Finally, we document the core regularity of the postal system: the tight clustering of inter-station distances around the official rule of approximately thirty kilometers. This mechanical spacing, largely orthogonal to geography or prior settlement, provides the quasi-experimental variation exploited in our identification strategy.

3.1 Post Routes: Historical Sources and Digitization

Our empirical analysis begins by reconstructing the historical geography of Imperial Russia's postal relay network. The primary source is N. Ruban's *Foreign and Russian Road Guide* (1777), the first comprehensive registry of all postal stations (*yams*) together with the distances between consecutive stops. The register enumerates routes radiating from Moscow and St. Petersburg in all directions, listing each stop, its administrative jurisdiction, and the number of versts to the next station. We complement this catalogue with I. F. Truskott's *Postal Map of the Russian Empire* (1760), which provides a visual reference for route continuity and spatial coverage.

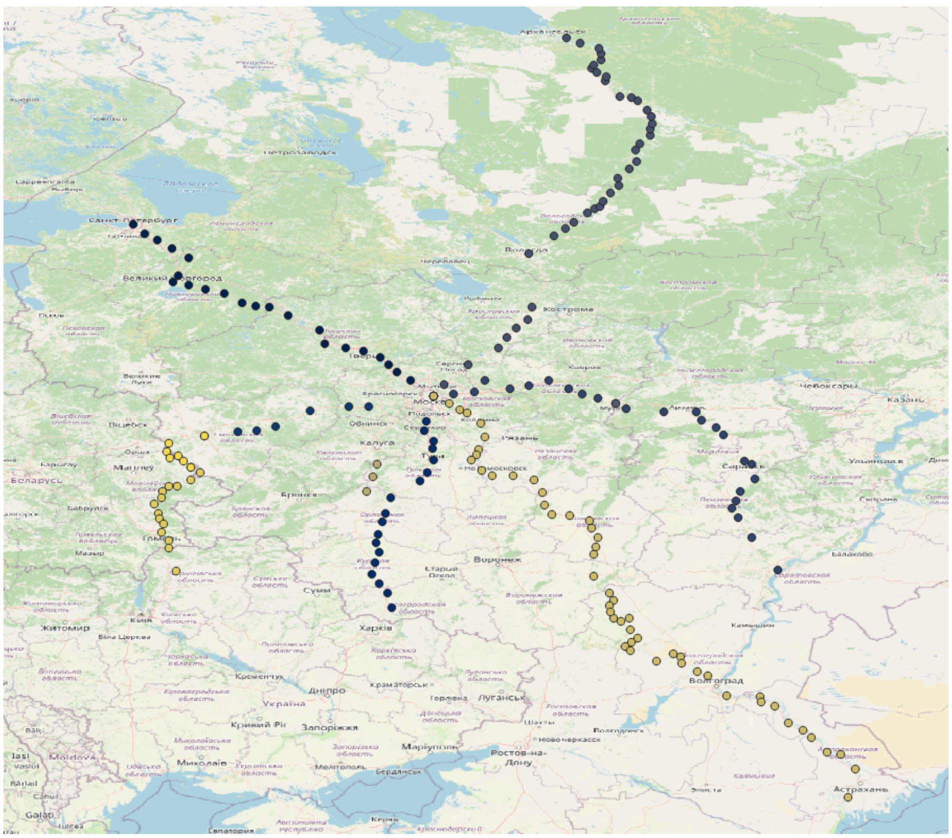

Figure 2: Sample of digitized post routes and stations from Ruban (1777)

Notes: The figure plots eight of the fiftenn digitized routes, namely those radiating from or around Moscow. Red points denote individual post stations. Routes and coordinates are reconstructed from Ruban's *Foreign and Russian Road Guide* (1777) and cross-checked against Truskott's *Postal Map of the Russian Empire* (1760). See Appendix B for a detailed list of routes and station counts.

Each station name was transcribed and transliterated into modern Russian orthography, matched to contemporary coordinates using the GeoNames and Wikimapia gazetteers, and cross-checked against historical atlases when it was impossible to identify the location otherwise. Our working sample consists of fifteen major routes extending from Moscow toward the north, west, south, and east of the empire, partially mapped in Figure 2.[5]

3.2 Constructing the Empirical Grid

To link historical routes to modern outcomes, we convert the network into a high-resolution spatial grid. Figure 3 shows the step-by-step construction of our units of observation. For each pair of consecutive stations, we draw a straight line connecting their coordinates – referred to as a *ray*. Each ray is divided into points spaced every 500 meters, around which we draw circles of 500 meter radius. The overlapping circles define our basic observational units, or *cells*. Each cell represents a segment of the historical corridor where a potential settlement could arise. This procedure yields roughly 40,000 cells across all rays. Each cell inherits its ray's historical context and modern geographic features, ensuring comparability across locations subject to the same administrative design.

(a) From station pairs to a straight-line ray

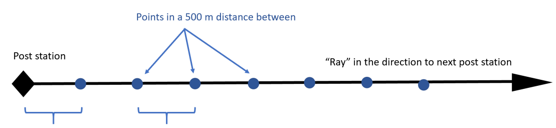

(b) Sampling points every 500 m along the ray

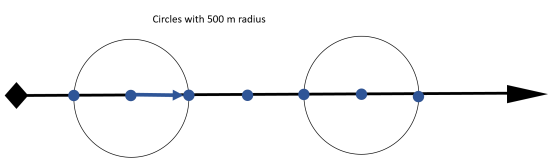

(c) 500 m circles around points ⇒ overlapping cells

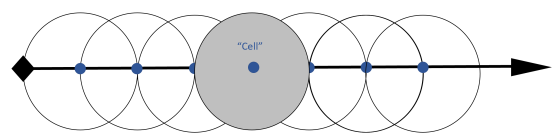

(d) Trim to 10–50 km and merge lights & controls

Figure 3: Construction of the Unit of Observation

Notes: Panel (a) illustrates the straight-line ray connecting each pair of consecutive post stations. Panel (b) places sampling points at 500 m intervals along the ray. Panel (c) draws 500 m-radius circles around each point, yielding overlapping cells that ensure continuous coverage of the historical corridor. Panel (d) shows the final analytic sample: we exclude the first 10 km after each station and truncate at 50 km to avoid overlap between adjacent segments, then merge each cell with modern night-time lights and geographic/infrastructure covariates.

3.3 Outcome Variable: Night-Time Light Intensity

Our measure of contemporary economic activity is night-time light intensity from the harmonized DMSP/VIIRS dataset of Li, Zhou, Zhao, and Zhao (2020), which provides consistent global luminosity estimates from 1992–2018. We use data for 2008, aggregated to 500 meter resolution. Night lights offer a finely disaggregated proxy for local income and population density, particularly valuable for territories where census or GDP data are un-

[5]This map focuses on the surroundings of Moscow. Several routes farther east are not displayed here. A route-by-route summary, including the number of stations digitized on each, is provided in Appendix B (Table A.2)

available or inconsistent. Each cell in the historical grid is matched to the luminosity value of its centroid pixel. Radiance values range from 0 (darkest) to 62 (brightest).

Following the approach in Henderson, Storeygard, and Weil (2012), Chen and Nordhaus (2011), and Donaldson and Storeygard (2016), satellite-based luminosity data provide a globally comparable and relatively noise-free measure of economic activity. Their primary limitations – saturation in bright urban cores and contamination from gas flares or fires – are minor in our setting, where most sites remain small towns or rural settlements.[6]

Modern luminosity visibly spikes near historical post-station sites, offering an intuitive first indication of their persistent influence on spatial development.

3.4 Control Variables and Additional Data

Each cell is matched with a comprehensive set of geographic and infrastructural controls. First-nature characteristics include elevation from the Shuttle Radar Topography Mission (SRTM, 90 m) and terrain ruggedness computed following Nunn and Puga (2012) using the SRTM surface (Jarvis, Reuter, Nelson, and Guevara, 2008). We measure caloric suitability using the global agro-climatic index of Galor and Özak (2016), and crop-specific suitability for potatoes, rye, and maize, amongst others, from the IIASA/FAO GAEZ dataset (Fischer, Nachtergaele, Prieler, van Velthuizen, Verelst, and Wiberg, 2012). Hydrological proximity is measured with HydroSHEDS river networks (Lehner, Verdin, and Jarvis, 2008). Pre-existing settlements are proxied by the presence of monasteries founded before 1800, following O'Neill, Bryant, Gupta, and Kiyan (2024). Second-nature features capture modern transport infrastructure: railways and highways from OpenStreetMap (2023 release), rasterized to 500 m resolution (OpenStreetMap contributors, 2023). All continuous variables are standardized within each route to ensure comparability. Because rays traverse diverse terrain, these controls absorb potential correlations between geography and station placement.

