Geography, Technology, and the Extent of Markets:
A Method with an Application to the Wealth of Nations

Roy Elis, Stephen Haber, and Jordan Horrillo*

Abstract: People make decisions based on the extent of their market, but analyses of the resulting phenomena are performed on latitude-longitude grid cells. This generates threats to identification because market extents vary over space and time, but grid cells are of uniform size and are fixed in time. We develop a method to measure market extents as contiguous, technology-conditioned cost surfaces, such that they vary based on features of the Earth's surface and a state of technology. We develop two novel methods to analyze the resulting spatial data. We demonstrate the utility of these new methods of measurement and analysis by testing the predictions of Adam Smith's theory that economic development is bounded by the extent of the market, and the extent of the market is bounded by features of the Earth's surface. The results indicate that our methods have applications beyond *The Wealth of Nations*.

We acknowledge support from Stanford's IRiSS, FSI, VPUE, Center for Global Poverty and Development, and Hoover Institution, the UPS Foundation, and the John Templeton Foundation. Melda Alaluf, Ava Aydin, Nicholas Baldo, Andrew Brooks, Roger Cain, Victor Cárdenas Codriansky, Ian Claras, Kevin Cook, Eva Davis, Kara Downey, Lauren Felice, Adriane Fresh, Anne Given, Qianmin Hu, Kirill Kalinin, Abigail Kessler, Scott Khamphoune, Melissa Lee, Michelle Lee, Ian Lewis, Cole Lupoli, Annamaria Prati, Cara Reichard, Felipe Rodriguez, Armand Seru, Jeffrey Tran, Natalia Vasilenok, Xinyi Wang, and Tracy Williams provided invaluable research assistance. Avidit Acharya, Gustavo del Angel, Mark Dincecco, Justin Grimmer, Chad Jones, Dorothy Kronick, Joel Mokyr, Jean-Laurent Rosenthal, Paul Sniderman, and William Summerhill made helpful comments on an earlier draft.

---

* Roy Elis, Chime Financial Inc. (roy.elis@chime.com); Stephen Haber—corresponding author—Stanford University (haber@stanford.edu); Jordan Horrillo, Stanford University (horrillo@stanford.edu)

Section 1: Introduction

Human beings make economic decisions based on the extent of markets. The collective behavior that emerges from those decisions generates a wide range of phenomena. Econometric analyses of those phenomena were, until recently, carried out on political jurisdictions (e.g., Masters and McMillan 2001; Easterly and Levine 2003). Scholars have increasingly shifted to latitude-longitude grid cells as the unit of analysis because of their higher degree of resolution (e.g., Nordhaus 2006; Motamed, Florax, and Masters 2014; Henderson et. al. 2018; Rossi-Hansberg and Zhang 2026). Researchers are aware of the lack of congruence between a grid cell and the theoretical concept of a market—a contiguous geographic space bounded by the cost of shipping goods across it relative to the value of the surplus generated by trade, conditional on the technologies available at a point in time and the features of the earth's surface at that point in time. Nevertheless, they conduct analyses on grid cells because, as Nordhaus (2006) explains: "From a practical point of view, there is no alternative to a grid measurement system…"

Relying on grid cells, rather than market extents, generates three threats to identification: standard errors may be biased; spatial trends may remain undetected; and controls must be added to mitigate spatial autocorrelation (Conley and Kelly 2025). Plainly put, a market might be smaller than a grid cell, larger than a grid cell, or be composed of a ribbon of territory that winds its way through the contiguous interiors of numerous grid cells. Making inference more difficult still, grid cells are fixed in time, but market extents are not; they expand as technological change reduces transport costs.

We address these limitations of grid cells by developing a method to measure markets as contiguous, technology-conditioned cost surfaces. Our method estimates market extents based on the features of the earth's surface at a point in time, the state of technology at that point in time, and an energy budget that has verisimilitude for that point in time. Our approach facilitates the identification of spatial trends, allows researchers to estimate integration across markets, and permits the model to estimate the effects of changes in transport technologies.

We develop two novel methods for the analysis of the resulting spatial data. Our first method generates GIS layers of contiguous markets whose physical characteristics are correlated. This allows us to detect large scale spatial trends. Our second method estimates random forest regressions on stratified samples of non-overlapping markets. This allows us to conduct multivariate analyses on observations that vary in size, while at the same time maintaining independence across observations.

We demonstrate the utility of these new methods of measurement and analysis by operationalizing a 250-year-old idea that is central to economics but whose predictions have never been properly tested; the theory in Smith (1776) that the level of economic development is bounded by the extent of the market, and the extant of the market is bounded by three features of the Earth's surface; navigable rivers, jagged coastlines that permit access to the sea, and the fertility of soils. As Smith (1776) famously put it:

> England, on account of the natural fertility of the soil, of the great extent of the sea-coast in proportion to that of the whole country, and of the many navigable rivers which run through it, and afford the conveniency of water carriage to some of the most inland parts of it, is perhaps as well fitted by nature as any large country in Europe, to be the seat of foreign commerce, of manufactures for distant sale, and of all the improvements which these can occasion. (Smith, 1776, Book 1, chap 3).

The contrast Smith (1776) drew to the markets of other world regions were stark.

> The Sea of Tartary is the frozen ocean which admits of no navigation, and though some of the greatest rivers in the world run through that country [Central Asia], they are at too great a distance from one another to carry commerce and communication through the greater part of it. There are in Africa none of those great inlets, such as the Baltic and Adriatic seas in Europe, the Mediterranean and Euxine [Black] seas in both Europe and Asia, the gulfs of Arabia, Persia, India, Bengal, and Siam, in Asia, to carry maritime commerce into the inland parts of that great continent; and the great rivers of Africa are at too great a distance from one another to give occasion to any considerable inland navigation. (Smith, 1776, Book 1, chap. 3).

To operationalize the theory in Smith (1776) we draw 41,435 markets from initial points placed 60 kilometers apart in Eurasia and Africa (including their major islands and archipelagos) conditional on 18th century transport technologies and on the 18th century state of rivers, lakeshores, sea ice, and sea access.[1] This allows us to draw 18th century market extents at a high degree of resolution over a large geographic expanse, while retaining computational efficiency such that our model –and extensions to it—can be run by researchers on a desktop computer. Because testing the predictions of Smith (1776) requires us to estimate levels of economic development in the 18th century at the level of the market, we construct a novel dataset on urban populations with coverage in 1700 and 1800. In addition, because one of the factors specified by Smith (1776) was soil fertility, we make an improvement to extant approaches to using the FAO GAEZ 4.0 database to estimate historical levels of potential agricultural output by capturing the effect of traditional flood irrigation in river valleys which was crucial for farmers in arid regions.

We stress that our goal in testing the predictions of Smith (1776) is to demonstrate the utility of our methods of measurement and analysis, rather than to estimate the causal relationship between geographic factors and the distribution of economic development in the 18th century. Given the state of 18th century knowledge, Smith (1776) took account of only three factors. We hope that other researchers will employ the methods we have developed to test the predictions of other theories about the world distribution of income, as well as to test theories about other outcomes of interest in which the market is the theoretically appropriate unit of analysis. We will therefore (upon publication) make all the GIS shape files, computer scripts, and data available to other scholars so that they may tune our geospatial model to other applications.

---

[1] We exclude the Americas, Australia, and New Zealand from our analysis because their technological isolation from Eurasia since the beginning of the Holocene until Europeans arrived in the 1500s, 1700s, and 1800s, respectively, and their subsequent demographic collapses because of the introduction of new diseases, gave rise to economic dynamics different from those that were in operation in Eurasia and Africa during the centuries leading up to Smith (1776).

The rest of this paper proceeds as follows. Section 2 explains the conceptual basis for our measure of market extents. To make the discussion concrete, we illustrate those concepts by focusing on the factors that were of concern to Smith (1776), conditional on the technologies and the features of the Earth's surface in the 18th century. Section 3 explains how we parameterize our geospatial model. It also explains our methods to estimate urban populations as well as our improvements to the estimation of historical soil fertility. Section 4 employs the model and data developed in Section 3 to test the predictions of Smith (1776) by generating clusters of contiguous markets whose physical characteristics are correlated. Section 5 employs the model and data developed in Section 3 to test the predictions of Smith (1776) by estimating random forest regressions on stratified samples of non-overlapping markets. Section 6 concludes by discussing extensions to our model.

Section 2. The Market as the Unit of Analysis

Human beings must expend energy to move goods across the Earth's surface. How far they can move those goods is, therefore, bounded by an energy budget, the features of the Earth's surface that force them to spend that budget overcoming gravity, friction, and drag, the performance characteristics of their transport technology, and the gross weight of their cargo. It follows that the geographic extent of a market at a point in time can be estimated if a researcher provides a GIS Path Distance Allocation tool with an initial point, a rasterized Digital Elevation Model (DEM), a rasterized cost surface that gives the physical cost of moving a specified weight one meter through each raster cell using the transport technologies available at that point in time, a set of assumptions about the physical limits of those transport technologies, a physical model that translates elevation gain into work expended, and a budget constraint that has historical verisimilitude.

The next section of this paper explains how we derive those parameters to estimate 18th century market extents. Our goal in this section is to show readers how using market extents as the unit of analysis yields insights that would be difficult to capture using grid cells as the unit of analysis.

Map 1 provides a visualization of three markets generated by our geospatial model using the most efficient combination of three 18th century transport technologies to move one metric ton of bulk goods with an energy budget set at 40 megajoules (the amount of energy consumed by an 80-kilometer wagon haul over flat, nearly dry ground, explained in detail in Section 3). The top frame shows the market drawn from an initial point close to Paris. The middle frame shows the market drawn from an initial point close to Kayseri, Turkiye. The bottom frame shows the market drawn from an initial point close to N'shenge, the capital of the Kuba Kingdom in the present-day Democratic Republic of the Congo. Our estimate of the extent of each market is shown in green. The navigable rivers running through the markets (conditional on 18th century technologies and river obstacles), which influence market extents, are shown in blue. The terrain slopes that influence market extents, are shown in shades of black. The natural harbors of each market (which proxy for the concept of access to the sea in Smith, 1776) are shown as red dots. The three maps are drawn to the same scale. The box to the right of each map shows the numerical values for each market's size, potential staple crop production per square kilometer, and natural harbors.

Map 1: The Market Extents of Paris, Kayseri, and N'shenge

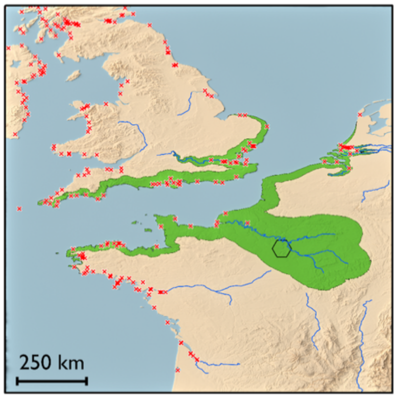

Paris, France

Area (1000 km²)	121.7
Agricultural Potential (Mkcal/km²/yr)	490.0
Harbors (count)	84

Kayseri, Turkey

Area (1000 km²)	19.5
Agricultural Potential (Mkcal/km²/yr)	362.2
Harbors (count)	0

Nshenge, DRC

Area (1000 km²)	0.5
Agricultural Potential (Mkcal/km²/yr)	308.3
Harbors (count)	0

Map 1 makes clear why Paris was one of the wealthiest cities in the world in the 18th century, why Kayseri was a caravanserai, and why N'shenge was a large village. Shippers in Paris could move bulk goods up and down the Seine and its tributaries, allowing them to then make use of the English Channel, the North Sea, the Thames River in England, the Western Scheldt River in the Netherlands, and the Rhine–Meuse–Scheldt

Delta in Belgium. Once Paris-based shippers landed their cargoes, they could then move them long distances overland using horse-drawn wagons, because of the extraordinarily flat terrain of the North European Plain, as shown in the top frame of Map 1. The Paris market was roughly six times larger than the Kayseri market, which had no navigable river access and was bounded by mountains to its south and east. The Paris market was 257 times larger than the N'shenge market, whose shippers had to use human porters, rather than horse-drawn wagons, to carry bulk goods overland, because endemic tsetse flies carried trypanosomes that killed draft animals (Alsan 2015). As Map 1 shows, overland transport from N'shenge was so costly that shippers exhausted the 40 megajoule energy budget before they reached the banks of the Sankuru and Kasai rivers, just to their north and west, respectively. This fact explains why N'Shenge was so isolated that the first Europeans did not step foot in it until the 1880s, roughly 400 years after they first arrived on the Central African coast (Vansina 1978: 78, 236).

Large size was not the only advantage of the Paris market. As Map 1 shows, its potential agricultural production per square kilometer (the primary source of energy to fuel market activity) was 35 percent greater than that of the Kayseri market and 59 percent greater than the N'shenge market. The combination of larger size and greater soil fertility meant that the Paris market could produce primary sector energy output eight times that of Kayseri and 409 times that of N'shenge. In addition, shippers in Paris could access the sea from 84 natural harbors, as compared to zero for Kayseri and N'shenge.

Map 1 makes clear one of the advantages of the market as the unit of analysis over grid cells; markets vary in size, but grid cells do not. Let us draw out this point by focusing on the Paris market as compared to a half-degree hexagonal grid cell centered on Paris (shown in Map 1). First, the Paris market is nearly 50 times the size of the area covered by the grid cell. In addition, the market as the unit of analysis shows that Paris-based shippers could access the sea through 84 natural harbors. A grid cell-based analysis would instead include a parameter for the distance as the crow flies from the centroid of the Paris cell to the nearest natural harbor or coastline, Le Havre, 177 kilometers away (e.g., Motamed, Florax, and Masters. 2014; Henderson et. al. 2018). A contrast running in the other direction can be drawn about a half degree grid cell and the N'shenge market, in which the market is only one-eighth the size of the grid cell.

Map 1 points to an additional advantage of employing the market as the unit of analysis; it allows researchers to capture network effects that would be difficult to capture by an approach based on grid cells. Paris was one of the wealthiest cities in the world in the late 18th century, with a population of 550,000, not because the half degree grid cell around it contained fertile soils and a river segment, but because it was a node in an extended network of markets and cities. Other cities that fell within the boundaries of the Paris market included London (population 959,000), Brussels (population 66,000), Antwerp (population 62,000), Gent (population 55,000), Rotterdam (population 53,000), and Bruges (population 31,000). Recall, moreover, that our model draws markets from initial points located 60 kilometers apart which means that there are 98 other initial points located inside the Paris market (e.g., initial points close to London, Brussels, Antwerp, Gent, Rotterdam, and Bruges) which generate their own markets (not shown) that both overlap the boundaries of the Paris market and extend beyond it. By comparison, the Kayseri market contains only seven other initial points, while that of N'shenge contains none. We will have more to say about this in Section 4, but for now we highlight that the physical characteristics of many of those 98 markets in the Paris network are correlated with the physical features of the Paris market. In short, an approach that uses markets as the unit of analysis can detect a large-scale spatial trend that, as we discuss in Section 4, is found in only one other region of Eurasia and Africa.

Section 3. Model Parameterization

3.A. Estimating Potential Market Sizes and Shapes

We begin by placing 41,561 hexagonal grid cells on a plate carrée projection of the terrestrial surface of Eurasia and Africa, including their major islands and archipelagos, whose centroids are set at a uniform distance of 60 kilometers from one another.[2] We

[2] We subsequently exclude 126 of these either because they do not have a GAEZ4 raster to measure potential agricultural production or because the initial points landed on the

choose 60 kilometers because it allows shippers of bulk goods traversing a level terrain from a centroid to the closest neighboring centroid using a horse and wagon without exceeding the 40 megajoule energy budget. If the centroid of a grid cell falls in a river, lake, or ocean, or if a grid cell is cut by one of those bodies of water, we snap the centroid to the nearest riverbank, lakeshore, or coast under the assumption that, all else equal, human beings would settle close to that body of water. We follow a similar protocol if a grid cell is cut by the boundary of tsetse fly endemicity (discussed below); we snap the centroid to the nearest point outside of the tsetse-endemic area on the assumption that, all else equal, human beings would settle where they could employ draft animals. We refer to these amended centroids as "initial points."

From each initial point we then develop vector and raster datasets that allow us to estimate the accumulated costs of moving a metric ton of bulk goods over a cost surface using the most efficient combination of an 18th century (pre-steam) river boat, horse-drawn wagon, and human porter using the Path Distance Allocation tool in ArcGIS Pro 3.3. From any initial point, the tool begins calculating the accumulated transport costs at that point and works its way outwards until the accumulated transport cost reaches the budget constraint. That is, we estimate least-cost paths, conditional on 18th century transport technologies, in which the uneven spatial distribution of navigable water, flat terrain, and tsetse fly endemicity drive variation in the distance one metric ton could be moved. We set costs and the budget in terms of physical parameters (friction, energy, work), rather than economic costs (dollars per ton-mile).

We supply the Path Distance Allocation tool with six inputs that model bulk transport in the 18th century. The first is a point dataset locating each initial point. The second is a rasterized Digital Elevation Model (DEM) giving the mean elevation of every 1 km × 1 km cell that allows the tool to charge transportation costs for elevation change (fighting gravity) and measure true distance across a three-dimensional surface (the hypotenuse rather than the horizontal leg when traveling uphill or downhill). The third is

---

slopes of steep mountains, such that potential markets drawn from them are trivially small (less than 30 square kilometers). This leaves us with 41,435 markets.

a rasterized cost surface giving the physical cost of moving one metric ton of goods one meter through each raster cell under the three technologies discussed below. The fourth is a physical model translating elevation gain into work expended. The fifth is a set of assumptions about the physical limits of wheeled transport (the maximum grade up which a horse can pull a load). The sixth is a budget constraint expressed in units of energy.

3.A.1 Transport Technologies

To estimate market extents, we need to specify the transport technologies that shippers could draw upon and identify the features of the Earth's surface on which those technologies could have been applied during the 18th century.

3.A.1.A Overland Transport

Wagon transport was the most efficient option for the overland transport of bulk goods in the 18th century. The basic design and construction of horse-drawn, four-wheeled wagons were largely the same the world over because knowledge about them had been disseminating since their first appearance in Bronze Age Eurasia. To parameterize our model, we draw on information about the Conestoga Wagon, built by German immigrants in 18th century Pennsylvania, because its wheel dimensions, weight, load sizes, and performance characteristics across various road conditions have been documented by historians.

We estimate the amount of energy required for a loaded Conestoga wagon to overcome friction (rolling resistance) to maintain a constant speed on flat land and then account for the additional energy required to overcome elevation changes. Wheel diameter and wheel width are crucial factors in determining the rolling resistance of a wagon. Based on Shumway (1964), we estimate an average rear wheel diameter of 59.9 inches, an average front wheel diameter of 44.8 inches, and average wheel widths for both front and back of 3.4 inches. Rolling resistance, for any given wheel diameter and width, is a function of the gross weight of the vehicle because the road surface deforms in proportion to the gross weight as spread over the surface of the wheel in contact with the road. We assume a wagon weight of 3,250 pounds and a cargo of 7,000 pounds based on data in Shumway (1964: 66).

Road surface conditions could impact the energy cost of wagon transport tremendously; poor conditions could increase rolling resistance by an order of magnitude, relative to good conditions. Based on data in Baker (1918, ch. 1) we assume a resistance of 68.5 pounds per ton based on the average for freight wagons on a nearly dry earth road in good condition. This decision biases upwards the size of markets in tropical environments, where heavy rains turned roads to mud, thereby biasing against the hypothesis in Smith (1776) that England and Continental Europe had large market extents relative to Africa. This decision also biases downwards the size of potential markets in Western Europe, where masonry roads constructed by the Roman Empire allowed for less road surface deformation than nearly dry earthen roads, again biasing against the hypothesis in Smith (1776).

We estimate 492 newtons as the force required to overcome the total resistance (per metric ton of cargo) of a loaded Conestoga wagon on a flat, nearly dry, earthen road in good condition. Applying that force over 1 meter implies 492 newton-meters (joules) as the amount of energy expended per metric ton. We encode this value for wagon-based overland travel in the friction surface used by the Path Distance Allocation tool.

To account for the energy required to overcome elevation changes, conditional on gross weight of the wagon and the gradient of the path, we draw on data in Baker (1918). Following Baker (1918) we limit wagon transport to a 10 percent grade, at which point a horse can barely generate enough power to transport itself over any appreciable distance. Because terrain slopes change on geologic time scales, we take terrain slopes as they exist today. Map 2 displays Eurasian and African terrain slopes in exaggerated shaded relief, showing the regions with flat terrains that would have facilitated the use of horse drawn wagons, the hilly regions that would have rapidly consumed the energy budget, and the steep regions that would have made it impossible to employ horse-drawn wagons. We treat slope costs as symmetric in the direction of travel: ascending and descending a given grade incur equal impedance because for a heavily laden 18th century wheeled vehicle descent was not energy-recovering; braking and controlling a loaded wagon on a steep downgrade is itself costly. Our protocol therefore prevents the least-cost path algorithm from treating downhill segments as a source of "negative-cost" energy.

In regions of Africa south of the Sahel and north of the Veldt where tsetse flies were endemic, wagon transport would not have been an available option because the trypanosomes they carried killed horses, mules, and oxen (Alsan 2015). We therefore assign human porters to any raster cell in which tsetse flies were endemic, drawing on a predicted distribution of tsetse flies in raster format developed by the FAO.[3]

We apply some basic mechanics to estimate the energy cost of moving one metric ton of bulk goods one meter using human porters, as compared to horse-drawn wagons. Human porters, because they carried loads on their shoulders or foreheads had to forego the mechanical advantage afforded by a wheel and axle. A Conestoga Wagon with a 60-inch diameter wheel typically had a five-inch axle, yielding a mechanical advantage of 12. Human porters would, however, gain some advantage over a horse drawn wagon by being able to take paths too narrow for a wagon. We therefore scale down the 12:1 mechanical advantage of horse drawn wagons over human porters to 10:1, thereby biasing upwards the size of markets in Africa's tsetse fly endemic regions. Horses are also more efficient than human beings, because they exert power and balance loads over four legs, rather than two. We do not include this differential in our estimates but note that its elision underestimates the advantages of horse-drawn wagons over porters, thereby creating a bias in favor of markets in tsetse fly endemic regions. We therefore assume 4,920 newton-meters to move one metric ton one meter and encode this value into the friction cost raster. We account for elevation changes using the same parameters we use for wagon transport.

[3] http://www.fao.org/geonetwork/srv/en/main.home?uuid=f8a4e330-88fd-11da-a88f-000d939bc5d8 The tsetse distribution index runs from 0 to 1, but the distribution of values is bimodal. To determine the breakpoint, we consulted country histories to determine whether horses were employed and found that regions with values below 0.1 historically employed horses, while those above 0.1 did not. If a raster cell has an index value greater than 0.1 we therefore require transport through it to use human porters.

3.A.1.B Water Transport

In dealing with water-based modes of transportation, we make several simplifying assumptions. We do not distinguish river travel from travel on lakes, seas, and oceans. That is, we assume that the same watercraft were used as on rivers. Given that the ships used for coastal and ocean transport were larger and more efficient than the boats used in riverine transport, this assumption leads us to overestimate the energy costs of ocean transport, thereby biasing against the hypothesis advanced in Smith (1776). In addition, we assume away the effect of river and ocean currents, with the exception that we code stretches of rivers with sufficiently strong currents, as indicated by contemporary sources per our discussion below, as non-navigable. We limit coastal and ocean transport to zones that were ice-free prior to climate change for at least six months per year, per our discussion in section 3.A.1.C, below.

We estimate the amount of energy required to overcome drag in moving a loaded boat, at a constant speed, through still water. We rely on historical data to determine the physical characteristics of typical boats, load sizes, drafts, and drag coefficients.

Estimating the physical cost of transportation over water requires that we fix a pre-steam transportation technology. We chose the keelboat and the white pirogue used by Lewis and Clark to sail, row, and pole their way up the Missouri River in 1804 because those boats were typical of the small vessels used for riverine transport around the world in the 18^{th} century and detailed knowledge of their shapes, drafts, and cargo capacity has been preserved. We note that while these boats were less capacious than river barges or flat boats, they were better designed to navigate smaller rivers both upstream and downstream.

Drag is the main source of resistance encountered by an object moving through water at low speeds, and is determined by the formula F_D=1/2 ρv^2 CD A, where FD is the drag force, ρ is the density of fluid (in this case water), v is the relative velocity of the object, CD is the drag coefficient, and A is the cross-sectional area of the object submerged in water. The value of ρ for water at 4°C is approximately 1000 kg/m3. We assume a velocity 3 miles per hour. We calculate the cross-sectional area based on the shapes and drafts of the keelboat and pirogue from Boss (1993), Mussulman (ND-A), and Mussulman (ND-B). We employ a drag coefficient of 0.295 based on the shape the keelboat and pirogue, which were rounded in front, squared at back.

We estimate a drag force of 56.6 newtons (per metric ton of cargo) for the keelboat and a drag force of 49.1 newtons (per metric ton of cargo) for the pirogue. Averaging the two, we arrive at the estimate of 53 newtons of drag force (per metric ton of cargo) for waterborne transport. Applying that force over 1 meter implies 53 newton-meters, or joules, as the amount of energy expended per metric ton in hauling cargo through water, and encode this value in the friction surface used by the Path Distance Allocation tool.

3.A.1.C Navigable Waterways

We estimate as closely as we can the navigability of the world's waterways in the 18th century. We code all lakes as navigable, except those above the Arctic Circle (66 degrees north latitude), because they would have been frozen for at least six months per year. We restore the shorelines of Lake Texcoco, Lake Zumpango, Lake Chalco, Lake Zaltocan, Lake Xochimilco, Lake Chad, and the Aral Sea to their approximate extents in the 18th century.[4] We remove man-made lakes and reservoirs. We code all oceans as navigable, except those that were frozen (at a threshold of 60 percent) per the analysis of Vanhatalo et. al. (2021) on the probability of a ship becoming beset in sea ice. We draw sea ice coverage from the NOAA, Gridded Monthly Sea Ice Extent and Concentration dataset, and use the observation for April, 1850 because it represents the earliest year in the dataset (i.e., outside the range of global warming), with April reflecting the conditions that would persist until the onset of the following winter.[5]

Estimating the navigability of rivers requires several steps. The invention of dynamite in 1867 and earth moving machinery in the 1920s allowed human beings to

---

[4] We draw on the Natural Earth free vector and raster maps data, available at https://www.naturalearthdata.com/downloads/10m-physical-vectors/10m-lakes/

[5] https://nsidc.org/data/g10010/versions/1

widen and deepen river channels, dredge sand bars, and dig canals around waterfalls and rapids at an unprecedented scale. One cannot, therefore, simply rely on present-day maps of navigable rivers. We start with a universe of cases consisting of major perennial rivers from the Natural Earth dataset (http://www.naturalearthdata.com/), as well as the CIA World Databank II (now incorporated into the Global Self-Consistent, Hierarchical, High-resolution Geography Database http://www.soest.hawaii.edu/pwessel/gshhg/). We then combed historical sources, most particularly travelers' accounts, that describe navigation and attempted navigation on those rivers using pre-steamboat technologies. We code as navigable any stretch of river that could be traveled upstream and downstream at least six months per year using boats of a similar size and shape discussed in 3.A.1.B above. That is, we remove from major present-day rivers the stretches that could not be navigated during the 18th century because of ice, rapids, waterfalls, sandbars, and marshes. We exclude all identifiable canals, except for China's Grand Canal, because it was constructed between the 7th and 12th centuries, making it infeasible to redraw the likely path of the Yellow River had there been no canal.

Map 2 shows the stretches of rivers that were navigable using 18th century technologies. It also displays the terrain slopes that either made the use of horse drawn wagons difficult (in brown) or impossible (in black), per Section 3.A.1.A, above.

Map 2: The Navigable Rivers of Eurasia and Africa in the 18th Century

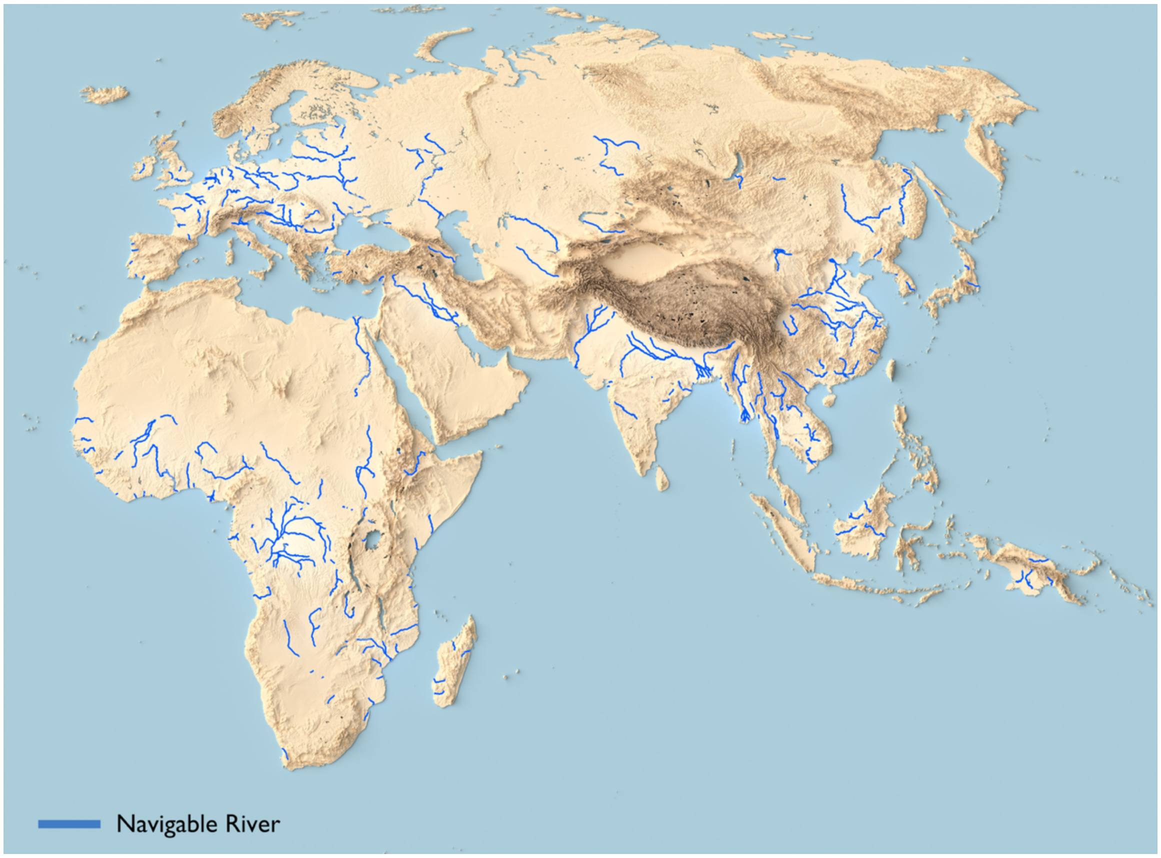

3.A.2 The Budget Constraint

The final input used for calculating the size and shape of a market is the budget constraint (the amount of energy available to shippers). We set the energy budget at 40 megajoules, the amount of energy it would take to transport a metric ton of goods 80 kilometers on a perfectly flat, nearly dry, earthen road using a horse and wagon. We choose this distance based on information in Fogel (1964: 75-79) about the distance of typical wagon hauls for staple agricultural goods in the United States. We conduct sensitivity analyses in Sections 4 and 5, in which we vary the 40 megajoule energy budget by two megajoule increments running from 30 to 50 megajoules, and find that our results are not sensitive to these tuning parameters.

3.B. Access to the Sea

We operationalize the insights in Smith (1776) about access to the sea by counting the number of natural harbors in every market. We follow Henderson et. al. (2018) by using the 1953 (first) edition of the U.S. Navy's *World Port Index*. We code as natural harbors those the U.S. Navy categorized as a coastal natural harbor, riverine natural harbor, or river basin natural harbor, but exclude riverine and river basin natural harbors that were located on river stretches that were not navigable with 18th century technologies. We also exclude coastal natural harbors that were located on bodies of water that were beset by sea ice, per 3.A.1.C., above. We supplement this data with information on Africa's natural harbors from Deasy (1942), thereby increasing the number of natural harbors in Africa to bias against the hypothesis in Smith (1776). Because U.S. Navy (1953) does not include Caspian Sea natural harbors, we draw on information from the 1911 *Encyclopedia Britannica* to obtain identify Caspian Sea natural harbors. Map 3 displays the results. It shows that there was only one region of Eurasia and Africa, the North European Plain, in which there existed a dense network of navigable rivers that flowed to the sea, flat terrains, and natural harbors.