3.5 Sample Regularities

The mechanical design of the postal system is reflected in the distribution of inter-station distances. Across all routes, 39 percent of consecutive stations are spaced between 25 and

[6] A large validation literature confirms that night lights correlate strongly with local GDP, population, and household welfare across countries and time. See, among others, Mellander, Stolarick, Lobo, and Matheson (2015), Lessmann and Seidel (2017), and the surveys by Donaldson and Storeygard (2016) and Gibson, Olivia, and Boe-Gibson (2020). Elasticities of luminosity with respect to income or population are typically close to unity, and the signal performs well even in developing and low-light environments.

35 km apart, and nearly two-thirds fall within 20–40 km. This tight clustering corroborates historical regulations prescribing station intervals of 30–40 versts (32–43 km), and it forms the empirical basis for our identification strategy.[7] The distribution confirms that station placement reflected standardized logistical rules rather than variation in terrain or settlement density.

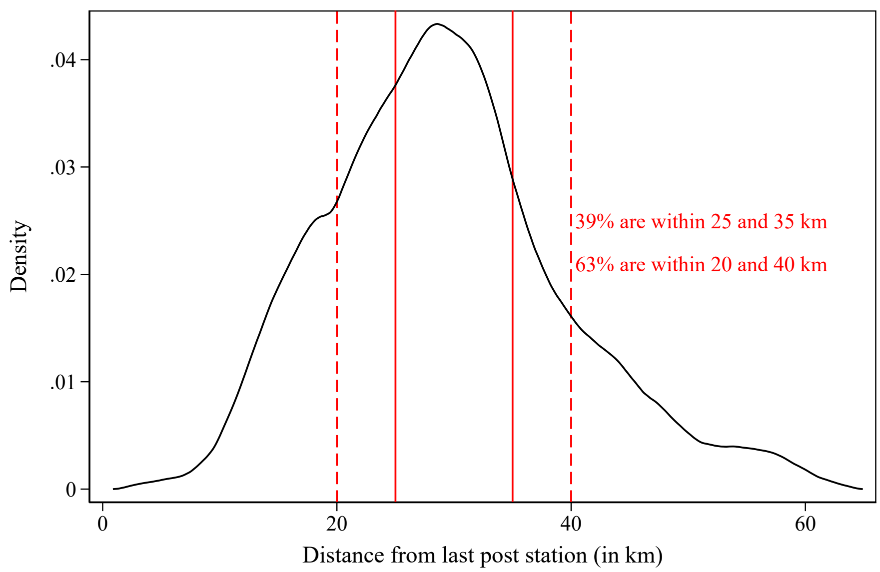

Figure 4: Distribution of Distances Between Consecutive Post Stations

Notes: Kernel density of inter-station distances (in kilometers) for all digitized routes from Ruban (1777). Nearly two-fifths of all intervals fall within the 25–35 km range corresponding to a day's horse relay. Distances are measured as straight-line (geodesic) separations between successive post stations.

Taken together, the dataset combines a rule-based pre-industrial network with high-resolution modern measures of economic activity and geography. The regularity of station spacing, combined with the limited role of natural features in route selection, provides the foundation for the quasi-experimental empirical design developed in the following section.

4 Empirical Strategy and Core Results

The preceding section established the empirical foundations of the analysis: a rule-based network of postal stations spaced at nearly fixed intervals across largely unsettled territory. We now turn from description to inference. This section lays out the empirical strategy used to quantify the long-run causal imprint of those mechanically placed stations on modern economic activity. We first formalize the identification framework, showing how the

[7] A verst equals 1.067 kilometers. The slightly smaller modal spacing in our data (around 30 km) arises because we measure the direct linear distance between stations rather than the winding postal road length.

network's fixed spacing provides quasi-random variation in settlement placement along historical transport corridors. We then present the main OLS and instrumental-variable estimates, interpret their magnitude and meaning, and assess robustness to alternative specifications. Together, these results demonstrate how an administrative rule for horse relays translated into a durable pattern of city location.

4.1 Identification Framework

Our goal is to estimate the causal effect of rule-based station placement on long-run population density. The empirical analysis exploits the mechanical spacing of postal relays to isolate quasi-random variation in settlement sites along historical postal routes. Because the network's reliance on horses dictated that a new post station was to be established at fixed intervals from the previous one, the distance alone provides an exogenous predictor of where a station *should* have been placed, irrespective of geography or pre-existing population.

To quantify this effect, we estimate variants of the following baseline specification at the cell level:

$$\text{Light}_c = \beta\,\mathbb{1}(\text{PostStation}_c) + \gamma X_c + \mu_r + \varepsilon_c, \tag{1}$$

where Light_c measures the average night-time light intensity in cell c, $\mathbb{1}(\text{PostStation}_c)$ is an indicator equal to one if the cell contains a historical postal relay, X_c denotes a vector of first- and second-nature controls described in Section 3, and μ_r are ray fixed effects absorbing all unobserved characteristics common to each station-to-station segment. Standard errors are clustered at the route level to account for spatial correlation in residuals along each route, and for the spatial autocorrelation introduced by constructing the overall cells.

The coefficient β captures the average difference in modern economic activity between locations that historically hosted a postal station and other cells along the same historical route, conditional on geography and infrastructure. Since the route alignment itself was non-random, ray fixed effects ensure identification arises solely from within-ray variation – comparing cells subject to the same logistical constraints but differing in whether they coincided with a horse-changing post.

4.2 Instrumental-Variables Design

Although historical accounts suggest that stations were located according to administrative distance rules rather than geography, there remains the possibility that officials occasionally adjusted locations to nearby settlements, river crossings, or terrain features. To purge such residual correlation, we instrument the presence of a station with a purely mechanical predictor: an indicator for whether a cell lies within the modal relay distance of 25–35 kilometers from the preceding post. The first-stage relationship can be written as

$$\mathbb{1}(\text{PostStation}_c) = \pi \, \mathbb{1}(\text{Range}_c) + \rho X_c + \mu_r + \upsilon_c, \tag{2}$$

where $\mathbb{1}(\text{Range}_c)$ equals one if cell c lies within the expected spacing interval. Because this range was determined by horse endurance rather than local conditions, it provides an exogenous source of variation in the probability of station placement along each ray. The resulting two-stage least squares estimate identifies the local average treatment effect (LATE) of a postal station's presence for cells whose assignment to treatment is induced by the mechanical spacing rule.

The exclusion restriction requires that the 25–35 km range affects modern brightness only through its impact on the likelihood of hosting a historical relay. This assumption is plausible because the empirical variation exploited in our design arises only within rays, not across them. Route fixed effects absorb any differences in geography, regional infrastructure, market access, or administrative priority, leaving the 25–35 km spacing rule as the only systematic determinant of where a station could have been placed. Moreover, the rule-based distance was not chosen to optimize transport costs or settlement spacing but to minimize horse fatigue; any residual correlation between the mechanical interval and modern outcomes would therefore have to operate through the establishment of the station itself.

4.3 Main Results

To visualize the core relationship, Figure 5 plots average night-time light intensity along historical postal routes as a function of distance from the preceding station. Brightness rises sharply around the 25–35 km range – the interval corresponding to the prescribed horse-relay distance – and declines symmetrically before and after. This pattern provides a visual counterpart to our identification strategy: modern economic activity peaks precisely where the historical spacing rule dictated a stop, consistent with the view that administrative placement seeded persistent settlement.