Map 3: The Natural Harbors, Navigable Rivers, and Terrain Slopes of Eurasia and Africa

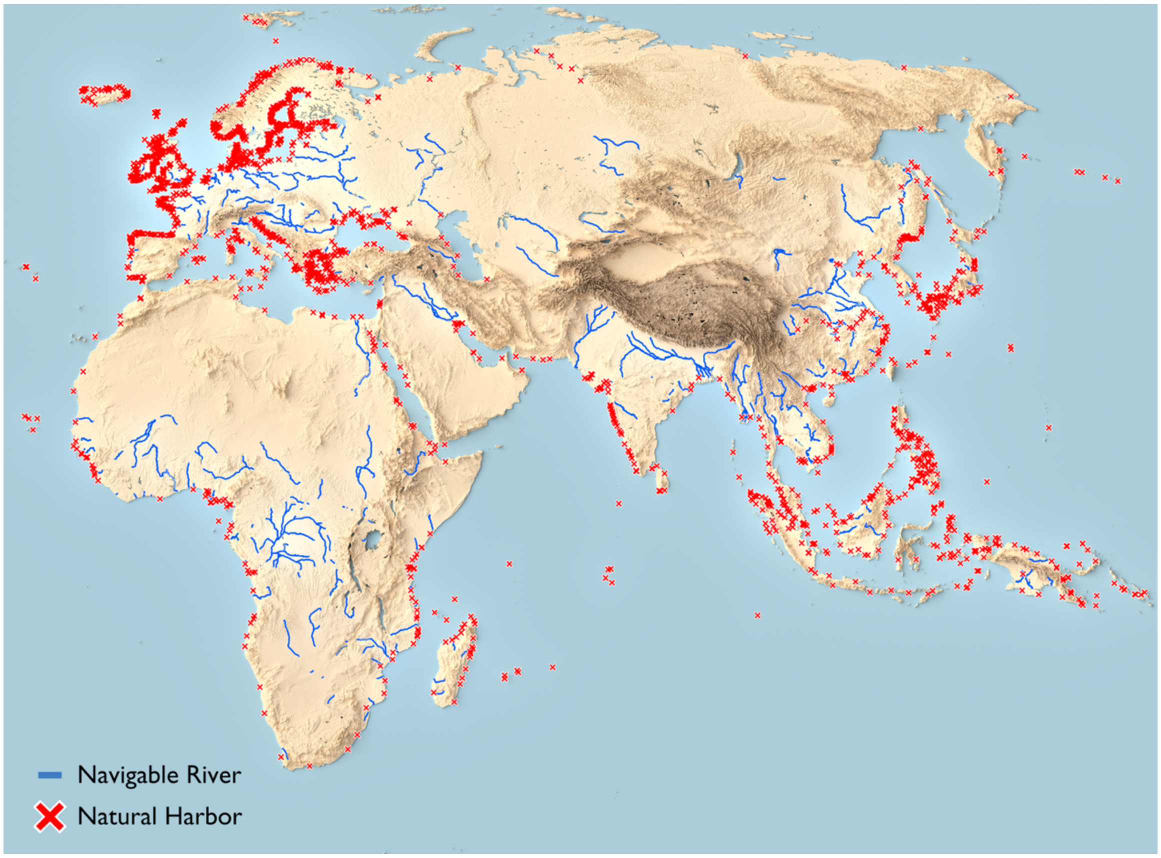

3.C. Potential Soil Productivity

We estimate the depth of each of the 41,435 markets by following the insights in Smith (1762-63: 223) regarding soil fertility: "The soil must be improveable, otherwise there can be nothing from whence they might draw that which they should work up and improve. That must be the foundation of their labour and industry." The intuition is that a potential market had a carrying capacity based on the food energy that it could produce, because the technologies that made it economically feasible to move staple foods vast distances, such as steamships and railroads, had not yet been invented (Allen 2000). We follow a long line of researchers by estimating potential agricultural production in the past by drawing on the Global Agro-Ecological Assessment for Agriculture (GAEZ), but we make a refinement to the extant estimation methods as regards the gains in productivity afforded by traditional flood irrigation in low rainfall river valleys.[6] We focus on 22 calorie dense staple crops and draw the data from version 4.0 of GAEZ.[7] To approximate the potential yields of those crops using the technologies of the 18th century, we set the GAEZ 4.0 conditions to rain-fed, low inputs.[8] We convert metric tons to calories for each of the crops based on U.S. Department of Agriculture (2016). To approximate the productive capacity of river valleys in low rainfall environments that historically benefitted from traditional flood irrigation (such as occurred along the Nile) we set the GAEZ 4.0 conditions to gravity-fed irrigation, low inputs across a 10-kilometer buffer on either side of any river (navigable or not). Our model then selects the GAEZ 4 parameters (rain fed/low inputs, or gravity-fed irrigation/low inputs) that produces the most calories.[9] Map 4 shows the results, presented at the 1km x 1km level.

---

[6] See, for example, Galor and Ozak (2016); Henderson et. al. (2018); Henderson, Storeygard, and Weil (2022); Mayshar, Moav, and Pascali (2022). For a critique of GAEZ to estimate production potential for historical periods, see Rhode (2024).

[7] We focus on wheat, sorghum, rye, wetland rice, dryland rice, pearl millet, foxtail millet, oats, maize, barley, buckwheat, peas, green gram, chickpeas, cowpeas, soybeans, groundnuts, bananas/plantains, sweet potatoes, white potatoes, cassava, and yams.

[8] Low inputs are operationalized in GAEZ 4.0 as traditional cultivars, labor intensive techniques, no application of nutrients or chemical fertilizers, no use of chemicals for pest and disease control, and minimum conservation measures.

[9] We find that rivers flowing through moderate and high rainfall environments receive no increase in production from gravity-fed irrigation. Thus, our tuning only affects arid lands located along rivers that historically relied on flood irrigation.

Map 4: Potential Staple Crop Production Using Pre-Green Revolution Methods

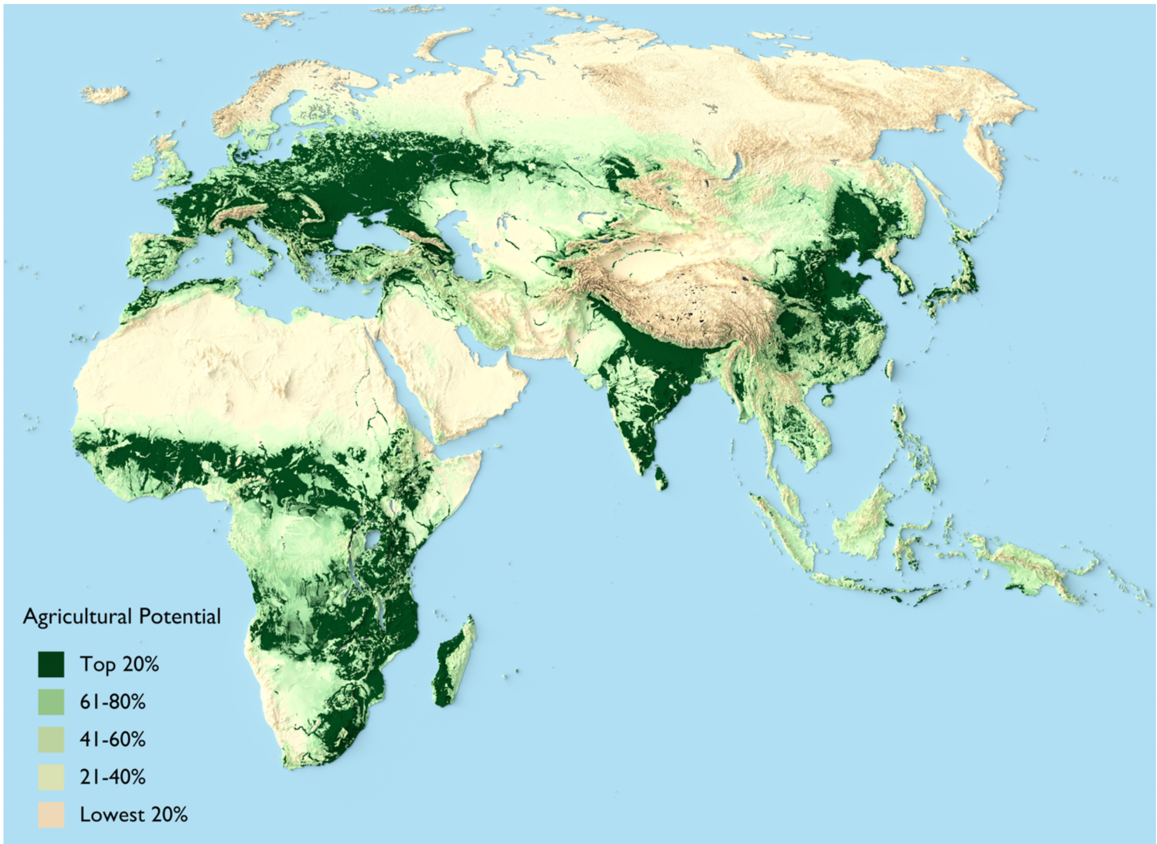

3.D. Estimating Urban Populations

We follow standard practice in the economics literature by measuring the level of economic development based on the size and density of urban populations (e.g., Acemoglu, Johnson, and Robinson, 2002; Motamed, Florax, and Masters 2014). We build a dataset that identifies, geolocates, and provides population estimates in 1700 and 1800 for any city that meets a threshold size of 20,000 inhabitants.[10] We set 20,000 as the

---

[10] Many 18th century city centers were located on the banks of rivers, lakes, and oceans, but as those cities grew their city-centers shifted inland. The present-day geo-coordinates for cities (typically fixed at their city hall or similar office) may not, therefore, accurately capture their historical proximity to navigable water. We therefore "snap" the geo-coordinates of a present-day city to the banks of the nearest historically navigable body of water, if it is within 10 kilometers of that body of water.

threshold to avoid selection bias that might result from the higher resolution of city-size estimates in Western Europe and Japan because they have received more attention by historians and archaeologists than other world regions. We construct the dataset by integrating information from online datasets compiled by other scholars, monographs by historians and archaeologists, specialist encyclopedias, historical dictionaries, and colonial gazetteers. We document the sources and methods in Supplemental Appendix A.

Section 4. Data Analysis by Market Clusters

4.1 Data Description

Each of the 41,435 markets has seven characteristics: size (in square kilometers); potential agricultural output per square kilometer (in millions of kilocalories); a count of natural harbors; a count of cities in 1800; an urban population in 1800, a count of cities in 1700; and an urban population in 1700. Table 1 presents summary statistics.

Table 1: Summary Statistics

Variable	N	Mean	SD	Min	P25	Median	P75	Max
Market Area (km^2)	41,435	34,002.5	36,268.2	30.9	19,002	20,799.5	29,117.3	267,230.8
Agricultural Potential ($Mkcal/km^2/year$)	41,435	205.2	230.4	0	0	132.1	355.3	1,525.1
Harbors (count)	41,435	7.7	31.4	0	0	0	0	359
Cities in 1700 (count)	41,435	0.5	2.1	0	0	0	0	29
Urban Population in 1700	41,435	38,343.5	179,571.2	0	0	0	0	2,331,000
Cities in 1800 (count)	41,435	0.8	3.1	0	0	0	0	45
Urban Population in 1800	41,435	58,038	246,579.5	0	0	0	0	3,322,000

Markets range from a minimum of 30.9 square kilometers to a maximum of 267,231 square kilometers, with a median size of 20,800 square kilometers. The minimum corresponds to markets drawn from initial points located on mountain slopes or located in tsetse endemic area. The maximum corresponds to markets located on broad coastal plains pierced by navigable rivers that flow unimpeded to oceans and seas. The median corresponds to markets on nearly flat terrain, without access to navigable water, and outside the zone of tsetse fly endemicity. The distributions for potential agricultural production and natural harbors are heavily skewed. More than 30 percent of markets could not produce any staple crops at all, and the most productive markets could produce more than 70 times what the median market could produce. Similarly, 78.5 percent of markets did not have access to the sea via a natural harbor, while those that had sea access frequently had access to dozens, and sometimes hundreds, of natural harbors. The distribution of urban population, whether measured by the number of inhabitants or the number of cities, is heavily skewed as well; 83.4 percent of markets had no cities and thus no urban populations at all in 1800. Supplemental Appendix B, Figures B1 and B2, show the univariate and bivariate distributions of these variables.

Figure 1 displays a scatter plot of the size of urban populations as a function of markets' size, soil fertility, and access to the sea. Market size is plotted on the X axis. Soil fertility (measured as potential staple crop production per square kilometer) is plotted on the Y axis. We divide the data into two panels: the left panel displays markets that did not have access to the sea via a natural harbor; the right panel displays markets that did have sea access. We color the datapoints based on the size of their urban populations. Blue points represent markets in which no city reached the threshold size of 20,000 inhabitants. Red points represent markets in which there was at least one city that reached the threshold size of 20,000. We then shade the points toward yellow as total urban populations increase (the sum of the populations of all cities in a market that reached the threshold size of 20,000) such that markets with urban populations greater than 1 million are displayed in gold.

We call your attention to the top right quadrant of the left panel, which represents markets of large size with fertile soils that did not have access to the sea. This is the only region of the left panel that displays a significant number of markets with urban populations. Notably, the datapoints shade toward red, indicating that these markets tended to have total urban populations at the lower end of the size distribution.

We now call your attention to the upper right quadrant of the right panel, which represents markets of large size with fertile soils and access to the sea. This is the only region of the panel that displays a significant number of markets with urban populations. Notably, the markets with the largest urban populations (those that shade toward gold) are located in the top right corner of this top right quadrant. There appears to be a non-linear interaction between market size, soil fertility, and sea access that gives rise to high levels of economic development. Supplemental Appendix B provides the univariate and bivariate distributions of these variables (Figures B1 and B2).

Figure 1: Market Size, Soil Fertility and Sea Access, Colored by Urban Population in 1800

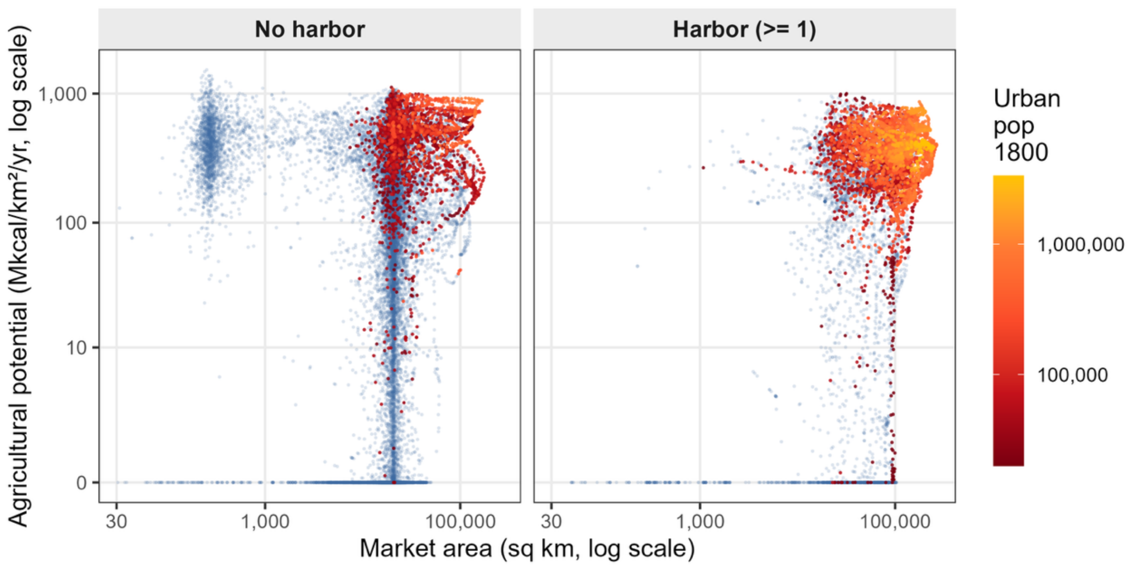

As a conceptual matter, markets do not have sharp boundaries. Our approach to drawing markets therefore allows their boundaries to overlap. Where transport costs were low and markets had large extents, the boundaries of a market could be overlapped by scores (and sometimes hundreds) of other markets. For example, the Paris market, discussed in Section 2 above, was overlapped by 98 other markets, suggesting a high degree of integration across markets. Where transport costs were high, and hence markets had very limited extents, the boundaries of a market were not overlapped by any other market. The N'shenge market, discussed in Section 2 above, provides an example. When this occurs, it suggests that that market was an isolated pocket of economic activity.

Overlapping units of analysis presents a challenge for statistical inference because of the problem of double counting. We therefore take two complementary approaches to the data. In this section, we relax the unit of analysis from the individual market to clusters of contiguous markets whose size, soil fertility, and sea access are correlated. We refer to these units as market clusters. In this approach, all city populations that fall into a market cluster are counted once, and, by design, cannot be counted in any other market cluster. This approach to the data has the advantage of approximating the phenomenon of interest to Smith (1776); the identification of the regions of the world with high potential for economic development. The disadvantage of this approach is that, because it reduces the number of observations with high-type characteristics, it only permits low powered statistical tests. In Section 5, we therefore retain the individual market as the unit of analysis but estimate random forest regressions on thousands of stratified samples constructed such that no two markets in any sample overlap one another's boundaries. In this approach, all city populations that fall into a market are treated as belonging to that market and cannot be double counted because any overlapping market is not drawn into the sample. As we show below, both approaches to the data—market clusters and random forest regressions—point to the same conclusion; 18th century economic development was concentrated in the regions predicted by Smith (1776).

4.2 Constructing Market Clusters

To generate market clusters, we start by dividing the size, soil fertility, and sea access variables into bins. We classify as having fertile soils those markets whose potential kilocalorie output per square kilometer exceeds that which we observe for regions of Eurasia and Africa that relied entirely on pastoralism as the food source (e.g., the Mongolian Plateau). The cut point is the top 30 percent of the distribution. We classify large markets as those whose size exceeds the area that could have been traversed using a horse and wagon to move one metric ton of goods on perfectly flat, nearly dry ground with our model's 40 megajoule energy budget. The implication is that large markets are those that could take advantage of at least some water-based transport. The cut point is the top 15 percent of the distribution. We classify as having sea access those markets that contained at least one natural harbor. The cut point is the top 22 percent of the distribution. This procedure places each market into one of eight categories in a 2×2×2 matrix, ranging from large markets with fertile soils and access to the sea, to small markets with infertile soils and no access to the sea.

Within each of the eight categories we then dissolve markets that touch one another into contiguous regions, such that, to a first approximation, we identify clusters composed of adjoining markets whose physical characteristics are spatially correlated. In drawing these regions, higher ranked clusters take precedence over the next highest ranked cluster. For example, if a cluster of large markets with fertile soils and sea access overlaps the boundary of a cluster of large markets with fertile soils but no sea access, the former takes precedence over the latter. This process yields a single GIS layer of market clusters tiling Eurasia and Africa, such that every raster cell and every city is uniquely assigned to a market cluster.[11]

Map 5 presents the results. The market clusters that Smith (1776) would predict to have had the highest probability of economic development (large markets with fertile soils and sea access) are shown in dark green. The North European Plain provides an example. The market clusters that Smith (1776) would predict to have had the lowest probability of economic development (small markets with infertile soils and no sea access) are shown in beige. The Sahara Desert provides an example. The market clusters with intermediate probabilities of economic development are shown in shades of green, with higher probabilities appearing as progressively darker shades of green. For example, the market clusters that Smith (1776) would have given the second highest probability for economic development (large markets with fertile soils, but no sea access) are shown in medium green. Central Europe provides an example.

[11] We merge clusters within the same category whose boundaries are within five kilometers of one another on the assumption that such proximity reflects functional integration. The dissolve also produces six unassigned boundary slivers of less than 150 km² which we disregard.