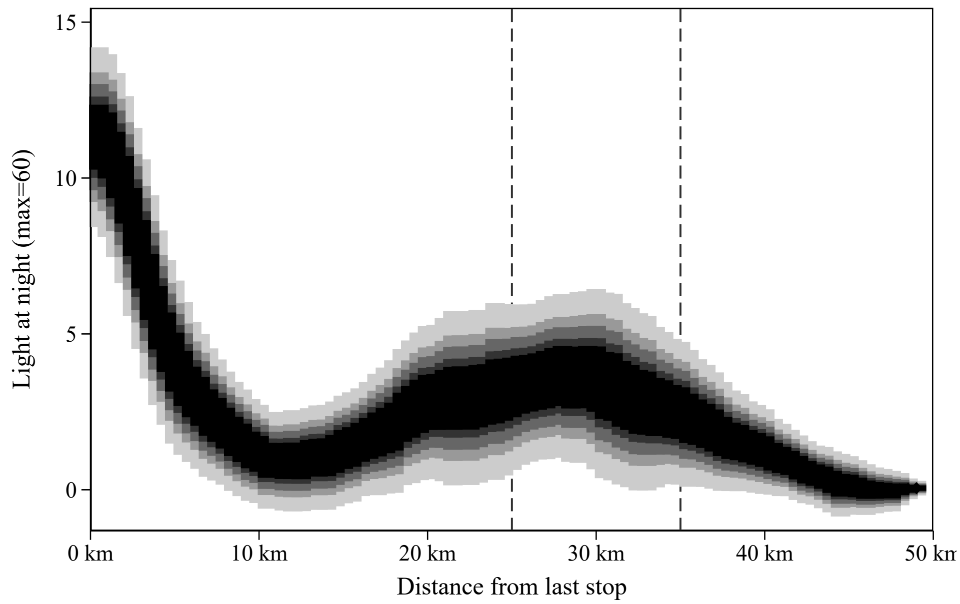

Figure 5: Modern Night-Time Light Intensity by Distance from Previous Post Station

Notes: Each point shows average VIIRS night-time light intensity in 500 m cells, binned by distance from the preceding station, after removing ray fixed effects. The shaded area marks the modal horse-relay interval (25–35 km). Night lights peak precisely at this distance, consistent with persistent development at historical post sites.

Before turning to formal estimates, Figure 6 demonstrates that no comparable discontinuities appear in observable geographic features. Elevation, caloric soil suitability, river proximity, and suitability for maize remain essentially flat across the 25–35 km range, suggesting that the brightness spike is not driven by correlated first- or second-nature factors.[8]

Table 1 reports the corresponding OLS, reduced-form, and instrumental-variable estimates of Equation (1). All specifications include ray fixed effects, and successively add first- and second-nature controls.

In the simplest OLS specification in Column 1 of Panel A, without additional controls and without ray fixed effects , the coefficient on $\mathbb{1}(\text{PostStation}_c)$ is positive and highly significant. Cells coinciding with former postal sites are today brighter by one sample mean (8.2) compared to those before or after the modal range. Adding controls for topogra-

[8]The same broad pattern holds for any other feature of local natural geography we located. Results available upon request.

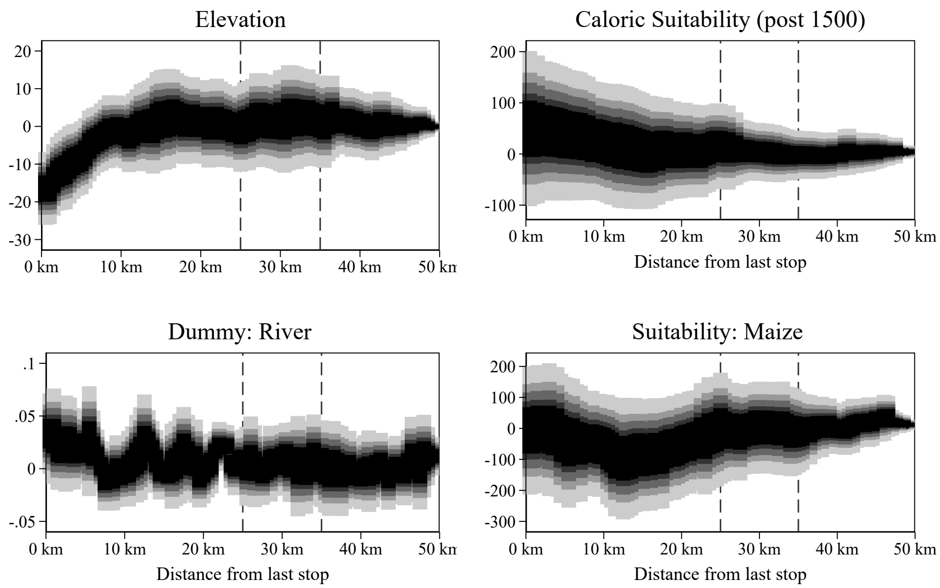

Figure 6: Balancing Tests: Natural Geography by Distance from Previous Station

Notes: Each line plots the average value of the indicated control (elevation, river proximity, soil suitability) by distance from the preceding station, normalized within ray. None of these variables exhibits a discontinuity at the 25–35 km interval, confirming that the brightness peak is not driven by local geography or infrastructure.

phy, soil, rivers, and modern infrastructure slightly attenuates the coefficient but leaves its magnitude and significance essentially unchanged. The similarity of the estimates with and without controls indicates that neither geography nor later transport investments can explain the clustering of modern activity around historical station sites.

The reduced-form regression of Panel B, which replaces the actual station indicator with the instrumental variable $\mathbb{1}(\text{Range}_c)$, confirms that brightness peaks precisely at the historical horse-changing interval. The coefficient in the most demanding specification in Column 4 of Table 1, 1.86, corresponds to roughly 23% of the sample mean. The log specification in Table A.3 suggests that this amounts to approximately a 30% increase compared to the regions before and after the modal range.

The corresponding IV estimate, depicted in panel C, is an order of magnitude larger than the OLS coefficient, consistent with attenuation bias in the latter due to measurement error in identifying exact post-station coordinates. Because the IV captures the effect for locations whose treatment status was determined by the spacing rule, it reflects a local average treatment effect (LATE) for marginal station sites. The first stage is strong: cells within 25–35 km are significantly more likely to contain a station, with an F-statistic exceeding conventional thresholds.[9]

Table 1 also demonstrates that these results remain stable across specifications. In the most demanding regression including ray fixed effects and all geographic and infrastructural covariates, the IV coefficient remains statistically significant at the one percent level.

Table 1: Effect of Historical Postal Stations on Modern Economic Activity

	(1)	(2)	(3)	(4)
Panel A: OLS				
Post Stop	8.876***	8.404***	8.071***	6.972***
	(0.704)	(0.652)	(0.499)	(0.388)
R^2	0.01	0.59	0.62	0.64
Observations	33,259	33,259	30,583	30,583
Panel B: Reduced Form				
$\mathbb{1}(25 <$ Dist. $< 35)$	1.801***	1.801***	2.001***	1.861***
	(0.524)	(0.524)	(0.562)	(0.545)
Panel C: Instrumental Variables				
Post Stop	105.399***	105.399***	101.515***	99.604***
	(25.650)	(26.551)	(24.071)	(24.880)
First-Stage F-Stat	19.53	19.53	25.13	21.76

Notes: Dependent variable is VIIRS night-time light intensity in 500 m cells along rays from each historical post station. Each column after column (1) includes ray fixed effects; columns (3)–(4) add first- and second-nature controls, respectively. The instrument for the presence of a station is an indicator for cells lying within 25–35 km of the preceding stop. Robust standard errors clustered at the route level. All regressions exclude the first 10 km after each station and truncate at 50 km.

4.4 Interpretation and Mechanisms

The pattern shown in Figure 5 and quantified in Table 1 implies a lasting causal imprint of administrative design. Locations that hosted an eighteenth-century postal relay – sites originally chosen for purely logistical reasons – remain significantly more developed today than neighboring points along the same route. The magnitude of the effect is economically meaningful: an increase of roughly one-third in night-light intensity corresponds to substantially higher local population density and income per unit area.[10]

[9] Anderson Rubin intervals (not reported) do not include zero, suggesting that our first stage is sufficiently strong (Lee, McCrary, Moreira, and Porter, 2022).

[10] Using the elasticity of GDP with respect to night lights of around 1 from Henderson et al. (2012), a 30% brightness differential translates into roughly 30% higher local economic output. Comparable elasticities for population density are reported in Donaldson and Storeygard (2016).

Because the average brightness of cells along the routes is low, these effects correspond not to large cities but to the emergence of small towns and peri-urban centers around former post sites.

The larger IV coefficients relative to the OLS estimates are consistent with both measurement error and local heterogeneity. Geolocation uncertainty in some station coordinates mechanically attenuates OLS estimates toward zero, while the IV design identifies the *local average treatment effect* for those settlements whose establishment was determined by the mechanical spacing rule. In this sense, the IV estimate isolates the causal effect of rule-based placement on long-run development, independent of pre-existing settlement or geography.