Map 5: The Market Clusters of Eurasia and Africa

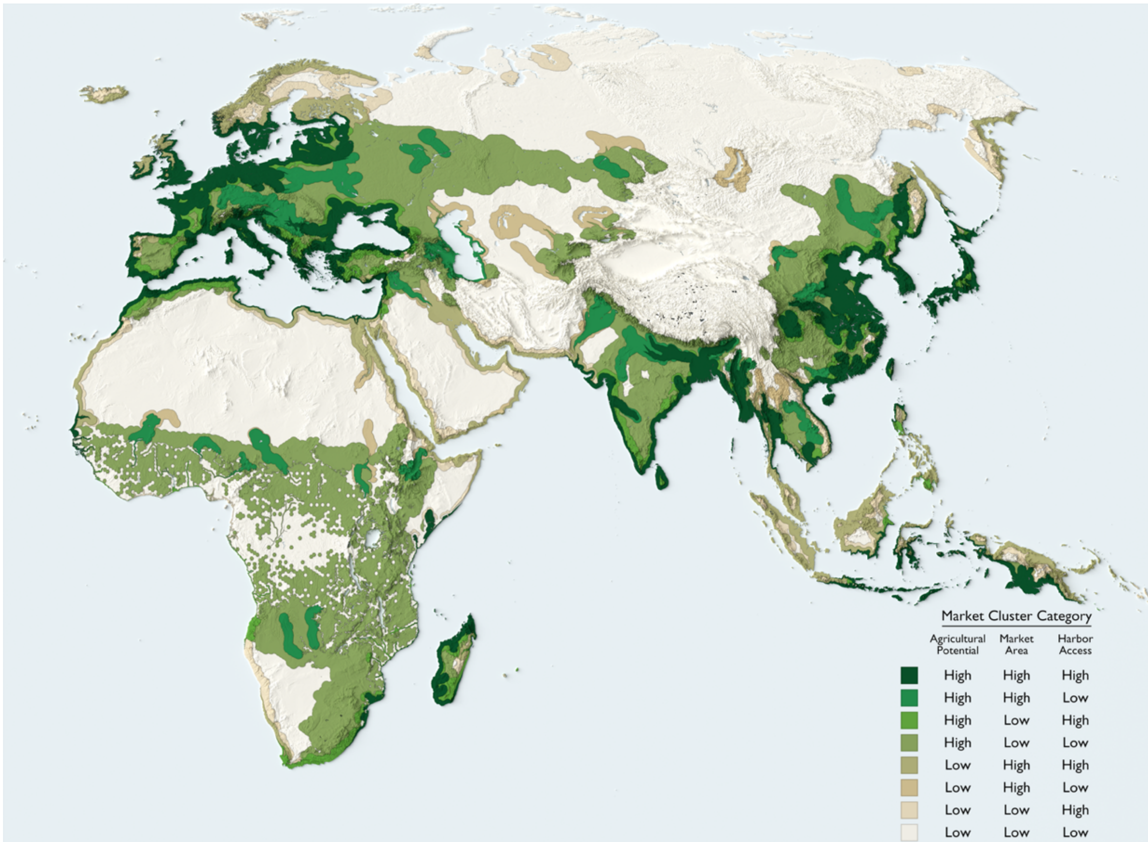

Map 5 reveals a striking pattern. There were basically two immense market clusters (shown in dark green) which Smith (1776) would have predicted to be the most economically developed regions of the world in the 18th century. One of those stretched across Europe, from the Baltic Rim, across the North European Plain and England, southwards along the coast of France, Spain, and Portugal, and then along the Mediterranean coasts of Southern Europe and North Africa, ending on the coasts of the Black Sea. The second stretched across East and South Asia, from the main island of Japan in the northeast to the western coastline of the Indian subcontinent. Along the way, it took in the North China Plain, Taiwan, the Pearl River Delta, and the river valleys of the lower and middle Yangtze, Mekong, Salween, Irrawaddy, Brahmaputra, and Ganges rivers. Because these clusters were composed of markets in the top 15 percent of the size distribution, each of the markets overlapped the boundaries of numerous adjoining markets in its cluster. The average market in the European cluster contained the initial points of 109 other markets in that cluster, while the average market in the Asian cluster contained the initial points of 74 other markets in that cluster, indicating that these market clusters were composed of highly integrated markets.

Map 6 adds cities of at least 20,000 inhabitants in 1800 to Map 5, showing them as gold circles. Larger cities are represented by progressively larger circles, such that the largest cities (those with at least 500,000 inhabitants) are shown as five times the size of the smallest cities. The map shows that economic development was concentrated in the regions predicted by Smith (1776); the highly integrated, European and Asian market clusters composed of large markets with fertile soils and sea access displayed in Map 5. To the degree that there are concentrations of cities located outside these two clusters, they are in market clusters composed of large markets with fertile soils that are just inland—usually upriver—from the clusters composed of large markets, with fertile soils and sea access. These are shown in Maps 5 and 6 as medium green. The upper Mekong, upper Ganges, and Indus river valleys provide examples.

Note that much of Africa falls in the category of market clusters composed of small markets with infertile soils and no sea access. There are no cities in those market clusters. To the degree that Map 6 shows regions with concentrations of cities, there are two. The first is the interior portions of West Africa that display pockets of market clusters composed of large markets with fertile soils, but without sea access. These are the consequence of the navigable portions of the Niger and Volta rivers with fertile soils that are north of the zone of tsetse fly endemicity. The second is the lower stretches of those same rivers near the Bight of Benin, where the soils are fertile, but markets were small because tsetse fly transmitted trypanosomiasis precluded the use of draft animals in overland transport once goods were offloaded at riverbanks.

Map 6: Market Clusters Overlaid with Cities in 1800

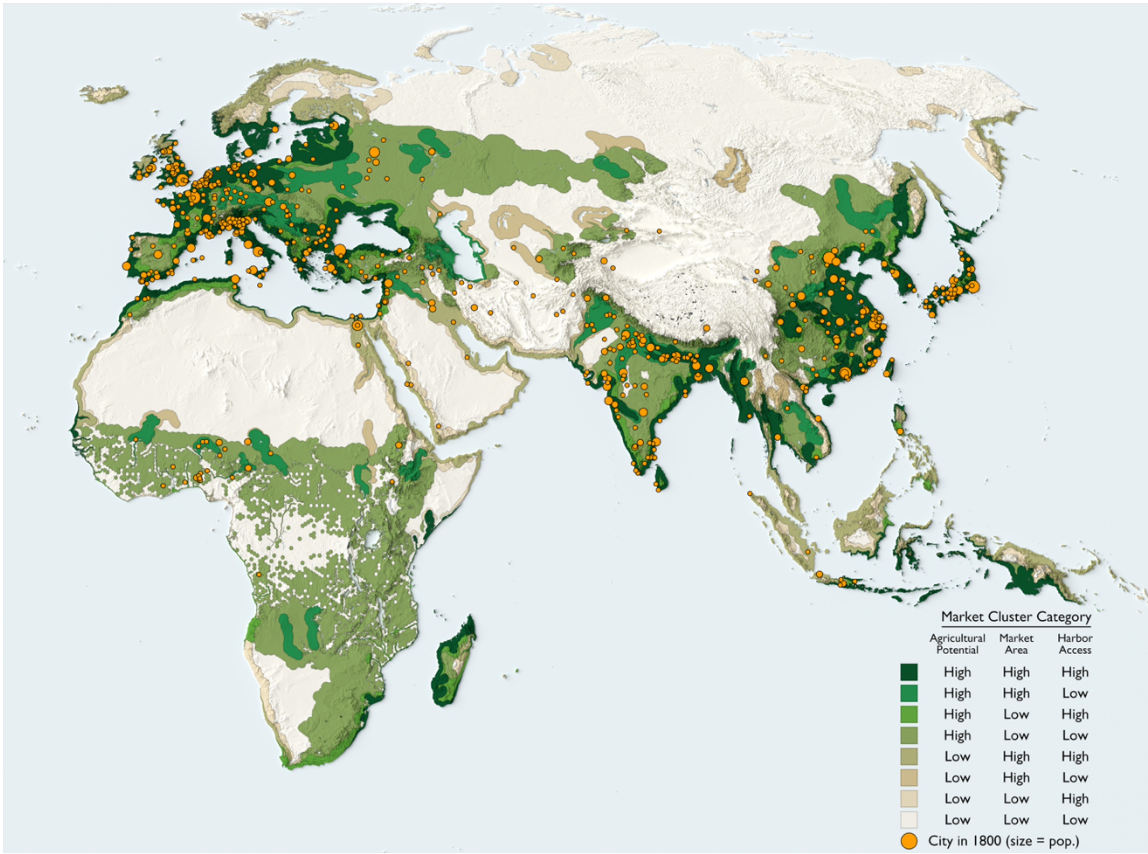

Table 2 aggregates the market clusters displayed in Map 6 by their size-fertility-sea access category, allowing us to measure the degree to which a category of market clusters exhibited a level of economic development (as measured by its urban population) above, below, or equal to the percent of the Earth's terrestrial surface that it covers. A value of 1 in the resulting location quotient indicates that a cluster category's level of economic development is as spatially concentrated as that of Eurasia and Africa as a whole. Higher values indicate greater concentration of economic development. Lower values indicate a lower concentration of economic development. The top row of Table 2, for example, aggregates all market clusters composed of large markets with fertile soils and sea access (the dark green regions of Map 6). It shows that this category of market clusters represented only 10.2 percent of the land area of Eurasia and Africa but contained 71.0 percent of its urban population, yielding a location quotient of 6.96. That is, the regions of the world that Smith (1776) predicts would be the most economically developed were seven times wealthier than would be expected if economic development were randomly distributed. Importantly, consistent with the predictions in Smith (1776), as we move down the rows in Table 2, to the categories of market clusters that include only two of the three high type factors, to the categories that include only one of them, to the category that includes none of them, location quotients tend to fall exponentially.

Table 2: Market Cluster Characteristics Predict Economic Development in 1800

Fertile Soils	Large Market	Sea Access	% of Eurasian and African Land	No. Potential Markets	No. Cities 1800	Urban Pop. (thousands)	% of Eurasian and African Urban Pop.	Location Quotient
Yes	Yes	Yes	10.2	2,255	333	22,599	71.0	6.96
Yes	Yes	No	4.5	601	54	2,633	8.3	1.84
Yes	No	Yes	3.0	1,291	32	1,120	3.5	1.17
No	Yes	Yes	6.2	2,772	17	996	3.1	0.50
Yes	No	No	23.4	8,284	82	3,553	11.2	0.48
No	Yes	No	2.0	588	2	43	0.1	0.07
No	No	Yes	3.8	2,425	2	58	0.2	0.05
No	No	No	46.8	23,219	24	841	2.6	0.06
Total			100.0	41,435	546	31,842	100.0	1.00

4.3 Sensitivity Analyses

One potential concern is that perhaps the results displayed in Table 2 are the product of a single region, such as Europe. Map 6 suggests that is not the case, but to address this concern more fully, we divide Map 6 into a 3 by 3 grid and redo the analysis in Table 2 by serially dropping one of the nine gridded regions. We display the grid of regions, along with the sensitivity analysis tables and figures, in Supplemental Appendix C, Figure C1 and Table C1. They show that after dropping each gridded region sequentially the relationship between the level of economic development and the category of market clusters composed of large markets with fertile soils, and sea access continues to hold. Of note is that when the Northwest (European) block is dropped, the category of market clusters made up of large potential markets with fertile soils and sea access yields a location quotient of 8.83, which is higher than that in the base model (6.96).

Another potential concern is that perhaps a single variable explains the results in Table 2, rather than the interaction of the three variables hypothesized in Smith (1776). We therefore repeat the analysis presented in Table 2, but instead of interacting the three factors we generate single-factor clusters and two-factor clusters (using the same cut points from Section 4.2, above) and compare their location quotients to the location quotient of the three-factor region from our base model (6.96, per Table 2). We show the results in Supplemental Appendix C, Table C2. None of the three variables by themselves generates a location quotient as large as the location quotient of the category of the three interacted variables. The bivariate interactions also fail to produce a location quotient as large as the three-way interaction.

Finally, perhaps the results displayed in Table 2 are the product of tuning our geospatial model to an energy budget of 40 megajoules. We therefore substitute energy budgets from 30 to 50 megajoules in two megajoule increments and rerun the analyses. We report the results in Supplemental Appendix C, Table C3. We find that our base results are not sensitive to this retuning. The category of market clusters composed of large markets with fertile soils and sea access has location quotients ranging from 6.16 to 8.6, indicating that the regions hypothesized by Smith (1776) were roughly six to eight times wealthier than would have been the case had economic development been randomly distributed.

4.4 Robustness Tests

Perhaps the results displayed in Table 2 are an artifact of estimating location quotients using urban populations measured circa 1800, roughly two decades after the publication of Smith (1776)? Because the studies by historians and archaeologists that we rely upon to build our cities dataset tend to estimate city populations at 50 or 100-year intervals, it is not feasible to present results that focus precisely on 1776. We can, however, re-run the analyses presented in Table 2 (and Supplemental Appendix C, Table C1) substituting urban populations circa 1700. We present the results in Supplemental Appendix C, Table C4. They show that our results for 1800 are robust to the substitution of circa 1700 urban population data. In fact, the location quotients for the category of market clusters composed of large markets with fertile soils and sea access tend to be larger in 1700 than in 1800.

Section 5. Random Forest Analysis

Drawing clusters of markets allows us to test the predictions of Smith (1776) by identifying large-scale spatial trends. It only permits, however, low-powered statistical tests: binning continuous variables at arbitrary cut points and dissolving overlapping markets into contiguous clusters leaves too few observations in the top categories for reliable inference. We therefore estimate random forest regressions on stratified samples of non-overlapping potential markets, preserving the full continuous variation across all 41,435 markets. Random forests are a predictive tool, widely adopted in the social sciences for their ability to recover non-linear interactions without imposing a functional form on the data (Efron and Hastie 2016; Mullainathan and Spiess 2017; Athey and Imbens 2019). We follow the standard random forest workflow at every step where it applies—tree ensembling, cross-validated tuning of standard hyperparameters, held-out evaluation, impurity-based variable importance, partial dependence plots, and permutation-based significance testing—and depart from it only where the overlapping structure of our units of analysis renders the textbook procedure of randomized bootstrapping with replacement invalid.

A textbook random forest generates diversity across its trees by bagging; each tree is fit to a bootstrap resample of the rows, drawn with replacement (Breiman 1996, 2001). Bagging assumes the rows are independent, which our units of analysis are not. Market boundaries overlap in space, so a bootstrap draw would routinely place a market and its overlapping neighbor in the same tree, double counting the cities they share. We therefore replace the bootstrap with a source of inter-tree diversity valid for overlapping units; each tree is fit to a stratified draw of non-overlapping markets, with diversity coming from spatial variation across draws rather than from resampling within them. This preserves what bagging provides—an ensemble of decorrelated, high-variance trees whose average is a low-variance predictor—while respecting the structure of our data.

We employ stratified samples because uniform random selection of non-overlapping markets would bias heavily toward smaller markets. Each time a market is drawn, it removes from the candidate pool every other market that overlaps it; large markets overlap with more candidates and are correspondingly harder to fit into a sample without conflict. The resulting ensemble of decision trees would then be drawn from samples that under-represent the large markets of interest to Smith (1776).

We therefore implement a two-phase stratified priority selection process. We divide markets into two strata by area; those in the top 15 percent of the distribution (per Section 4.2, above), and those in the rest of the distribution. The algorithm draws markets in the top stratum randomly but rejects those whose boundaries overlap the boundary of any market that has already been drawn into the sample. Once no additional top stratum markets can fit into the sample without overlap, the algorithm draws from the bottom strata, picking at random, but rejecting those whose boundaries overlap the boundary of any market that has already been drawn into the sample. This selection process increases the probability that each sample distribution reflects the population distribution of market sizes. Map 7 displays an example of what one draw looks like.[12]

[12] As a robustness check, we draw a parallel set of samples using area-weighted selection, in which markets larger than the modal area receive a probability of inclusion that is proportional to market area, while markets smaller than the mode are fixed at the modal area. We note that the two methods yield nearly identical results. We present the results in Supplemental Appendix D, Table D1.

Map 7: A Single Sample of 400 Non-Overlapping Markets

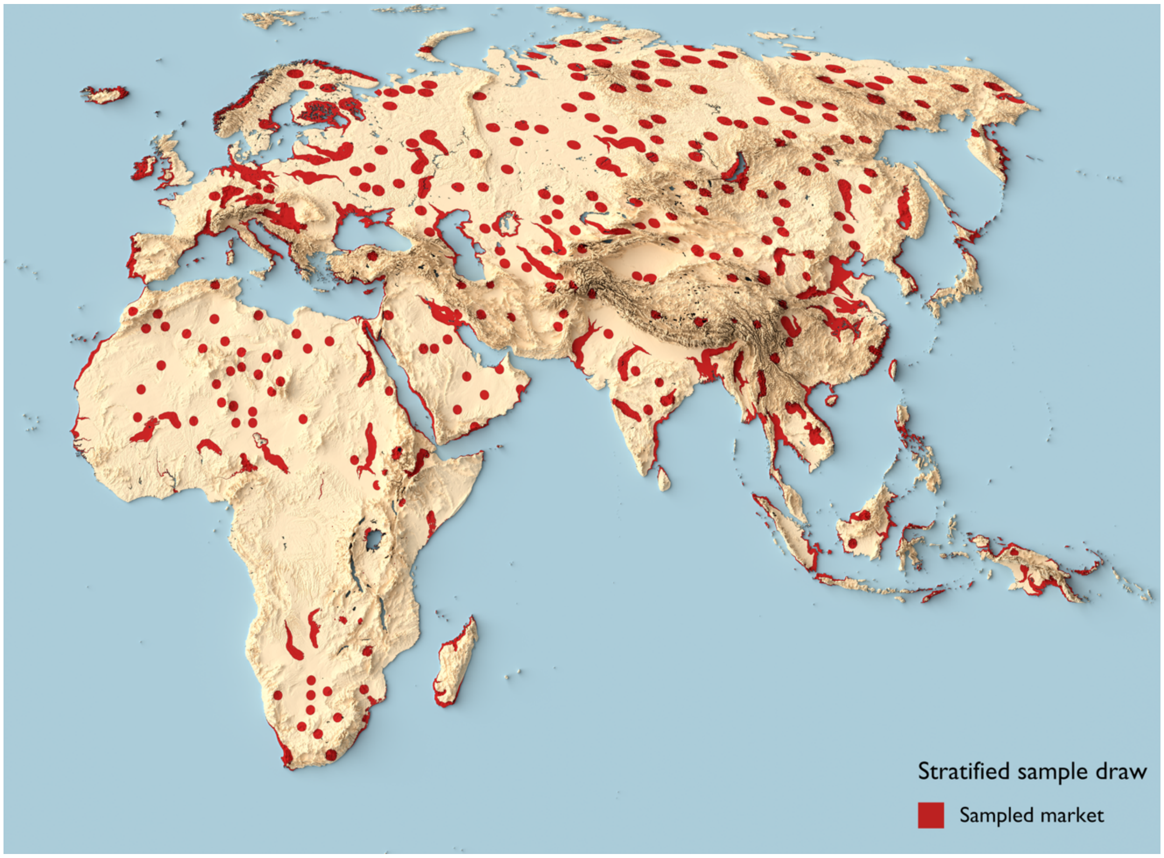

We draw 6,000 samples, each of which contains 400 markets.[13] We log the data to minimize the impact of outliers (Supplemental Appendix D, Table D3). We split the samples into 4,800 training samples and 1,200 holdout samples that we reserve for out-of-sample evaluation. We select hyperparameters through five-fold cross-validation (Arlot and Celisse 2010), tuning the two hyperparameters conventionally tuned in a

[13] We select 400 markets per sample based on experiments with sample sizes of 350, 400, and 450. We reran the full pipeline—hyperparameter tuning, 4,800-tree ensemble, and 1,200-draw holdout evaluation—for each of the three sample sizes. We found that the mean R^2 of the predicted values from the 4,800 ensemble versus the 1,200 holdout sample improves from a sample size of 350 to 400, and then declines modestly to 450. This suggests that 400 markets is the appropriate balance point between sample size and composition. We present the results of in Supplemental Appendix D, Table D2.

random forest: the number of predictors per split and the minimum leaf size (Probst, Wright, and Boulesteix 2019). We report the tuning results in Supplemental Appendix D, Figure D1. We iterate over six candidate configurations, with maximum tree depth fixed at seven (a non-binding constraint given three predictors and approximately 400 observations per tree, since the natural stopping depth is $\log_2(400/\text{min\_node\_size}) \approx 4–5$). Each sample represents a different spatial constellation of non-overlapping markets and therefore generates a unique decision tree. Diversity within the ensemble of decision trees comes from spatial variation across samples; different constellations of markets produce different trees. The 1,200 hold-out samples are each evaluated against the full 4,800-tree ensemble. For each holdout sample, the ensemble computes a RMSE and R^2 based on the values predicted by the 4,800 training samples against the actual values of the dependent variable in that sample. The reported holdout RMSE and R^2 are the means across all 1,200 holdout samples (Supplemental Appendix D, Figure D2). If the holdout RMSE and R^2 are close to the training RMSE and R^2 the implication is that the model is not overfitting to the specific spatial configurations in the training samples.

We evaluate spatial generalization through leave-one-block-out (LOBO) cross-validation. We rerun the full pipeline nine times, each time dropping every market whose initial point falls within one block of the 3×3 spatial grid that we drew in Section 4.3. In each iteration we train 4,800 trees on the markets in the eight remaining blocks and evaluate them on the 1,200 holdout samples, also filtered to exclude markets whose initial points fall in the dropped block. Jackknife standard errors across the nine blocks provide 95 percent confidence intervals for holdout RMSE and R^2. Holdout performance is stable across LOBO iterations, and the full-sample estimates fall within the jackknife confidence intervals, indicating that the results are not driven by any single geographic region. This follows the spatial block jackknife methodology in Künsch (1989) and the spatial cross-validation recommendations of Roberts et. al. (2017).[14]

[14] The nine LOBO analyses share 8/9ths of the data, so their performance estimates are correlated and a naive standard deviation across them understates the sampling

We conduct a permutation test, following Ojala and Garriga (2010), to check whether the observed R^2 from the random forest could arise by chance—that is, whether the random forest can manufacture apparent predictive power from our predictors even when they bear no true relationship to the outcome. We randomly shuffle the values of the dependent variable across observations, breaking any genuine link between predictors and outcome while leaving the predictors and all marginal distributions unchanged. We then rerun the pipeline on the shuffled data, using a 480-tree training ensemble and 120-draw holdout—a 10 percent down-sample of the main pipeline that is computationally feasible to repeat many times. We repeat this 100 times to estimate a null distribution of holdout R^2 (Supplemental Appendix D, Table D4). The permuted ensembles produce holdout R^2 clustering near zero (mean 0.010, maximum 0.047), far below the observed 0.518 ($p < 0.01$), indicating that the model captures a genuine predictive relationship rather than an artifact of the model's flexibility. The distinct concern that spatial autocorrelation could inflate this relationship is addressed separately through the leave-one-block-out (LOBO) analysis.

5.1 Base Results

Table 3 presents the results of the random forest regressions. The ensemble accounts for 52 percent of the variance in logged urban population in 1800. Cross-validated performance on the training draws ($R^2 = 0.518$) and performance on the 1,200 held-out draws ($R^2 = 0.518$) are nearly identical, indicating that the ensemble is not overfitting to the particular spatial constellations on which it was trained. All three of the factors hypothesized by Smith (1776) contribute to the model's predictions: market size accounts for 66 percent of the mean decrease in impurity across the ensemble, soil fertility 26 percent, and access to the sea 8 percent (Table 3, Panel C).

As a benchmark, we estimate a linear regression on the same three logged predictors, using the same training draws and the same holdout evaluation protocol as the

---

variability. The jackknife formula applies an (n−1) scaling factor to correct for this overlap.

random forest. The linear model achieves a holdout R^2 of 0.35, against the forest's 0.52; the gap measures the contribution of the threshold effects and interactions visible in Figure 1, which a linear additive specification does not capture.

Table 3: Random Forest Ensemble Results: Economic Development in 1800

Panel A: Ensemble Specification		
Dependent variable	$\ln(1 + \text{population in 1800})$	
Predictors	$\ln(1 + \text{area})$; $\ln(1 + \text{caloric productivity})$; $\ln(1 + \text{harbors})$	
Potential markets	41,435	
Training / holdout draws	4,800 / 1,200 (400 markets per draw)	
Hyperparameters	`mtry`=3, min. node size=15, max depth=7	
Panel B: Ensemble Performance	R^2	RMSE
Cross-validated, training draws	0.518 (0.003)	2.790 (0.014)
Held-out draws	0.518 (0.037)	2.790 (0.145)
OLS benchmark, held-out draws	0.346 (0.026)	3.251 (0.101)
Permutation null (100 shuffles of DV)	0.010 (max 0.047)	3.999
Panel C: Variable Importance (share of impurity reduction)		
Market area (log)	65.6%	
Caloric productivity (log)	26.5%	
Natural harbors (log)	7.9%	
Panel D: Spatial (LOBO) Cross-Validation		
Holdout R^2, range across nine dropped blocks	0.451 – 0.549	
Holdout R^2, Northwest block dropped	0.451 (0.042)	
Jackknife 95% CI for holdout R^2	[0.33, 0.71]	

The partial dependence plots shown in Figure 2 trace the isolated effect of each predictor while averaging over the other two. All three effects are positive and monotonic. Market size and soil fertility only affect economic development, however, at the right tails of the distributions; beyond the 81st percentile of the distribution by size and beyond the 60th percentile for soil fertility. The random forest, in other words, recovers from continuous data nearly the same cut points that we imposed a priori in the market cluster analysis of Section 4.2 where "large" meant the top 15 percent of area and "fertility" meant the top 30 percent of potential agricultural productivity. Sea access, however, operates in a more linear manner; predicted economic development rises from the very first natural harbor and continues to climb, with diminishing returns, across the full range.

We test whether the model's performance depends on any single region of Eurasia and Africa—in particular, whether the results are simply a description of Europe. We conduct a Leave One Block Out Analysis using the same 3x3 grid as in Section 4.3. We summarize the results in Panel D of Table 3 and report them iteration by iteration in Supplemental Appendix D, Table D5. When we drop all markets in the northwest (Europe) block and retrain the full ensemble on the remaining eight blocks, holdout R^2 falls only modestly, from 0.518 to 0.451. Across all nine iterations, holdout R^2 ranges from 0.451 to 0.549, the full-sample estimate falls within the jackknife 95 percent confidence interval [0.33, 0.71], and the ranking of variable importance—area first, productivity second, harbors third—is unchanged. In short, we find that the relationship between the three factors specified by Smith (1776) and the level of economic development is not the product of any one region.

Figure 2: Partial Dependence Plots

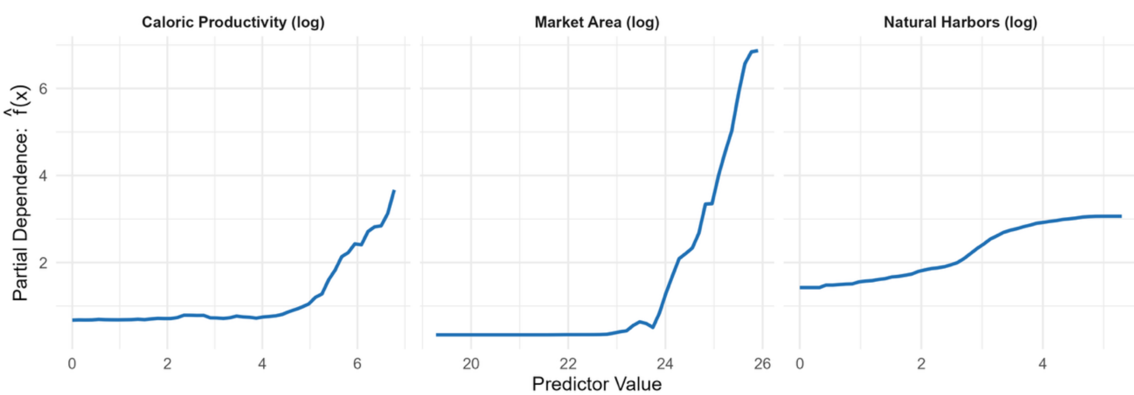

For each of the 41,435 markets, the ensemble generates a prediction by passing the market's three characteristics—market size, potential agricultural productivity, and natural harbors—down each of the 4,800 trees, averaging the 4,800 resulting estimates of logged urban population. Ranking all markets by these predicted values, Map 8 plots the markets in the top 5 percent of the predicted distribution in dark purple and the next 10 percent in light purple. Cities in 1800 are plotted the same way as in Map 6, allowing readers to see the regions in which the model does and does not successfully predict economic development. Map 8 shows that the model does remarkably well in predicting the regions with high levels of development in 1800: cities (especially large cities) cluster in the dark purple (top 5 percent) regions. There are also cities in the light purple regions (the next 10 percent) but they are less densely packed and tend to have smaller populations than in the dark purple region. Regions that the model predicts to be in the top five percent of the distribution of economic development, but that lacked cities in 1800, include the coast of Norway, the Baltic coasts of Finland and Estonia, the north coast of the Black Sea, the Philippine archipelago outside of the main island of Luzon, the Indonesian island of Sulawesi, and the south coast of Papua New Guinea. There are two regions in which the model fails to predict economic development that was present in 1800. The first is a set of cities in Western Russia. Examination of their histories indicates that they were fortress towns built to repel the Golden Horde, only one of was sizable, Moscow. The second are the isolated cities spread across western China, Tajikistan, Uzbekistan, Afghanistan, Iran, and Turkey. Examination of their histories indicates that, with only a few exceptions, these were Silk Road trading hubs; they developed because of commerce, but it was trade in silk and tea moving overland from China to Europe before the discovery of a sea route from Europe to Asia in 1498.

Map 8: Predicted Versus Actual Levels of Economic Development, Circa 1800

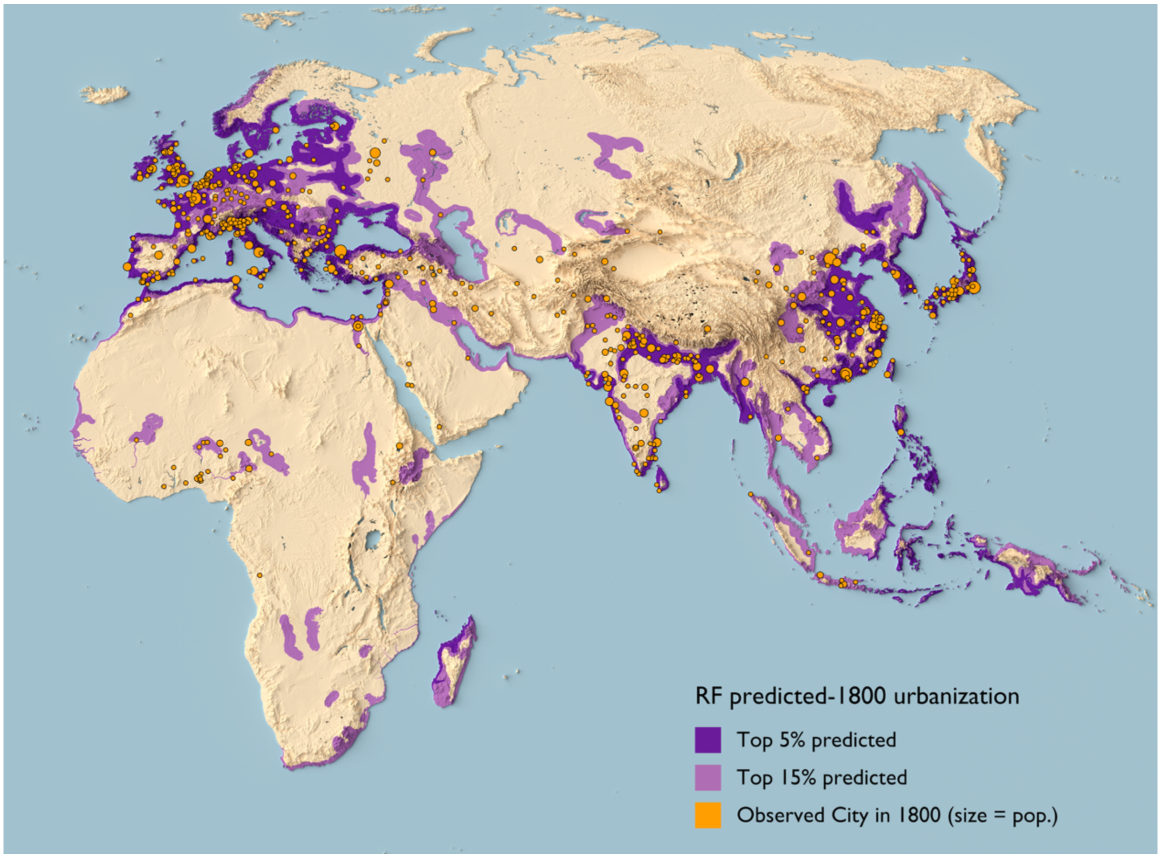

5.2 Robustness Test

As we did in Section 4.4, above, we re-run our analysis by substituting urban populations circa 1700. We present the results in Supplemental Appendix D, Table D6 and Figures D3 and D4. They show that our base results for 1800 are robust to the substitution of circa 1700 urban population data. The random forest results are likewise insensitive to changes in the energy budget (Supplemental Appendix D, Table D7).

Section 6. Conclusion

One of the core concerns of economics is how human beings, operating through markets, allocate scarce resources across competing ends. The econometric analyses of the resulting phenomena tend not, however, to employ markets as the unit of analysis. Rather, they employ administrative jurisdictions, which are based on political criteria, or, increasingly, on latitude-longitude grid cells, which do not have a theoretical basis. The result is a lack of congruence between a phenomenon and the units of analysis to study it. Plainly put, market extents vary over space and time, but grid cells are of uniform size and are fixed in time. This generates multiple threats to inference, including biased standard errors, the need to introduce covariates to control for spatial correlation, and undetected spatial trends.

We address these limitations of grid cells as units of analysis by developing a method to measure markets as contiguous, technology-conditioned cost surfaces. Our method estimates market extents based on the features of the earth's surface at a point in time, the state of technology at that point in time, and an energy budget that has verisimilitude for that point in time. We also develop novel methods to analyze the resulting geospatial data: GIS layers of contiguous markets whose physical characteristics are correlated, such that they constitute market clusters; and random forest regressions on stratified samples of non-overlapping markets.

We demonstrate the utility of these methods of measurement and analysis by testing an idea that is central to economics, but that has not before been properly tested: Adam Smith's theory that economic development is bounded by the extent of the market, and the extent of the market is bounded by physical features of the Earth's surface. When we analyze the data by generating market clusters, we find that 10 percent of the land mass of Eurasia and Africa accounted for 71 percent of economic development (as measured by urban populations). When we analyze the data by estimating random forest regressions, we find that the factors hypothesized by Smith (1776) account for just over half of the distribution of economic development (as measured by urban populations).