What mechanisms might explain the persistence of this administrative imprint? At the time of their founding, post stations created small but durable concentrations of economic activity: stable travel demand, employment for the *yamschiki* and *okhotniki*, fodder markets, and infrastructure such as stables, wells, and rest houses. Once established, these outposts lowered fixed costs of settlement and became focal points for subsequent trade, craft production, and agricultural exchange. Over time, local agglomeration forces – knowledge spillovers, transport complementarities, and scale economies in service provision – reinforced their viability even after the postal system's logistical purpose faded. The result is a pattern of enduring small-scale urbanization whose roots lie in bureaucratic design rather than in natural geography.

Section 6 examines these mechanisms more systematically. We test whether the enduring success of former post stations can be attributed to favorable geography or later infrastructure, or whether the administrative "city seeds" grew largely independently of both.

4.5 Summary of Robustness

The baseline results are stable across a battery of checks detailed in Appendix C. First, alternative outcome definitions (binary and log lights) and coarser clustering schemes deliver comparable estimates in both OLS and IV, indicating that neither functional form nor inference hinges on modeling choices (Table A.3). Second, varying the mechanical relay interval used in the instrument (20–40 km; 15–45 km) and expanding or shrinking the along-ray analysis window (0–50 km; 0–80 km; baseline 10–50 km) shows large and precisely estimated coefficients throughout (Table A.4). Third, route-by-route estimates show largely positive effects, confirming that no single corridor drives the results (Figure A.1).

Finally, in Appendix D, we document that the results replicate at higher levels of spatial aggregation. Specifically, we create three segments of each ray (one at the "instrument" range of 25-35 km, and one each before and after that), and show that light spikes at the instrument range, and indeed operates over it.

Together, these exercises support the conclusion that the brightness peak at the historical horse-relay interval reflects a pervasive feature of the postal network rather than a modeling artifact.

5 Ruling out Alternative Explanations

The previous section established a strong and robust relationship between the placement of eighteenth-century postal stations and modern economic activity. While the rule-based nature of the network provides a credible source of quasi-random variation, it remains necessary to verify that this persistence does not reflect pre-existing settlement patterns, agglomeration around large cities, or generic regularities in the spatial distribution of towns. This section addresses each of these potential alternative explanations in turn.

5.1 Pre-Existing Settlements

A first concern is that postal stations may not have been placed in uninhabited areas but instead located at or near existing population centers. If so, the observed brightness peaks could simply reflect long-standing settlement continuity rather than the causal effect of administrative station placement. Although large parts of eighteenth-century Russia were sparsely populated, some pre-Muscovite religious foundations and trading posts existed along major routes. To address this possibility, we match the historical postal network to an independent dataset of monasteries founded before 1800 (O'Neill et al., 2024). These monasteries represent the earliest durable settlements for which precise geocodes are available across the empire. We then re-estimate the baseline and IV specifications including a dummy variable indicating whether a cell contains, or lies adjacent to, a pre-1800 monastery. As shown in Figure 7, which shows the baseline brightness profile from Figure 5 controlling for local features (Panel a) and the version further controlling for monastery presence (Panel b), the inclusion of this control does not alter the main pattern.

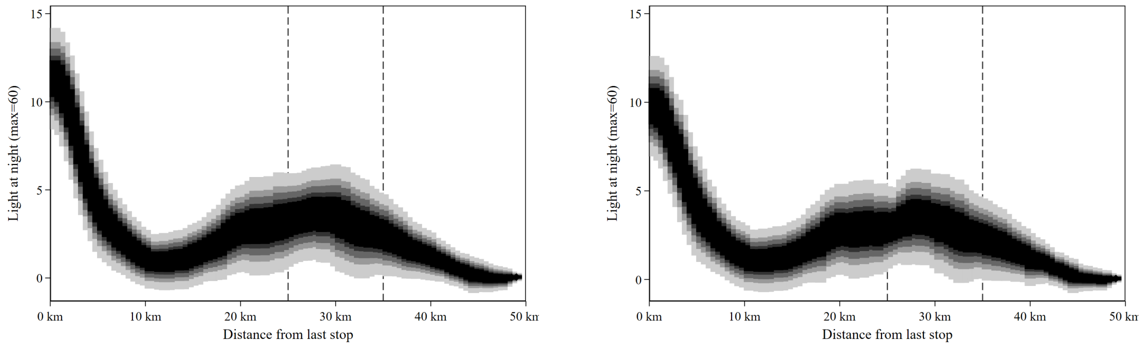

(a) Controlling for geography and infrastructure

(b) Further controlling for pre-1800 monastery presence

Figure 7: Controlling for Pre-Existing Settlements

Notes: Each panel plots average night-time light intensity along historical rays, expressed as distance (km) from the preceding post station. Panel (a) includes all first- and second-nature controls; panel (b) further controls for the presence of any monastery founded before 1800. The brightness peak around 30 km persists in both cases.

These results suggest that the emergence of towns at post-station sites cannot be attributed to the re-use of pre-existing settlements but instead reflects new urban growth induced by the postal infrastructure itself.

5.2 Proximity to Major Urban Centers

A second possibility is that the results merely capture spillovers from major cities such as Moscow or St. Petersburg. To evaluate this concern, we sequentially exclude (i) all routes that terminate in Moscow or St. Petersburg, and (ii) any station that later evolved into a large urban center. The estimated effects remain very comparable to the baseline, both in magnitude and statistical precision (Table 2).

The robustness of the results to these exclusions indicates that the persistence we document is not confined to metropolitan corridors but extends to peripheral and regional routes. Even in remote areas, brightness today peaks precisely at the historical horse-changing interval, confirming that the spatial pattern is a general property of the postal network rather than a by-product of urban proximity.

5.3 Generic Spacing Patterns or "Systems of Cities"

A third explanation could be that settlements naturally form at regular intervals, independent of any postal or administrative design. Classical theories of urban hierarchies and central-place systems posit equilibrium distances between towns that depend on transport costs and communication range (Henderson, 1974; Fujita et al., 1999). If so, the observed 30-kilometer spacing could reflect a general "law of urban spacing" rather than the specific institutional legacy of the horse-relay network.

Table 2: Robustness: Excluding Routes to Major Cities

	(1) Baseline	(2) Drop Moscow Routes	(3) Drop St. Petersburg Routes	(4) Drop Top-5% Largest Stations
Panel A. OLS				
Post Stop	7.082***	6.874***	8.879***	19.681***
	(0.627)	(0.846)	(2.255)	(2.728)
R^2	0.63	0.66	0.75	0.71
Observations	30,583	15,146	1,952	22,546
Panel B. Reduced Form				
$\mathbb{1}(25 <$ Dist. $< 35)$	1.859***	3.302***	5.128**	1.214***
	(0.415)	(0.641)	(1.935)	(0.444)
Panel C. IV				
Post Stop	99.080***	133.452***	210.284**	357.914**
	(24.493)	(30.163)	(96.692)	(147.709)
First-Stage F-Stat	21.15	33.28	3.97	6.20

Notes: Dependent variable is VIIRS night-time light intensity in 500 m cells along rays from each historical post station. Columns exclude, respectively, (2) all routes leading to Moscow, (3) those leading to St. Petersburg, and (4) the 5% of stations that later became large cities. All regressions include ray fixed effects and first- and second-nature controls. Standard errors clustered at the ray level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

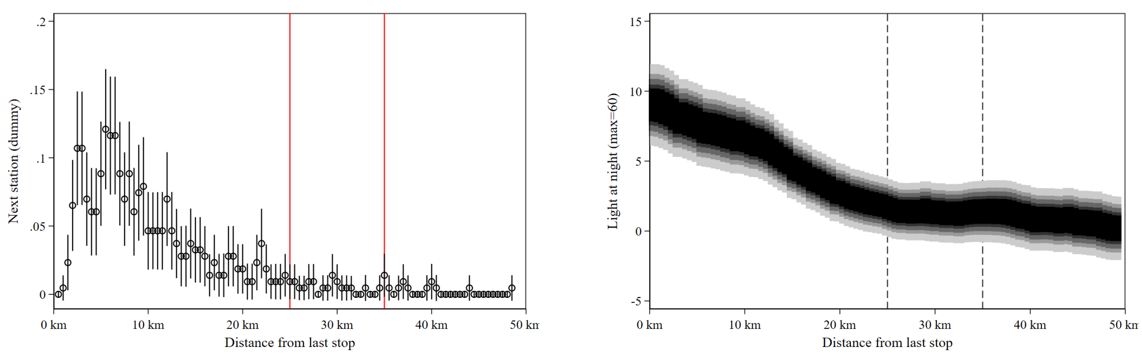

(a) Inter-station distances, Trans-Siberian Railroad

(b) Modern night lights along railroad routes

Figure 8: Placebo: The Trans-Siberian Railroad

Notes: The Trans-Siberian Railroad (completed 1901) provides a placebo test for generic spacing patterns. Panel (a) shows that railroad stops are spaced roughly every 10 km, corresponding to steam-locomotive water stops; panel (b) reveals no brightness spike at the 25–35 km interval characteristic of horse relays.