Our results indicate that the methods of measurement and analysis we offer have applications beyond the distribution of economic development in the 18th century. The extent of the market is central to a broad range of human decision making; from marriage transitions to investments in human capital, from the location of production facilities to investments in political and legal institutions. Importantly, the methods we introduce allow researchers to measure market extents—and the phenomena to which they give rise—over time, because as technologies change, market extents change. To point to an example, the development of the railroad allowed market extents to expand because it consumed less energy per ton-kilometer than horse-drawn wagons and because it increased energy budgets by harnessing a fossil fuel energy source. Research to date on the impact of railroads has focused on estimating social savings at the national level (e.g., Fogel 1964, Summerhill 1997) or on estimating the effects of railroads at the level of localities or economic sectors within a country (e.g., Donaldson and Hornbeck 2016, Donaldson 2018). An approach based on market extents allows researchers to ask a different question: how did the differential diffusion of railroad technology affect the distribution of market extents around the planet and the resulting phenomena that emerged from it? The market extent approach would allow similar questions to be posed about the differential effects of the diffusion of other transport technologies.

The all-in-all objective is to develop and test an account of the conditions under which technological change itself occurs. Sokoloff (1988), for example, demonstrated that canals were corridors of innovative activity. We believe that the methods of measurement and analysis we introduce offer a general framework to establish how the distribution of market extents, conditional on a state of technology at a point in time and the features of the Earth's surface at that point in time, affected subsequent technological innovation.

References

Acemoglu, Daron, Simon Johnson, and James Robinson. 2002. "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution," The Quarterly Journal of Economics 117: 1231-94.

Allen, Robert, 2000, "Economic Structure and Agricultural Productivity in Europe, 1300-1800," Economic History Review 4: 1-26.

Alsan, Marcella. 2015. "The Effect of the TseTse Fly on African Development." American Economic Review 105: 382–410.

Arlot, S., and A. Celisse. 2010. "A Survey of Cross-Validation Procedures for Model Selection." Statistics Surveys 4: 40–79.

Athey, S., and G. Imbens. 2019. "Machine Learning Methods That Economists Should Know About." Annual Review of Economics 11: 685–725.

Baker, Ira Osborn, 1918, A Treatise on Roads and Pavements, 3rd Edition (John Wiley and Sons).

Boss, Richard. 1993. "Keelboat, Pirogue, and Canoe: Vessels Used by the Lewis and Clark Corps of Discovery." Nautical Research Journal 1:27.

Breiman, L. 1996. "Bagging Predictors." Machine Learning 24: 123–140.

Breiman, L. 2001. "Random Forests." Machine Learning 45: 5–32.

Conley, Timothy, and Morgan Kelly. 2025. "The Standard Errors of Persistence." Journal of International Economics 153: 1-22.

Deasy, George. 1942. "The Harbors of Africa." Economic Geography 18: 325-42.

Donaldson, David. 2018. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." The American Economic Review 108: 899-934.

Donaldson, David, and Richard Hornbeck. 2016. "Railroads and Economic Development: A Market Access Approach." The Quarterly Journal of Economics 131:799–858.

Easterly, William and Ross Levine. 2003. "Tropics, Germs, and Crops: How Endowments Influence Economic Development." Journal of Monetary Economics 50: 3-40.

Efron, Bradley, and Trevor Hastie. 2016. Computer Age Statistical Inference: Algorithms, Evidence, and Data Science (Cambridge University Press).

Fogel, Robert. 1964. Railroads and American Economic Growth: Essays in Econometric History (Johns Hopkins University Press).

Food and Agriculture Organization of the United Nations. 2016. Global Agro-Ecological Zones, Accessed February 1, 2016. (http://gaez.fao.org/Main.html?ticket=ST-14587-gObOUJ3M5AJwqmJTXs4w-cas#)

Friedman, J. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." Annals of Statistics 29: 1189–1232.

Galor, Oded and Omer Ozak, 2016, “The Agricultural Origins of Time Preference,” American Economic Review 106: 3064-3103.

Henderson, J. Vernon, Tim Squires, Adam Storeygard, and David Weil. 2018. “The Global Distribution of Economic Activity: Nature, History, and the Role of Trade.” The Quarterly Journal of Economics 133: 357–406.

Henderson, J. Vernon, Adam Storeygard, and David Weil. 2022. “Land Quality.” https://bpb-us-w2.wpmucdn.com/sites.brown.edu/dist/1/24/files/2022/08/Land-Quality-August-2022.pdf accessed May 12, 2026.

Masters, William, and Margaret McMillan. 2001. “Climate and Scale in Economic Growth.” Journal of Economic Growth 6: 167-86.

Mayshar, Joram, Omer Moav, and Luigi Pascali, 2022, “The Origin of the State: Land Productivity or Appropriability?” Journal of Political Economy 130: 1091-1144.

Motamed, Mesbah, Raymond Florax, and William A. Masters. 2014. “Agriculture, Transportation and the Timing of Urbanization: Global Analysis at the Grid Cell Level.” Journal of Economic Growth 19: 339–36.

Mullainathan, S., and J. Spiess. 2017. "Machine Learning: An Applied Econometric Approach." Journal of Economic Perspectives 31: 87–106.

Mussulman, Joseph. ND-A. “The Pirogues.” Retrieved August 26, 2022 from https://lewis-clark.org/boats/pirogues/

Mussulman, Joseph. ND-B. “The Barge: Barge, Keelboat, or Just 'the boat'?” Retrieved August 26, 2022 from http://lewis-clark.org/boats/the-barge/

Nordhaus, William. 2006. “Geography and Macroeconomics: New Data and New Findings.” PNAS 103: 3510-17.

Ojala, M., and G. Garriga. 2010. "Permutation Tests for Studying Classifier Performance." Journal of Machine Learning Research 11: 1833–1863.

Probst, P., M. Wright, and A.-L. Boulesteix. 2019. "Hyperparameters and Tuning Strategies for Random Forest." WIREs Data Mining and Knowledge Discovery 9:e1301.

Rhode, Paul, 2024, "How Suitable are FAO-GAEZ Crop Suitability Indices for Historical Analysis?"

Rossi-Hansberg, Esteban, and Jialing Zhang, 2026, "Local GDP Estimates Around the World," Journal of Urban Economics 154: 1-11.

Shumway, George, Edward Durell, and Howard Frey, 1964, Conestoga Wagon 1150-1850: Freight Carrier for 100 Years of America's Westward Expansion (Early American Industries Association).

Smith, Adam, 1762-1763. Lectures on Jurisprudence

Smith, Adam. 1776 [1976]. An Inquiry into the Nature and Causes of the Wealth of Nations (University of Chicago Press).

Strobl, C., A.-L. Boulesteix, A. Zeileis, and T. Hothorn. 2007. "Bias in Random Forest Variable Importance Measures." BMC Bioinformatics 8:25.

Summerhill, William. 2003. Order Against Progress: Government, Foreign Investment, and Railroads in Brazil, 1854-1913 (Stanford University Press)

Vanhatalo, Jarno, Juri Huuhtanen, Martin Bergstrom, Inari Helle, Jussi Makinen, Pentti Kujala. 2021. "Probability of a ship becoming beset in ice along the Northern Sea Route –A Bayesian analysis of real-life data." Cold Regions Science and Technology 184.

Vansina, Jan. 1978. The Children of Woot: A History of the Kuba Peoples (University of Wisconsin Press).

United States, Department of Agriculture. 2016. National Nutrient Database for Standard Reference: Release 28. (United States Department of Agriculture, Agriculture Research Service).

United States, Department of the Navy, 1953. World Port Index. U.S. Government Printing Office.

Supplemental Appendices for *Geography, Technology, and the Extent of Markets*

Roy Elis, Stephen Haber, and Jordan Horrillo

Contents

Appendix A — Sources to Estimate Urban Populations Circa 1700 and 1800

We give priority to scholarly monographs that focus on individual cities and to the entries for individual cities in specialist encyclopedias. We then prioritize monographs that focus on specific countries. We then focus on scholarly monographs that provide coverage for specific regions in specific time periods. We then turn to studies by historians with broad, multi-region coverage. We then give priority to sources that provide world-wide coverage, many of which have are available online.

To facilitate replicability, and to aid other scholars in building upon this dataset, we adopt the protocol that each city-year observation is drawn from a single source that we document in our online Excel workbook.

Science is a cumulative enterprise. There are, without doubt, studies by archaeologists and historians that contain city data of which we are unaware. We are hopeful that users of this dataset will call our attention to those sources so that we may expand and update this dataset.

Sources Consulted

Adamec, Ludwig W. 2012. Historical Dictionary of Afghanistan, 4th Edition (Scarecrow Press).

Akhtar, Sohail. 2018. "The Historical and Archeological Significance Dera Ghazi Khan District through Ages." Archaeology & Anthropology Open Access 2:161-64.

A Landowner. 1853. Demarara After Fifteen Years of Freedom (T. Bosworth).

Babatundel, Ibrahim, Bako Iyanda, Raheem Mayowa and Abdulyekeen Ola. 2014. "Appraisal of Population Trends in Ilorin, Nigeria." Journal of Sustainable Development in Africa 16: 1-14.

Bacchi, Roberto. 1977. The Population of Israel (Hebrew University).

Bairoch, Paul, Jean Batou, and Pierre Chèvre. 1988. La population des Villes Européennes, 800-1850: Banque de Données et Analyse Sommaire des Résultats (Droz).

Bigon, Liora. 2016. French Colonial Dakar: The Morphogenesis of an African Regional Capital (Oxford University Press).

Biguzzi, Andrea. ND. "Worldcitypop.com" ['Database of City Populations from around the World over Time"]

Blanton, Richard, and Lane Fargher. 2010. Collective Action in the Formation of Pre-Modern States (Springer).

Bosker, Maarten, Eltjo Buringh, and Jan Luiten Van Zanden. 2013. "From Baghdad to London, Unraveling Urban Development in Europe, North Africa and the Middle East, 800–1800." The Review of Economics and Statistics 95: 1418–37.

Brown, Kenneth. 1976. The People of Salé: Tradition and Change in a Moroccan City, 1830-1930 (Manchester University Press).

Bruun, Malthe Conrad. 1823. Universal Geography, or A description of All the Parts of the World (Adam Black).

Canadian Institute of Ukrainian Studies, ND, Internet Encyclopedia of Ukraine, https://www.encyclopediaofukraine.com/

Cao Shuji. 2000a. A Demographic History of China, Volume 4; the Ming Period (Fudan University Press [in Mandarin]).

Cao Shuji. 2000b. A Demographic History of China, Volume 5; the Qing Period (Fudan University Press [in Mandarin]).

Cerman, Markus, and Herbert Knittler. 2004. "Town and Country in the Austrian and Czech Lands, 1450-1800" in S.R. Epstein ed., Town and Country in Europe, 1300-1800 (Cambridge University Press).

Cesaretti, Rudolf, José Lobo, Luís M. A. Bettencourt, Scott G. Ortman, and Michael E. Smith. 2016. “Population-Area Relationship for Medieval European Cities,” PLoS ONE 11.

Chandler, Tertius. 1987. Four Thousand Years of Urban Growth: An Historical Census (Edwin Mellen).

Cook, Thomas. 1904. Cook's Practical Guide to Algiers, Algeria, and Tunisia (Thomas Cook and Company).

Coquery-Vidrovitch, Catherine. 2005. The History of African Cities South of the Sahara: From the Origins to Colonization (Markus Weiner).

Cotton, James Sutherland, Richard Burn, and William Stevenson Meyer. 1908. Imperial Gazetteer of India (Oxford at the Clarendon Press).

Curto, José. 1999. "Anatomy of a Demographic Explosion: Luanda, 1844-1850." International Journal of African Historical Studies 32: 381-405.

Curto, José, and Raymond R. Gervais, 2001. “The Population History of Luanda during the Late Atlantic Slave Trade, 1781-1844.” African Economic History 29: 1-59.

DeCorse, Christopher, and Sam Spiers. 2023. “Tradition Summary: West African Regional Development.” HRAF. https://ehrafarchaeology.yale.edu/document?id=fa77-000.

Denbow, James. 2013. The Archaeology and Ethnography of Central Africa (Cambridge University Press).

Dickson, Kwamina. 1971. A Historical Geography of Ghana (Cambridge University Press).

Doeppers, Daniel. 1972. “The Development of Philippine Cities Before 1900.” The Journal of Asian Studies 31: 769-92.

Donkoh, Wilhelmina, Joseline. 2004. “Kumasi: Ambience of Urbanity, Tradition, and Modernity. Transactions of the Historical Society of Ghana 8: 167-83.

Ebong, Maurice Okon. 1980. Urban Dimensions in Calabar, Nigeria (Michigan State University, African Studies Center).

El Bushra, El Sayed. 1979. "Some Demographic Indicators for Khartoum Conurbation, Sudan" Journal of Middle Eastern Studies 15: 295-309.

Eldredge, Elizabeth. 2015. Kingdoms and Chiefdoms of Southeastern Africa (University of Rochester Press).

Encyclopaedia Britannica. Various editions.

Encyclopaedia Iranica, Online edition. ND. (Brill).

Encyclopaedia Islamica Online edition. ND. (Brill).

Falola, Toyin, and Matthew Heaton. 2008. A History of Nigeria (Cambridge University Press).

Francia, Luis. 2021. A History of the Philippines: From Indios Bravos to Filipinos (Abrams Press).

Freund, Bill. 2007. The African City: A History (Cambridge University Press).

Grabowsky, Volker. 2017. "Population Dynamics in Lan Na During the 19th Century." Journal of the Siam Society 105: 197-244.

Haim, Sylvia, and Elie Kedourie. 1988. Essays on the Economic History of the Middle East (Routledge).

Hendro, Eko Punto, Susanna Ratih Sari, Siddhi Saputro, Indriyanto. 2021. “Demak Kingdom: Study of Environmental Condition and Geographical.” E3S Web of Conferences 317: 1-10.

Hunter, William Wilson. 1886. The Imperial Gazetteer of India (Trubner and Company).

Issawi, Charles. 1988. The Fertile Crescent, 1800-1914: A Documentary Economic History (Oxford University Press).

http://www.jewishvirtuallibrary.org/jewish-and-non-jewish-population-of-israel-palestine-1517-present

Jennings, Ronald. 1976. “Urban Population in Anatolia in the Sixteenth Century: A Study of Kayseri, Karaman, Amasya, Trabzon, and Erzurum.” International Journal of Middle East Studies 7: 21-57.

Johnson, Ben. ND. “Merthyr Tydfil and the Welsh Men of Steel.” History Magazine. http://www.historic-uk.com/HistoryUK/HistoryofWales/Merthyr-the-Welsh-Men-of-Steel/

Kisangani, Emizet Francois, and Scott F. Bobb. 2010. Historical Dictionary of the Democratic Republic of the Congo (Scarecrow Press).

Kour, Z.H. 2005. The History of Aden, 1839-1972 (Routledge).

Kwon, Tai Hwan, Hae Young Lee, Yunshik Chang, and Eui-Young Yu. 1975. The Population of Korea (The Population and Development Studies Center).

Lahmeyer, Jan. (1999/2006). “Populstat.info.” [“Population Statistics: Historical Demography of All Countries, their Divisions and Towns.”]

Lambek, Michael. 2002. The Weight of the Past: Living with History in Mahajanga, Madagascar (Palgrave McMillan).

Law, Robin. 2000. “Ouidah: Pre-Colonial Urban Center in West Africa," in David Anderson and Richard Rathbone, eds, Africa's Urban Past (Heinemann).

Law, Robin. 2005. Ouidah: The Social History of a West African Slaving Port, 1727–1892 (Ohio University Press).

Lee, Jonathan. 1987. “The History of Maimana in Northwestern Afghanistan 1731-1893.” Journal of Persian Studies 25: 107-124

Lovejoy, Paul. 2011. Transformations in Slavery: A History of Slavery in Africa (Cambridge University Press).

Makgala, Christian. 2002. “A Note on the History of Serowe.” Botswana Notes and Records 34: 160-63.

McCarthy, Justin. 1976. “Nineteenth-Century Egyptian Population.” Middle Eastern Studies 12: 1-39.

Monroe, J. Cameron. 2011. “Urbanism on West Africa’s Slave Coast.” American Scientist 99:400.

Morse, Jedidiah, and Richard Cary Morse. 1821. A New Universal Gazetteer, Or, Geographical Dictionary (Sherman Converse).

Murray, John, 1858. A Handbook for Travelers in Syria and Palestine, Vol II.

Ouyyanon, Porphant.1997. "Bangkok's Population and the Ministry of the Capital in Early 20th Century Thai History." Southeast Asian Studies 35: 240-60.

Papinot, Edmond. 1910. Historical and Geographical Dictionary of Japan (Librairie Sansaisha).

Parkin, David. 2015. Neighbours and Nationals in an African City Ward (Routledge).

Porter, Jonathan. 1993. "The Transformation of Macau" Pacific Affairs 66: 7-20.

Raymond, Andre. 2002. Arab Cities During the Ottoman Period: Cairo, Syria, and the Magreb (Routledge).

Reba, Meredith., Femke Reitsma, and Karen Seto. 2016. “Spatializing 6,000 Years of Global Urbanization from 3700 BC to AD 2000.” Scientific Data 3: 1-16.

Richardson, James. 1853. Narrative of a Mission to Central Africa Performed in the Years 1850-51, Volume 2: Under the Orders and at the Expense of Her Majesty's Government (Chapman and Hall).

Rijke-Espstein, Tasha, 2023. Children of the Soil: The Power of Built Form in Urban Madagascar (Duke University Press).

Rozman, Gilbert. 1976. Urban Networks in Russia 1750-1800 and Premodern Periodization (Princeton University Press).

Russ, Jake. N.D. https://github.com/JakeRuss/bairoch-1988/blob/master/bairoch-1988.csv

Saitō, Seiji. 1984. "Urban Population during Edo Period", Chiiki Kaihatsu 9: 48–63 [in Japanese].

Samuel, Isaac. 2024. "A Complete History of the Old City of Gao ca. 700-1898." https://www.africanhistoryextra.com/p/a-complete-history-of-the-old-city

Sandal, Ersin, and Fatih Agiduzel. 2014. "Spatial Development of Tarsus and the Changes in Land Use," in Recef Ife, Turgut Onay, Igor Sharuho, and Emin Atasoy, Urban and Urbanization (St. Kliment Ohridski University) pp. 570-78.

Satow, Ernest Mason, and A. G. S. Hawes. 1884. A Handbook for Travellers in Central & Northern Japan (Kelly and Co.)

Schinz, Alfred. 1989. Cities in China (Gerbruder Borntraeger).

Shinn, David, and Thomas Ofcansky. 2004. Historical Dictionary of Ethiopia (Scarecrow Press).

Smith, R.S. 1962. Ijaiye: "The Western Palatinate of the Yoruba." Journal of the Historical Society of Nigeria 2: 329-49.

Sternstein, Larry. 1966. "The Distribution of Thai Centres at Mid-Nineteenth Century." Journal of Southeast Asian History 7: 66-72.

Thapa, Rajesh Bahadur, Yuji Murayama, and Shailja Ale. 2008. "Kathmandu." Cities 25: 45-57.

Thornton, Edward, and Arthur Naylor Wollaston. 1858. A Gazetteer of the Territories under the Government of the East-India Company (East India Company).

Thornton, John. 2013. "Kongo, Teke, and Loango: History to 1483." In Kevin Shingleton ed., Encyclopedia of African History (Routledge).

Thornton, John. 2021. "Revising the Population History of the Kingdom of Kongo." The Journal of African History 62: 201–12.

Toffin, Gerard. 1994. "The Farmers in the City. The Social and Territorial Organization of the Maharjan of Kathmandu." Anthropos 89.

Topalidis, Sam. 2022. “A History of Trabzon.” Pontos World. https://pontosworld.com/index.php/pontus/news/797-the-history-of-trabzon

Tsao, Li-ling. ND. “A Study on the Formation of Kaohsiung City, Taiwan.” Mimeo.

United States, Bureau of the Census. 1905. Census of the Population of the Philippine Islands, Taken Under the Direction of the Philippine Commission, 1903 (United States Government Printing Office).

Unknown Author. 1880. Gazetteer of the Bombay Presidency: Cutch, Palanpur, and Mahi Kantha. (Government Central Press).

Vansina, Jan. 1978. The Children of Woot: A History of the Kuba Peoples (University of Wisconsin Press).

Vansina, Jan. 2004. How Societies are Born (University of Virginia Press.)

Walker, John. (1810). The Universal Gazetteer: Being a Concise Description, Alphabetically… (Harvard University).

Wikipedia. ND. Various entries.

Wilkinson, David. 1993. “Spatial-Temporal Boundaries of African Civilizations Reconsidered, Part 1.” Comparative Civilizations Review 29: 52-90.

Wilkinson, David. 1994. “Spatial-Temporal Boundaries of African Civilizations Reconsidered, Part 2.” Comparative Civilizations Review 31: 46-60.

Wright, George. 1902. Asiatic Russia (McClure, Phillips & Co)

Wrigley, E.A. 1987. People, Cities, and Wealth: The Transformation of Traditional Society (Blackwell).

Appendix B — Supplemental Data Descriptions

Univariate distributions and bivariate relationships of the core market variables (market area, potential agricultural productivity, harbors, cities in 1800, urban population in 1800), for the Eurasia-and-Africa market sample (N = 41,435). Potential agricultural productivity is reported in Mkcal/km²/year, the same units used in the main paper. In Figure B1 the left-hand panels are drawn on linear scales and the right-hand panels on log(1 + x) scales, so that markets with no harbors, no cities, or no urban population remain in every panel; dashed lines mark medians and solid lines means. Harbors and cities are binned at unit width.

Figure B1. Univariate distributions of market characteristics, linear (left) and log (right) scales.

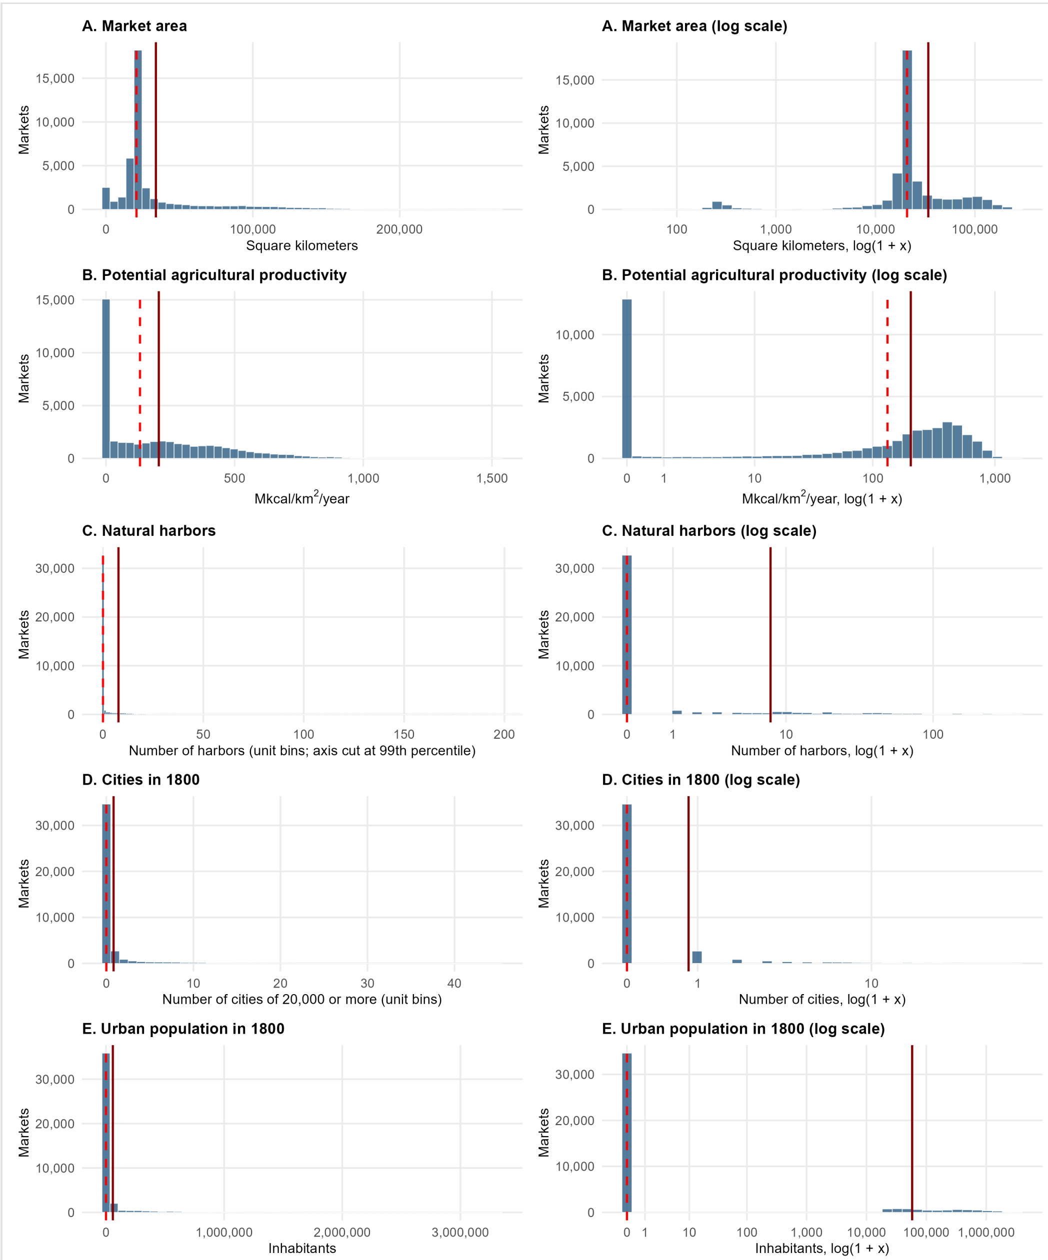

Figure B2. Bivariate relationships among market variables.

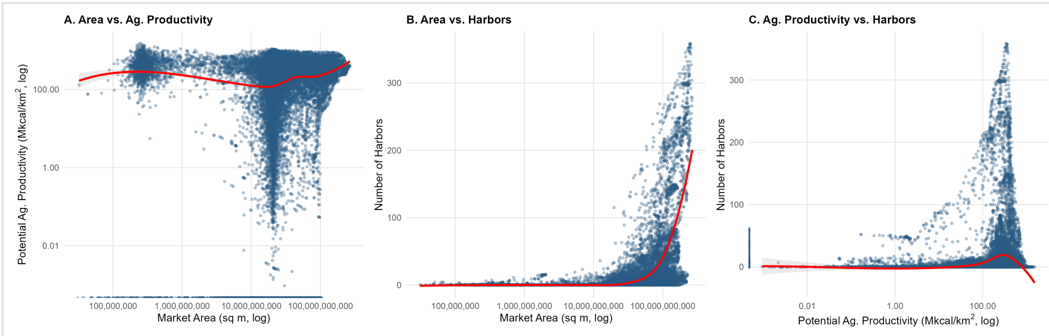

Appendix C — Supplemental Potential Market Cluster Figures

Leave-one-block-out (LOBO) spatial robustness for the potential-market-cluster / location-quotient analysis (the eight market categories of main-paper Table 2 and Map 5). A 3×3 grid is laid over Eurasia and Africa; each iteration drops every cluster whose centroid falls in one cell and re-estimates. The grid itself is shown as Figure C1 below — main-paper Map 6 shows the eight market cluster categories without the grid overlay. Exhibits are ordered as they are cited in the main text. Table C1 reports the leave-one-block-out results for urban population in 1800: for the full sample and for each dropped block, the category of market clusters with all three characteristics (large markets with fertile soils and sea access) and the category with none of them. Table C2 reports location quotients for single-factor and two-factor market clusters, each constructed independently with the cut points of main-paper Section 4.2, so that a market and its cities can appear in more than one row. Table C3 repeats main-paper Table 2 for energy budgets from 30 to 50 megajoules, and Table C4 repeats Table C1 for urban population circa 1700 (329 cities of at least 20,000 inhabitants).

Conventions. Fertile soils = potential staple-crop kilocalorie output per km² in the top 30% of markets; large market = market area in the top 15%; sea access = one or more natural harbors. All exhibits below are computed on the same sample as main-paper Table 2 (546 cities of at least 20,000 inhabitants in 1800). Throughout the paper and appendices, “% of Eurasian and African Land” is measured against the total market area — the dissolved footprint of all market clusters — so that land shares sum to 100 and the all-Eurasia-and-Africa location quotient equals 1.00 by construction. About 2.4% of the Eurasian and African landmass lies outside every market cluster, chiefly high-friction, river-poor terrain in Sub-Saharan Africa; it contains no city of 20,000 in 1800 (and one in 1700, Allada in present-day Benin) and is excluded from the denominator. In each LOBO iteration the denominator is recomputed over the remaining clusters.

Figure C1. The 3×3 leave-one-out grid over the market clusters of Eurasia and Africa.

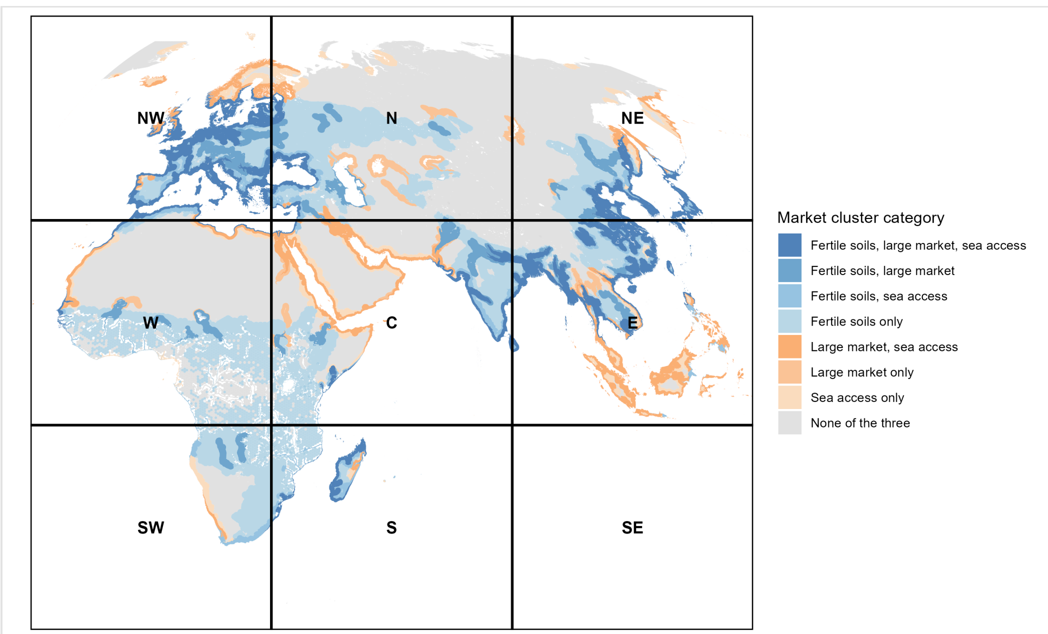

Table C1. Leave-one-block-out estimates for urban population in 1800 — market clusters with all three characteristics (large markets with fertile soils and sea access) and with none.

			Market clusters with all three characteristics				Location quotient	
Block dropped	Cities remaining	Urban pop. (000s)	Cities	Urban pop. (000s)	% of land	% of urban pop.	All three	None
None (full sample)	546	31,842	333	22,599	10.2	71.0	6.96	0.06
Southwest	546	31,842	333	22,599	10.8	71.0	6.59	0.05
South	546	31,842	333	22,599	9.9	71.0	7.19	0.06
Southeast	546	31,842	333	22,599	10.2	71.0	6.96	0.06
West	524	31,063	333	22,599	11.2	72.8	6.47	0.05
Center	477	28,513	333	22,599	11.1	79.3	7.16	0.05
East	391	19,631	191	10,946	5.2	55.8	10.65	0.08
Northwest	316	19,374	142	11,654	6.8	60.1	8.83	0.09
North	493	29,672	333	22,599	20.2	76.2	3.76	0.03
Northeast	529	30,956	333	22,599	10.7	73.0	6.85	0.06

Table C2. Location quotients of single-factor and two-factor market clusters, 1800.

Market Category Region	% of Eurasian and African Land	No. Potential Markets	No. Cities 1800	Urban Pop. (thousands)	% of Eurasian and African Urban Pop.	Location Quotient
Fertile soils + large market	14.7	2,856	399	25,910	81.4	5.54
Fertile soils + sea access	13.3	3,546	372	24,191	76.0	5.71
Large market + sea access	16.8	5,027	362	24,029	75.5	4.49
Fertile soils (top 30% of agricultural potential)	41.3	12,431	507	30,416	95.5	2.31
Large market (top 15% of area)	23.7	6,216	423	26,931	84.6	3.57
Sea access (one or more natural harbors)	23.9	8,743	391	24,973	78.4	3.28
All of Eurasia and Africa *(reference)*	100.0	41,435	546	31,842	100.0	1.00

Sensitivity to the energy budget

The potential markets in the main paper are drawn with a 40 megajoule energy budget. To check that the results of main-paper Table 2 are not an artifact of that tuning, the full market-cluster pipeline was repeated for every budget from 30 to 50 MJ in 2 MJ steps: the markets drawn at each budget were re-classified with the cut points recomputed within that budget (top 30% of agricultural potential, top 15% of area, one or more natural harbors), dissolved into contiguous clusters with the same precedence, ocean/lake erase and minimum-area rules as the main-paper clusters, and joined to the 1800 cities exactly as in Table 2. Land shares are therefore shares of total dissolved cluster area at every budget, so each row of Table C3 is on the same basis as Table 2, and the 40 MJ row reproduces Table 2 exactly.

Table C3. Sensitivity of the location quotient analysis to the energy budget, 30 to 50 megajoules.

			Market clusters with all three characteristics				Location quotient	
Energy budget (MJ)	Markets (Eurasia + Africa)	Dissolved clusters	Markets	% of land	Cities 1800	% of urban pop.	All three	None
30	41,385	2,940	2,127	7.8	298	67.1	8.60	0.05
32	41,404	2,817	2,140	8.1	306	68.0	8.40	0.05
34	41,414	2,728	2,181	8.7	312	68.6	7.88	0.05
36	41,422	2,646	2,204	9.2	314	68.9	7.49	0.05
38	41,428	2,590	2,249	9.8	327	71.0	7.25	0.06
40*	41,435	2,543	2,255	10.2	333	71.0	6.96	0.06
42	41,441	2,505	2,275	10.7	343	73.0	6.82	0.06
44	41,444	2,487	2,299	11.0	346	73.4	6.67	0.06
46	41,448	2,459	2,323	11.6	360	76.2	6.57	0.05
48	41,452	2,430	2,337	12.0	364	77.3	6.44	0.05
50	41,453	2,397	2,373	12.6	368	77.6	6.16	0.05

Urban population in 1700

Table C4 repeats the leave-one-block-out exercise of Table C1 with urban populations circa 1700 (329 cities of at least 20,000 inhabitants).