To test this hypothesis, we replicate our analysis for a distinct transport technology constructed under very different engineering constraints: the Trans-Siberian Railroad. Completed at the turn of the twentieth century, the Trans-Siberian line spanned more than 9,000 km (Liliopoulou, Roe, and Pasukeviciute, 2005), and engineering-history sources suggest station locations were heavily influenced by water-supply logistics, including pump houses and pipelines (Shilnikova, 2022). If our results merely captured a universal spacing regularity, similar brightness peaks should appear at characteristic intervals along the railway network.

Figure 8 plots the distribution of distances between consecutive railroad stops (panel a) and the corresponding modern light intensity along those routes (panel b). In contrast to the postal system, brightness peaks at before 10 kilometers – the typical distance between water stops for steam locomotives – and not at 30 kilometers. No secondary peak appears near the postal-relay interval. This placebo confirms that the spatial imprint we document is specific to the rule-based horse-relay system and not a general property of settlement formation or transport spacing.

5.4 Summary

Across all tests, the evidence consistently rejects alternative explanations based on pre-existing settlements, proximity to major cities, or universal settlement spacing. The persistence of brightness peaks at the horse-changing interval even when all of these factors are controlled for indicates that the causal channel operates through the original administrative design of the postal system. The following section investigates why some of these historically seeded locations developed into thriving towns while others remained small.

6 Why Do Some City Seeds Grow?

The results so far demonstrate that the initial placement of post stations acted as quasi-random "city seeds" whose locations have left a lasting imprint on Russia's urban geography. Yet not all seeds developed equally. Many post settlements remain small villages, while others have grown into towns and regional centers. This section investigates what explains this heterogeneity: did natural endowments or later infrastructure investments influence which stations prospered? Following the literature, we refer to these determinants as *first nature* and *second nature*, respectively.

6.1 Empirical Approach

To analyze the correlates of post-station growth in a coarser but more interpretable way, we shift the level of analysis to the level of city seeds, i.e, those intervals between 25 and 35 km distance from the last post relay station. To this end, we aggregate all variables for this range for each of the rays, and compare these intervals between different rays.

We estimate cross-sectional regressions of the form:

$$\text{Light}_r = \alpha + \gamma Z_r + \varepsilon_r, \tag{3}$$

where Light_r measures the average brightness in the 25-35 km range of ray r, and Z_r

denotes averaged local features such as elevation, caloric suitability, or proximity to modern infrastructure. Throughout, we restrict the sample to those cases in which a post station was indeed located in that range.

6.2 Results

Table 3: Which Local Features Further Urban Growth?

	(1)	(2)	(3)	(4)
Elevation (std.)	1.519		2.086	
	(2.482)		(2.067)	
Caloric suitability (std.)	2.895		2.986**	
	(1.793)		(1.493)	
Barley suit. (std.)	-10.172***		-7.965**	
	(3.699)		(3.788)	
Cabbage suit. (std.)	1.442		-0.719	
	(2.695)		(2.277)	
Oats suit. (std.)	1.577		2.901	
	(7.341)		(6.029)	
Rye suit. (std.)	4.027		0.588	
	(6.437)		(5.031)	
Wheat suit. (std.)	-0.865		0.662	
	(4.779)		(4.292)	
Potato suit. (std.)	4.651		3.982*	
	(2.852)		(2.365)	
River		4.860	4.769	
		(3.149)	(3.048)	
Road		5.585***	5.317***	
		(1.544)	(1.466)	
Railroad		10.615***	9.024***	
		(2.293)	(2.246)	
Highway		0.367	2.573	
		(9.401)	(10.015)	
R^2	0.16	0.33	0.41	
Observations	155	159	155	

Notes: The table presents results from cross-sectional regressions at the ray level. The sample is segments in the 25–35 km range that also hosted a post station. The dependent variable is the mean of the median night light within each segment, and the independent variables are standardized averages of first nature features and dummies for second nature features. Robust standard errors in parentheses.

Table 3 highlights two main findings. First, *first-nature* factors play a modest role. Elevation and overall caloric suitability are positively correlated with modern brightness, suggesting that more favorable terrain and general agricultural potential offered some long-run advantage. However, since the mean of the dependent variable for this sample is approximately 14.7, and the standard deviation is approximately 17.9, these associations are not substantial: a one-standard-deviation increase in caloric suitability increases light at night by around 0.16 standard deviations. The suitability for specific crops such as barley, rye, oats, or wheat shows no consistent pattern and is often insignificant or even negative. Hence, broad agricultural potential, rather than specialization in any single crop, appears to have mattered only weakly. Crucially, potato suitability is positively associated with higher luminosity. This is telling, since potatoes were adopted in Russia only after post station locations were chosen, and bureaucrats could hardly have selected places based on their suitability for crops that were not yet cultivated. Yet, rays on which the natural stopping rule happened to occur on land particularly suited to potato cultivation tended to grow much larger over time.

Second, *second-nature* characteristics such as rivers, roads, or railways make a real difference for current luminosity and therefore, settlement sizes. Since these variables are dummy, it is straightforward to gauge their relative importance. Having a (rail)road cross anywhere through the relevant segment increases light at night there by 0.3 (0.5) standard deviations.

Overall, these results suggest that first nature can only explain the differential success of seeds to a limited extent. The persistence of the postal network's imprint thus reflects an administrative rather than geographic origin of urban seeds, which was later reinforced by subsequent, second-nature transportation investments.

7 Conclusion

Why do cities emerge where they do? This paper revisits this foundational question by exploiting a unique natural experiment from the pre-industrial era. We show that an administrative transport network – originally designed to optimize horse relays across eighteenth-century Russia – left a lasting imprint on the spatial distribution of modern economic activity. By combining the mechanical regularity of historical postal routes with high-resolution satellite data, we provide rare quasi-experimental evidence on the origins and persistence of urban settlement.

Because post stations were placed at nearly fixed intervals of roughly 25–35 kilometers, independent of geography or population, their locations can be treated as exogenous "city seeds" scattered across an otherwise sparsely inhabited landscape. Two and a half centuries later, these sites remain systematically brighter in night-light data, indicating higher population density and income. The persistence of this pattern cannot be explained by geography, pre-existing settlements, or proximity to major cities, nor does it arise from generic spacing regularities in urban systems. Even after controlling for first-nature endowments and second-nature infrastructure, the causal effect of historical station placement remains large and statistically robust. Aggregated analyses further reveal that natural advantages – except for potatoes and overall caloric suitability – played little role in determining which station towns prospered, and that later transport infrastructure reinforced the underlying pattern already established before.

In sum, the study illustrates how an apparently minor logistical rule – how far a horse could travel before rest – shaped the long-run geography of one of the world's largest countries. Administrative decisions made for reasons of communication and control inadvertently determined where people would live, trade, and build for centuries to come.

Directions for Future Work. While this version of the draft has focused on eighteenth-century Russia as a uniquely well-suited setting to study rule-based infrastructure placement, this preliminary version of the draft will be extended and refined in several directions. The next iteration of this project will expand both the geographic scope and methodological depth of the analysis. First, we plan to explore *heterogeneity* across space, examining whether persistence differs between European and Asian Russia or between postal routes established under distinct administrative regimes. Second, we will compare the industrial and demographic composition of post-station towns to nearby control towns, providing new evidence on how initial administrative functions translated into specific economic trajectories. Third, we aim to broaden the scope to other historical transport systems – most notably the nineteenth-century U.S. railroads, where changes in locomotive technology altered the distance between mandatory water-refueling stops. This shift generated quasi-experimental variation over time in station spacing, offering a natural *difference-in-differences* setting to test how evolving transport technologies shape the spatial persistence of development. Finally, we will refine our empirical framework using least-cost-path algorithms, circular buffers, and cones instead of straight lines, and will use lines in other directions from each post station as placebos. Finally, we will try to find intermediate population estimates for the post relays, to further trace how and why some grew into cities and others not.

References

Acemoglu, D., J. Moscona, and J. A. Robinson (2016). State Capacity and American Technology: Evidence from the Nineteenth Century. *American Economic Review: Papers and Proceedings 106*(5), 61–67.

Ades, A. F. and E. L. Glaeser (1995). Trade and Circuses: Explaining Urban Giants. *Quarterly Journal of Economics 110*(1), 195–227.

Allen, T. and D. Donaldson (2020). Persistence and path dependence in the spatial economy. Technical report, National Bureau of Economic Research, Inc.

Aneja, A. and G. Xu (2022). Strengthening State Capacity: Postal Reform and Innovation during the Gilded Age. NBER Working Paper 29852.

Belyakov, A. (2022). The Organization of Embassy Service in the Russian Lands in the Thirteenth through Sixteenth Centuries. *Herald of the Russian Academy of Sciences 92*(Suppl 5), S419–S427.

Bleakley, H. and J. Lin (2012). Portage and Path Dependence. *Quarterly Journal of Economics 127*(2), 587–644.

Bosker, M. and E. Buringh (2017). City Seeds: Geography and the Origins of the European City-System. *Journal of Urban Economics 98*, 139–157.

Chen, X. and W. D. Nordhaus (2011). Using Luminosity Data as a Proxy for Economic Statistics. *Proceedings of the National Academy of Sciences 108*(21), 8589–8594.

Cheremukhin, A., M. Golosov, S. Guriev, and A. Tsyvinski (2017). The Industrialization and Economic Development of Russia through the Lens of a Neoclassical Growth Model. *Review of Economic Studies 84*(2), 613–649.

Dalgaard, C.-J., N. Kaarsen, O. Olsson, and P. Selaya (2022). Roman Roads to Prosperity: Persistence and Non-Persistence of Public Infrastructure. *Journal of Comparative Economics 50*(4), 896–916.

Davis, D. R. and D. E. Weinstein (2002). Bones, Bombs, and Break Points: The Geography of Economic Activity. *American Economic Review 92*(5).

Donaldson, D. and R. Hornbeck (2016). Railroads and American Economic Growth: A 'Market Access' Approach. *Quarterly Journal of Economics 131*(2), 799–858.

Donaldson, D. and A. Storeygard (2016). The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives 30*(4), 171–198.

Düben, C. and M. Krause (2023). The Emperor's Geography: City Locations, Nature and Institutional Optimisation. *The Economic Journal 133*(651), 1067–1105.

Fischer, G., F. Nachtergaele, S. Prieler, H. van Velthuizen, L. Verelst, and D. Wiberg (2012). Global Agro-ecological Zones (GAEZ v3.0): Model Documentation. IIASA and FAO. Data and documentation at https://gaez.fao.org/.

Frick, S. A. (2018). Urbanization, cities and economic growth: trends, recent evidence and policy implications. *Regions*. Online journal; accessed 2025-11-10; DOI:10.1080/13673882.2018.00001006.

Fujita, M., P. Krugman, and A. J. Venables (1999). *The Spatial Economy: Cities, Regions, and International Trade*. Cambridge, MA: MIT Press.

Galor, O. and Ö. Özak (2016). The Agricultural Origins of Time Preference. *American Economic Review 106*(10), 3064–3103.

Gibson, J., S. Olivia, and G. Boe-Gibson (2020). Night lights in economics: Sources and uses. *Journal of Economic Surveys 34*(5), 955–980.

Glaeser, E. L. and G. A. Ponzetto (2018). The political economy of transportation investment. *Economics of Transportation 13*, 4–26.

Gurlyand, I. Y. (1900). *Yamskaya gon'ba v Moskovskom gosudarstve do kontsa XVIII veka*. Moscow. Vols. I–II.

Hanlon, W. W., S. Heblich, F. Monte, and M. B. Schmitz (2022). A Penny for your Thoughts. Working paper; version dated May 30, 2022.

Heblich, S., S. J. Redding, and D. M. Sturm (2020). The Making of the Modern Metropolis: Evidence from London. *Quarterly Journal of Economics 135*(4), 2059–2133.

Henderson, J. V. (1974). The Sizes and Types of Cities. *American Economic Review 64*(4), 640–656.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012). Measuring Economic Growth from Outer Space. *American Economic Review 102*(2), 994–1028.

Henderson, J. V. and H. Wang (2007). Urbanization and City Growth: The Role of Institutions. *Regional Science and Urban Economics 37*(3), 283–313.

Jarvis, A., H. I. Reuter, A. Nelson, and E. Guevara (2008). Hole-filled SRTM for the Globe Version 4. CGIAR-CSI SRTM 90m Database. Available at http://srtm.csi.cgiar.org/.

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

Konoplev, A. (2021). Routes of Yam Relay in the European North of Russia in the 16th–17th Centuries. *9*, 178–187.

Lee, D. S., J. McCrary, M. J. Moreira, and J. Porter (2022). Valid t-Ratio Inference for IV. *American Economic Review 112*(5), 1601–1637.

Lehner, B., K. Verdin, and A. Jarvis (2008). New Global Hydrography Derived from Spaceborne Elevation Data. *Hydrological Processes 22*(19), 3863–3874.

Lessmann, C. and A. Seidel (2017). Regional inequality, convergence, and its determinants – A view from outer space. *European Economic Review 92*, 110–132.

Li, X., Y. Zhou, M. Zhao, and X. Zhao (2020). Harmonization of DMSP and VIIRS Nighttime Light Data from 1992–2018 at the Global Scale. *Scientific Data 7*(168).

Liliopoulou, A., M. Roe, and I. Pasukeviciute (2005). Trans siberian railway: From inception to transition. *European Transport / Trasporti Europei X*(29), 46–56. Open Access.

Markevich, A. (2019). A Regional Perspective on the Economic Development of the Late Russian Empire. SSRN Working Paper, last revised May 14, 2019.

Markevich, A. and E. Zhuravskaya (2018). The Economic Effects of the Abolition of Serfdom: Evidence from the Russian Empire. *American Economic Review 108*(4-5), 1074–1117.

Mellander, C., K. Stolarick, J. Lobo, and Z. Matheson (2015). Night-Time Light Data: A Good Proxy Measure for Economic Activity? *PLOS ONE 10*(10), e0139779.

Michaels, G. and F. Rauch (2018). Resetting the urban network: 117–2012. *The Economic Journal 128*(608), 378–412.

Nagy, D. K. (2022). Trade and urbanization: Evidence from hungary. *American Economic Journal: Microeconomics 14*(3), 733–790.

Nguyen, D. A., S. Brakman, H. Garretsen, and T. Kohl (2023). What's in a Name? Initial Geography and German Urban Development. Technical Report 10435, CESifo Working Paper Series.

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

Ochsner, C. (2025). Hostility, Population Sorting, and Backwardness: Quasi-Experimental Evidence from the Red Army after WWII. CERGE-EI Working Paper.

O'Neill, K., J. Bryant, S. Gupta, and O. Kiyan (2024). Relics of Empire: The Geography of Holy Places, Volume 1 – Orthodox Monasteries. UNF:6:B2SSwbSLftYpbnhPXQDNZw==.

OpenStreetMap contributors (2023). OpenStreetMap railway & highway data. https://www.openstreetmap.org/copyright. Data licensed under the Open Database License (ODbL).

Oxford Economics (2024). Global Cities Index 2024. Technical report, Oxford Economics. Accessed: 2025-11-10.

Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative Spatial Economics. *Annual Review of*

Economics 9, 21–58.

Ritchie, Samborska, and Roser (2024). Urbanization. https://ourworldindata.org/urbanization. Accessed: 2025-11-10.

Ruban, N. (1777). Inostrannyy i Rossiyskiy dorozhnyy putevoditel' [Foreign and Russian Road Guide]. First comprehensive register of postal stations in the Russian Empire.

Sasaki, Y. (2025). Postal Growth: How the State-Sponsored Post Affected Growth in Preindustrial France, 1500–1850. Discussion paper f-1253, CIRJE, University of Tokyo.

Schweiger, H., A. Stepanov, and P. Zacchia (2022). The Long-run Effects of R&D Place-based Policies: Evidence from Russian Science Cities. *American Economic Journal: Economic Policy 14*(3), 322–351.