Table C4. Leave-one-block-out estimates for urban population in 1700.

			Market clusters with all three characteristics				Location quotient	
Block dropped	Cities remaining	Urban pop. (000s)	Cities	Urban pop. (000s)	% of land	% of urban pop.	All three	None
None (full sample)	316	20,632	216	15,484	10.2	75.0	7.36	0.11
Southwest	315	20,602	216	15,484	10.8	75.2	6.98	0.11
South	316	20,632	216	15,484	9.9	75.0	7.60	0.11
Southeast	316	20,632	216	15,484	10.2	75.0	7.36	0.11
West	301	20,110	216	15,484	11.2	77.0	6.85	0.10
Center	292	19,024	216	15,484	11.1	81.4	7.35	0.11
East	222	12,419	126	7,411	5.2	59.7	11.39	0.17
Northwest	169	12,443	90	8,074	6.8	64.9	9.53	0.17
North	288	18,990	216	15,484	20.2	81.5	4.03	0.06
Northeast	309	20,204	216	15,484	10.7	76.6	7.19	0.11

Appendix D — Supplemental Random Forest Figures and Tables

Exhibits are ordered as they are cited in Section 5 of the main paper. Throughout, the dependent variable is ln(1 + urban population) and the predictors are log potential agricultural productivity, log market area, and log harbors; the main specification is main-paper Table 3. Standard deviations across holdout draws are given in parentheses throughout.

Sampling and estimation checks

Table D1. Robustness — stratified vs. weighted sampling (both models).

	Model 1: Urbanization 1800		Model 2: Urbanization 1700	
	Stratified	Weighted	Stratified	Weighted
Holdout R^2	0.518	0.465	0.525	0.468
	(0.037)	(0.047)	(0.043)	(0.052)
Holdout RMSE	2.790	2.742	2.456	2.392
	(0.145)	(0.157)	(0.149)	(0.166)
Trees	4,800	4,800	4,800	4,800
Holdout draws	1,200	1,200	1,200	1,200

Notes: "Stratified" is the two-phase stratified priority sampling of the main specification; "Weighted" draws markets with probability proportional to area. Both use 400 markets per draw and the same hyperparameters. The stratified 1800 column matches main-paper Table 3, Panel B.

Table D2. Draw-size sensitivity.

	Model 1: Urbanization 1800			Model 2: Urbanization 1700		
Draw size	Obs/tree	R^2	RMSE	Obs/tree	R^2	RMSE
400 (main)	390	0.518	2.790	390	0.525	2.456
		(0.037)	(0.145)		(0.043)	(0.149)
350	350	0.525	2.894	350	0.530	2.564
		(0.038)	(0.152)		(0.043)	(0.154)
450	450	0.510	2.706	450	0.519	2.368
		(0.036)	(0.140)		(0.043)	(0.144)

Notes: The main specification draws 400 markets per sample. Holding everything else fixed, drawing 350 or 450 moves holdout R^2 by less than a point.

Table D3. Logged vs. unlogged dependent variable.

	Model 1: Urbanization 1800		Model 2: Urbanization 1700	
	Logged	Raw	Logged	Raw
Dependent variable	$\ln(1 + \text{pop})$	Population	$\ln(1 + \text{pop})$	Population
Holdout R^2	0.518	0.439	0.525	0.351
	(0.037)	(0.064)	(0.043)	(0.093)
Holdout RMSE	2.790	128,421	2.456	101,804
	(0.145)	(23,454)	(0.149)	(18,617)
Trees	4,800	4,800	4,800	4,800
Holdout draws	1,200	1,200	1,200	1,200

Notes: "Logged" columns are the main specification; "Raw" columns use the unlogged dependent variable with otherwise identical settings. RMSE for the raw columns is in inhabitants and is not comparable to the logged columns; R^2 is.

Figure D1. Hyperparameter tuning.

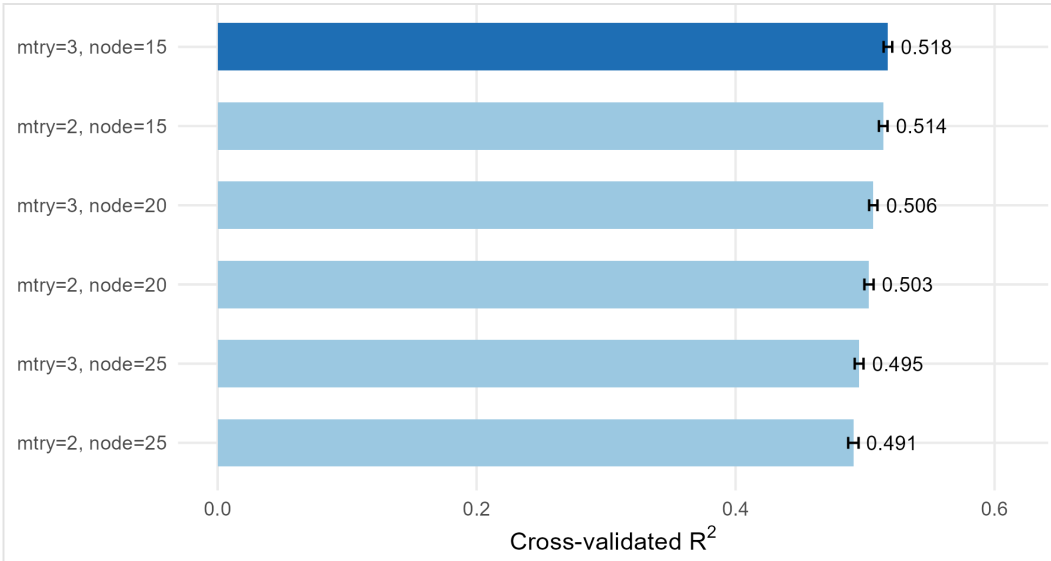

Notes: Cross-validated R^2 for the six candidate configurations; the selected configuration (three predictors per split, minimum node size 15) is shown in dark blue. Maximum tree depth is fixed at seven throughout.

Figure D2. Holdout performance across the 1,200 draws.

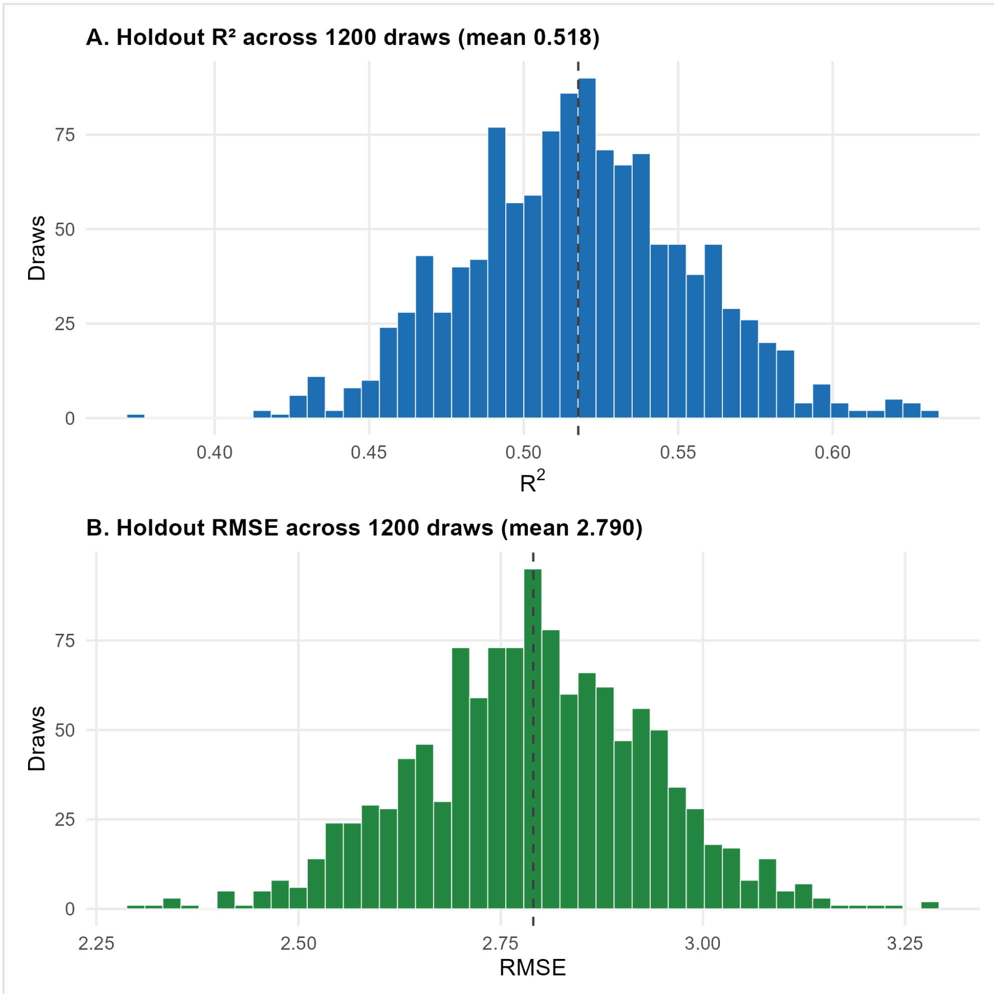

Notes: Dashed lines mark the means reported in main-paper Table 3, Panel B (R^2 0.518, RMSE 2.790).

Permutation test

Table D4. Permutation test results.

	Model 1: Urbanization 1800	Model 2: Urbanization 1700
Panel A: Holdout R^2		
Observed R^2	0.518	0.525
Permuted mean R^2	0.0098	0.0196
Permuted max R^2	0.0468	0.0508
p-value ($R^2 \geq$ observed)	< 0.010	< 0.010
Panel B: Holdout RMSE		
Observed RMSE	2.790	2.456
Permuted mean RMSE	3.9988	3.5257
Permuted min RMSE	3.9234	3.4691
p-value (RMSE $\leq$ observed)	< 0.010	< 0.010
Permutations	100	100
Trees per permutation	480	480

Notes: Each permutation shuffles the dependent variable, retrains a downsampled ensemble (480 trees, 120 holdout draws) and re-evaluates. The p-value is the share of the 100 permutations with R^2 at least as large (RMSE at least as small) as observed.

Spatial cross-validation

Table D5. RF combined summary (full sample + LOBO) — urbanization in 1800.

Iteration	Markets Dropped	Markets Remaining	Obs. per Tree	Holdout RMSE	Holdout R^2
Full Sample	—	41,435	400	2.790	0.518
Drop SW	1,113	40,322	387	2.805	0.526
Drop S	1,286	40,149	386	2.759	0.542
Drop SE	14	41,421	400	2.789	0.518
Drop W	5,203	36,232	338	2.965	0.504
Drop C	5,770	35,665	339	2.722	0.535
Drop E	3,712	37,723	370	2.616	0.496
Drop NW	3,314	38,121	378	2.693	0.451
Drop N	12,936	28,499	278	2.914	0.549
Drop NE	6,393	35,042	340	2.904	0.492
Jackknife 95% CI				[2.090, 3.491]	[0.330, 0.706]

Notes: Each row retrains the full 4,800-tree ensemble after dropping every potential market whose initial point falls in the indicated cell of the 3×3 grid of Figure C1, and evaluates on the 1,200 holdout draws filtered the same way. The jackknife interval follows Künsch (1989) with n = 9 blocks. Panel D of main-paper Table 3 summarizes these results.

Robustness test: urbanization in 1700

The paper reports levels in 1800. The same pipeline was estimated on the level of urban population in 1700 (329 cities of at least 20,000 inhabitants); those results are presented here.

Table D6. Random forest summary — urbanization in 1700.

Iteration	Markets Dropped	Markets Remaining	Obs. per Tree	Holdout RMSE	Holdout R^2
Full Sample	—	41,435	400	2.456	0.525
Drop SW	1,113	40,322	387	2.473	0.531
Drop S	1,286	40,149	386	2.427	0.550
Drop SE	14	41,421	400	2.455	0.525
Drop W	5,203	36,232	338	2.589	0.528
Drop C	5,770	35,665	339	2.431	0.533
Drop E	3,712	37,723	370	2.265	0.527
Drop NW	3,314	38,121	378	2.321	0.451
Drop N	12,936	28,499	278	2.671	0.527
Drop NE	6,393	35,042	340	2.557	0.489
Jackknife 95% CI				[1.671, 3.240]	[0.342, 0.708]

Notes: As Table D5, with dependent variable ln(1 + population in 1700).

Figure D3. Variable importance — urbanization in 1700.

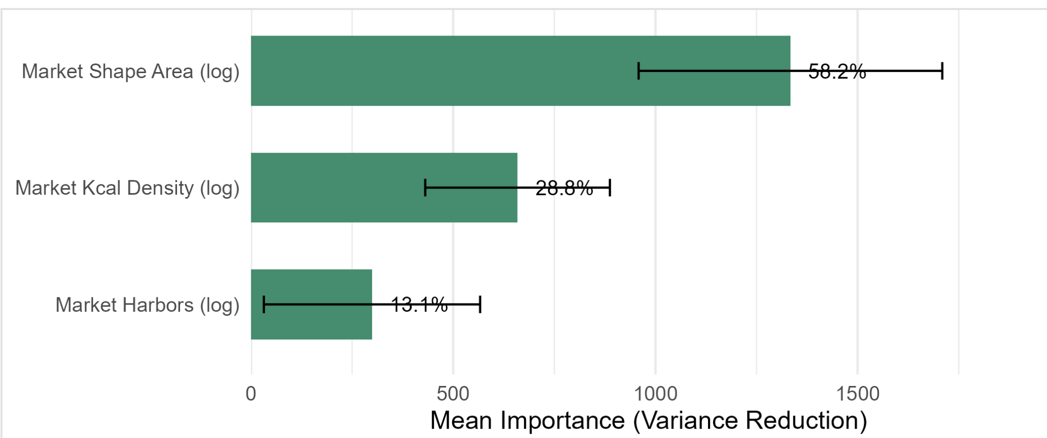

Figure D4. Partial dependence — urbanization in 1700.

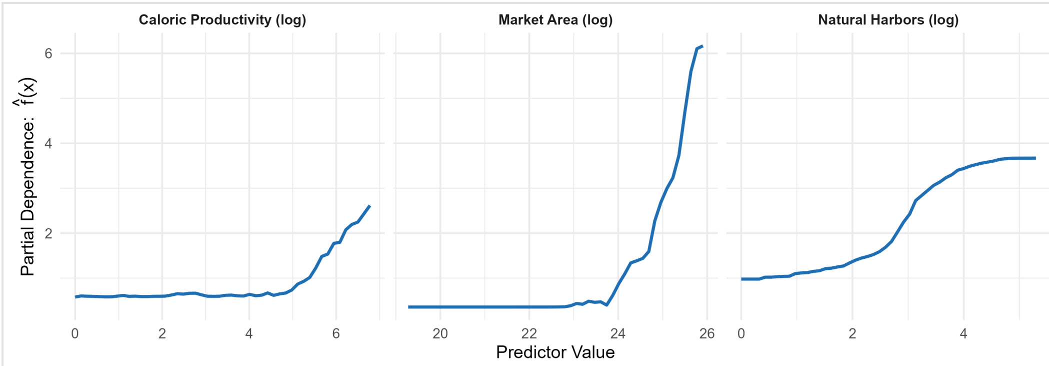

Sensitivity to the energy budget

The potential markets in the main paper are drawn with a 40 megajoule energy budget. The full market-drawing pipeline was repeated for budgets from 30 to 50 MJ in 2 MJ steps, and the main random forest (urbanization in 1800) was re-estimated on each set with the same predictors, draw design, and tuning grid. Table D7 summarizes the eleven fits. Holdout R^2 stays between 0.49 and 0.52 across the whole range and every budget selects the same hyperparameters; the 40 MJ specification is neither the best- nor the worst-fitting budget.

Table D7. Random forest sensitivity to the energy budget — urbanization in 1800.

Energy budget (MJ)	Eurasia + Africa markets	Tuned (mtry, min. node)	CV R^2	Holdout R^2	Holdout RMSE
30	41,385	(3, 15)	0.511	0.513	2.902
				(0.036)	(0.133)
32	41,404	(3, 15)	0.518	0.517	2.872
				(0.038)	(0.139)
34	41,414	(3, 15)	0.523	0.522	2.839
				(0.038)	(0.142)
36	41,422	(3, 15)	0.518	0.519	2.822
				(0.038)	(0.145)
38	41,428	(3, 15)	0.520	0.521	2.799
				(0.038)	(0.146)
40*	41,435	(3, 15)	0.519	0.518	2.786
				(0.036)	(0.145)
42	41,441	(3, 15)	0.505	0.502	2.796
				(0.037)	(0.146)
44	41,444	(3, 15)	0.503	0.506	2.758
				(0.037)	(0.145)
46	41,448	(3, 15)	0.494	0.495	2.759
				(0.038)	(0.149)
48	41,452	(3, 15)	0.495	0.493	2.757
				(0.038)	(0.148)
50	41,453	(3, 15)	0.492	0.492	2.759
				(0.038)	(0.145)

Notes: Each row re-estimates the main random forest, with the settings of Tables D1 and D5 and a 5-fold cross-validated tuning grid, on the potential markets drawn with that energy budget; the 40 MJ row (*) is the main specification. Per-budget variable importance was not retained by the per-budget runs and is therefore not reported; the 40 MJ importances are those of main-paper Table 3, Panel C.