Semenov, O. V. (2012). The Yam System and State "Court Affairs" in Western Siberia in the 17th Century. *Russian History* (3), 138–143.

Semenov, O. V. (2016). On the Time and Circumstances of the Emergence of the Yam System in the Russian Lands. *Izvestiya Ural Federal University. Series 2: Humanities 18*(1(148)), 103–119.

Semenov, O. V. (2019). Yamskaya Gon'ba in Siberia until the End of the 17th Century: Problem of Emergence and Functioning. *Cahiers du Monde Russe 60*(1), 147–174.

Shilnikova, I. V. (2022). Organization of water supply on the trans-siberian railway at the end of the 19th – beginning of the 20th century. *Vestnik Moskovskogo Universiteta. Series 8: History 2022*(5), 46–66. Open Access.

Sinan, M. and D. P. Derya (2023). Postal Organisation (Yam) in the Golden Horde. *11*(1), 37–61.

Truskott, I. F. (1760). Pochtovaya karta Rossiyskoy Imperii [Postal Map of the Russian Empire]. National Electronic Library of Russia, Map Collection.

Vinogradov, R. (2022). Chimeevo-1 – A 17th Century Yam Station. pp. 310–312.

Zhuravskaya, E., S. Guriev, and A. Markevich (2024). New Russian Economic History. *Journal of Economic Literature 62*(1), 47–114.

Abstrakt

Proč města vznikají právě tam, kde vznikají? Tato studie využívá pravidly determinovanou dopravní síť v Ruské říši ke zkoumání původu městských center. Poštovní systém *yams,* zavedený Mongoly ve třináctém století a udržovaný Moskevským státem, vyžadoval přepřahací stanice v pravidelných rozestupech pro výměnu koní. Tím vznikla infrastrukturní síť, jejíž rozestupy odrážely logistické požadavky spíše než geografii či předchozí osídlení. Digitalizujeme všechny stanice uvedené v ruském silničním průvodci z roku 1777 podél vzorku 15 hlavních tras a dělíme úseky (rays) mezi dvěma po sobě jdoucími stanicemi na buňky o velikosti 0,5 km. V moderních satelitních datech buňky nacházející se v historickém intervalu, v němž docházelo k výměně koní, vykazují zhruba o třicet procent vyšší intenzitu nočního světla než sousední buňky před tímto úsekem či za ním. Tento efekt je robustní vůči rozšíření o kontroly charakteristik první a druhé přirozenosti, fixním efektům jednotlivých tras i kontrole osídlení před rokem 1800, a u později vybudované Transsibiřské magistrály se neprojevuje. Další analýzy ukazují, že následný růst měst jen málo koreluje s geografickými danostmi, ale byl zesílen pozdějšími infrastrukturními investicemi, což naznačuje, že některé prvky ruské městské geografie byly založeny spíše na administrativních náhodách než na přírodních výhodách. Výsledky ilustrují, jak může prostorová nerovnost vznikat z arbitrárních historických koordinačních bodů a mít trvalé důsledky pro rozmístění ekonomické aktivity.

Online Appendix

Post(-Mongol) Roads to Path Dependence

Sebastian Ottinger
(CERGE-EI and IZA)

Elizaveta Zelnitskaia
(CERGE-EI)

A Historical Evidence on Post-Station Settlements

This appendix provides additional descriptive evidence on the continuity between historical post-station sites and modern urban settlements. Table A.1 lists a selection of towns that can be directly traced to *yams* established along Imperial Russia's postal routes. These examples illustrate how locations originally chosen to satisfy the logistical requirements of horse relays – typically spaced at 25–35 km intervals – frequently evolved into lasting administrative and population centers. Several such towns, including Kingisepp, Gavrilov-Yam, and Cheremkhovo, remain regional hubs today with populations exceeding 15,000–50,000 inhabitants. While anecdotal, these cases underscore the main mechanism documented in the paper: that administrative design, rather than natural geography, generated durable "city seeds." The complete list of digitized post stations, together with coordinates, distances, and data sources, will be available in the replication package accompanying this paper.

Table A.1: Post-Station Origins of Modern Towns in Russia

Historical Station (Name in 18th/19th c.)	Modern Name / Region	First Mention (Year)	Year as Post Station	Current Population	Approx. Distance to Next Station (verst)
Yamburg	Kingisepp, Leningrad Obl.	1384	Before 1384	47,312	24
Gavrilov Yam	Gavrilov-Yam, Yaroslavl Obl.	1545	-	17,648	-
Cheremkhovo	Cheremkhovo, Irkutsk Obl.	1722	Before 1722	50,586	39
Bokovskaya	Bokovskaya, Rostov Obl.	1873	Before 1873	4,550	-

Additional examples available upon request or in replication materials.

Notes: The table lists illustrative modern towns whose documented origins coincide with post-station sites. First-mention years reflect the earliest known references to settlements. Population figures correspond to the year 2020 and are taken from the dataset *Naselennye punkty Rossii: chislennost naseleniia i geograficheskie koordinaty* (Ministry of Health of the Russian Federation and ANO "TsPUR," 2021). Inter-station distances are derived from Ruban's (1777) *Foreign and Russian Road Guide*, measured in versts (1 verst $\approx$ 1.0668 km). *Sources:* Ruban (1777); Ministry of Health of the Russian Federation and ANO "TsPUR" (2021); Wikipedia.

B Routes and Station Counts

Table A.2: Digitized routes from Ruban (1777) and number of stations

Route (as in Ruban)	Direction	Endpoints (first → last)	Stations	Approx. length (verst)
St. Petersburg–Moscow	SSE	St. Petersburg → Moscow	*26*	*728*
Moscow–Belgorod	SSW	Moscow → Belgorod	*22*	*640*
Moscow–Smolensk	WSW	Moscow → Smolensk	*13*	*364*
Moscow–Saratov	SSE	Moscow → Saratov	*33*	*890*
Moscow–Arkhangelsk	NNE	Moscow → Arkhangelsk	*49*	*1293*
Murom–Kazan	E	Murom → Kazan	*17*	*466*
Koz'modem'yansk–Tobolsk	ENE	Koz'modem'yansk → Tobolsk	*56*	*1831*
Tobolsk – Irkutsk	SSE	Tobolsk → Irkutsk	*89*	-
Irkutsk – Khyagt	SSE	Irkutsk → Khyagt	*9*	*517*
Udinsk – Nerchinsk	NE	Udinsk → Nerchinsk	*7*	*651*
Nerchinsk–Nerchinsky Zavod	ESE	Nerchinsk → Nerchinsky Zavod	*6*	*248*
Moscow–Hlukhiv	SSW	Moscow → Hlukhiv	*17*	*550*
Moscow–Mozdok	SSE	Moscow → Mozdok	*70*	*2091*
Mscislaŭ–Chernihiv	SSW	Mscislaŭ → Chernihiv	*15*	*299*
Mscislaŭ–Smolensk	NNE	Mscislaŭ → Smolensk	*7*	*137*
		Total	*437*	*10705*

Notes: "Stations" counts the number of post relays digitized from Ruban (1777) along each route in our working sample. Approximate route length is calculated as the sum of distances between consecutive stations along the route, measured in versts (1 verst ≈ 1.0668 km), based on Ruban (1777) data. A part of the data for the Tobolsk – Irkutsk route comes from a different source, the historical map from Pyadyshev (1828), since the data from Ruban (1777) was unidentifiable for this part. Thus, the distance between stops for this route is omitted in the table. Direction labels follow the orientation from the first to the last endpoint of each route.

C Robustness Checks

This appendix assembles the full set of robustness exercises that underpin the baseline results. We organize the evidence into four parts. First, we show that alternative outcome definitions and clustering schemes leave the estimates essentially unchanged. Second, we vary the definition of the mechanical relay interval used in the instrument and expand/shrink the analysis window along rays; the results are stable. Third, we examine heterogeneity across post roads; the pattern is broad-based rather than driven by a single corridor. All specifications include ray fixed effects and, where indicated, the complete set of first- and second-nature controls. Standard errors are clustered at the ray level unless stated otherwise.

C.1 Alternative Outcome Definitions and Clustering

We re-estimate the main specifications using (i) the original continuous night-light measure, (ii) a binary indicator for any light (> 0), and (iii) the logarithm of light intensity (adding a small constant where necessary). We also report variants that cluster at other levels (adding stop and segment clusters to the baseline road clustering). Across all choices, both the OLS and IV coefficients remain positive, precise, and close in magnitude to the baseline, indicating that inference does not hinge on functional form or the clustering level.

Table A.3: Functional Forms and Clustering

Dep. var.:	Continuous lights (baseline)			Dummy: Light>0	Log(Light)
Panel A. OLS					
Post Stop	6.972***	6.972***	6.972***	0.194***	1.565***
	(0.388)	(0.768)	(0.839)	(0.028)	(0.166)
R^2	0.64	0.64	0.64	0.53	0.60
Observations	30,583	30,583	30,583	30,583	30,583
Panel B. Reduced-Form Regressions					
$\mathbb{1}(25 <$ Dist. $< 35)$	1.861***	1.861**	1.861**	0.034***	0.305***
	(0.545)	(0.651)	(0.656)	(0.009)	(0.067)
Panel C. IV Regressions					
Post Stop	99.604***	99.604***	99.604***	1.831**	16.327***
	(24.880)	(25.863)	(25.863)	(0.656)	(4.707)
First-Stage F-Stat	21.76	20.37	20.37	21.76	21.76
Additional Info (Panels A-C)					
Clustering	Road# (baseline)	+ Stop #	+ Segment #	Road#	Road#

Notes: All columns include ray fixed effects and the full set of first- and second-nature controls. The bottom line of the table indicates the clustering level used in each panel variant. The sample consists of cells 10–50 km from the preceding station along each ray.

C.2 Relay-Range Definitions and Ray Length

We next vary the instrument's range indicator to use wider (20-40 km and 15-45 km) horse-relay intervals and also alter the analysis window along rays (0–50 km and 0–80 km, in addition to the baseline 10–50 km). Estimates are nearly unchanged across these alternatives, confirming that the core effect is not an artifact of the exact 25–35 km cutoff nor of the truncation choice.

Table A.4: Alternative Relay Cutoffs

Dep. var.: Light at night in a 500 m cell on ray from last post station					
Panel A. OLS					
Post Stop	6.443***	7.335***	6.972***	6.972***	6.972***
	(0.299)	(0.360)	(0.768)	(0.768)	(0.768)
R^2	0.58	0.48	0.64	0.64	0.64
Observations	38,368	61,278	30,583	30,583	30,583
Panel B. Reduced-Form Regressions					
$\mathbb{1}(25 <$ Dist. $< 35)$	1.122**	2.165***	1.861***		
	(0.468)	(0.635)	(0.545)		
$\mathbb{1}(20 <$ Dist. $< 40)$				1.867***	
				(0.534)	
$\mathbb{1}(15 <$ Dist. $< 45)$					1.762***
					(0.493)
Panel C. IV Regressions					
Post Stop	95.163**	113.503***	99.604***	103.759***	134.538***
	(35.950)	(26.556)	(24.880)	(21.735)	(39.788)
First-Stage F-Stat	9.32	23.43	21.76	39.40	37.44
Additional Info (Panels A-C)					
Sample	0-50 km	0-80 km	10-50 km (base)		

Notes: All specifications include ray fixed effects and the full set of first- and second-nature controls. "Sample" rows indicate the ray-length window used in each block. Standard errors clustered at the ray level.

C.3 Heterogeneity Across Post Roads

Finally, we examine heterogeneity by road. Figure A.1 reports route-specific coefficients (with identical controls and fixed effects), revealing largely positive point estimates for individual routes, with overlap in confidence intervals. Routes close to Moscow and Saint Petersburg tend to have higher effects than average (e.g., Saint Petersburg – Moscow, Moscow – Smolensk, Moscow – Saratov, Moscow – Arkhangelsk, note that the ordering in Table

A.2 corresponds to the ordering here). In contrast, routes in very mountainous terrain (e.g., Irkutsk – Khyagt, route number 29) and routes near the present-day borders of Belarus with Russia and Ukraine, in the already more densely settled terrain (e.g., Mscislaŭ – Chernihiv, 57; Mscislaŭ – Smolensk, 62) exhibit even significant negative effects. This suggests that our empirical strategy works best in unsettled non-mountainous terrain where direct lines are a good approximation of the actual route and where horses indeed have to be interchanged at those specified intervals. Notably, all the routes with significantly negative coefficients are also relatively short.

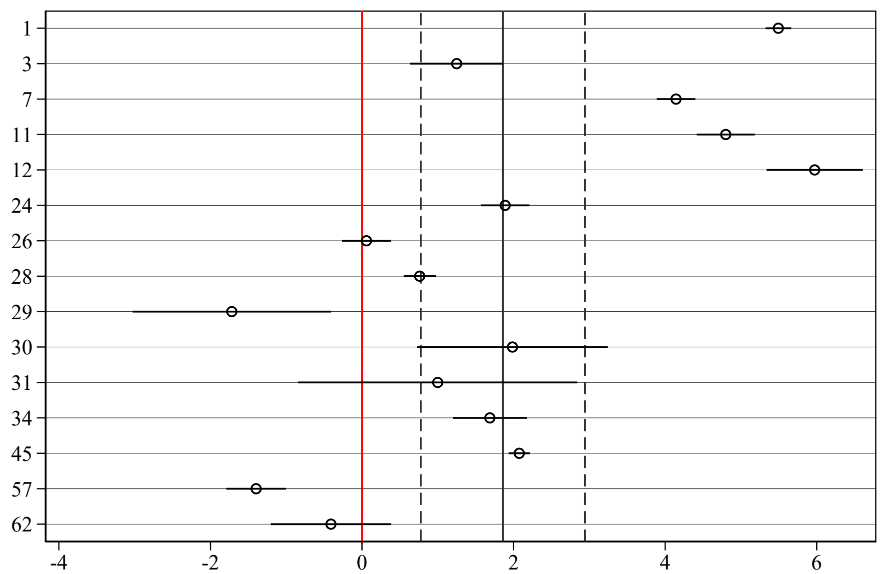

Figure A.1: Heterogeneity by Post Road

Notes: Route-level coefficients (and 95% CIs) from regressions analogous to the baseline specification, including ray fixed effects and all first- and second-nature controls.

D Segment-Level Baseline Regressions

This appendix repeats the baseline analysis of Section 4 at a coarser spatial resolution, aggregating the 500 m observation cells into ray segments of 10–25 km, 25–35 km, and 35–50 km from each preceding post station. The goal is to verify that the main results are not driven by local measurement noise or pixel-level variation in night-light data, and to provide a scale more directly comparable to historical travel distances.

For each segment s along ray r, we estimate the following reduced-form specification:

$$\text{Light}_{sr} = \beta\, \mathbb{1}(\text{Range}_{sr}) + \gamma X_{sr} + \mu_r + \varepsilon_{sr}, \tag{A.1}$$

where $\mathbb{1}(\text{Range}_{sr})$ equals one if the segment lies within 25–35 km of the previous station, X_{sr} is the vector of first- and second-nature controls, and μ_r are ray fixed effects. As in the main analysis, standard errors are clustered at the ray level.

Table A.5: Replicating the Baseline at the Ray-Segment Level

Dep. var..: Light at night in segment on ray from last post station

	(1)	(2)	(3)	(4)
Stage:	OLS	Red. Form	2SLS	
Dummy: Post Stop	4.520***		6.492***	4.165**
	(0.465)		(1.532)	(1.578)
Dummy: 25 < Range <35 km		1.611***		
		(0.489)		
R^2	0.72	0.70	0.07	0.37
Observations	1,263	1,263	1,263	1,161
First-Stage F-Stat			21.70	26.75

Note: All columns include ray FE, and all first nature controls. The sample consists of three segments per ray, where the second one, in which the post station is expected, is titled "Range' above. Robust standard errors clustered at the ray-level.

Table A.5 shows that the segment-level estimates mirror those obtained at the 500 m resolution. The brightness of segments located in the historical horse-changing range (25–35 km) is roughly 25 percent higher than that of adjacent segments before or after it.

Working Paper Series
ISSN 2788-0443

Individual researchers, as well as the on-line version of the CERGE-EI Working Papers (including their dissemination) were supported from institutional support RVO 67985998 from Economics Institute of the CAS, v. v. i.

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the CAS, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, Czech Republic
Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: https://www.cerge-ei.cz/

Editor: Byeongju Jeong

The paper is available online at https://www.cerge-ei.cz/working-papers/.

Electronically published December 3, 2025

ISBN 978-80-7343-614-8 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)