

NBER WORKING PAPER SERIES

SPATIAL INEFFICIENCIES IN AFRICA'S TRADE NETWORK

Tilman Graff

Working Paper 25951

<http://www.nber.org/papers/w25951>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

June 2019

The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Tilman Graff. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Spatial Inefficiencies in Africa's Trade Network

Tilman Graff

NBER Working Paper No. 25951

June 2019

JEL No. F1,O18,R4

ABSTRACT

I assess the efficiency of transport networks for every country in Africa. Using rich spatial data, I simulate trade flows over more than 70,000 links covering the entire continent. I maximise over the space of networks and find the optimal road system for every African state. My simulations predict that Africa would gain 1.1% of total welfare from better organising its national road systems. I then construct a novel dataset of local network inefficiency and I find that colonial infrastructure projects significantly skew trade networks towards a sub-optimal equilibrium. I also find evidence for regional favouritism and inefficient aid provision.

Tilman Graff

Busara Center for Behavioral Economics

Daykio Plaza, Ngong Ln

Nairobi

Kenya

tgraff@princeton.edu

1 Introduction

One of the key factors holding back Africa’s economic development is its inadequate infrastructure system. Sub-Saharan Africa’s coverage with paved roads is by far the lowest of any world region, with only 31 total paved road kilometres per 100 square kilometres of land, compared to 134 in other low-income countries (Foster and Briceño-Garmendia, 2010). In rural areas, more than two thirds of the population live further than two kilometres away from any all-season road (Teravaninthorn and Raballand, 2009). As a result, trade costs in Africa are the highest in the world, stifling interregional trade (Lima and Venables, 2001; Foster and Briceño-Garmendia, 2010; The Economist, 2015).

The main policy prescription by international development organisations for African governments has been to drastically expand the current road network. The World Bank has identified an annual infrastructure gap amounting to 9.6 billion US dollars and urges countries in Sub-Saharan Africa to spend almost one per cent of GDP on building new roads (Foster and Briceño-Garmendia, 2010). This reasoning is also reflected in the composition of development aid – in 2017, by far the largest share of World Bank lending to African countries was allocated to transport infrastructure projects (The World Bank, 2017). There is a clear consensus that Africa needs more roads.

In this paper, I investigate a neglected, yet powerful second source of spatial inefficiency in Africa’s transport system. I don’t ask if the continent has too *few* roads, but rather analyse whether the current infrastructure is *in the wrong place*. Do Africa’s roads connect the right areas to promote beneficial trade? How would a social planner design a perfect transport network which optimises welfare in a given country? Which African country is closest to its hypothetical optimum? And why are some locations systematically cut off from the national trade system?

I derive the unique optimal trade network for every country in Africa.¹ Using rich data from satellites and online routing services, I first construct an interconnected economic topography of more than 10,000 rectangular grid cells covering the entire continent. I then employ a simple network model to simulate trade flows through more than 70,000 links spanning all of Africa. In a second step, I use a variant of a recently established framework

¹I consider every member state of the African Union (which includes Western Sahara) – a total of 55 countries.

by Fajgelbaum and Schaal (2017) to optimise over the space of networks and find the optimal transport system given the underlying economic fundamentals for every African country. An intuitive thought experiment demonstrates this process: suppose the social planner were to observe the spatial distribution of roads, people, and economic activity in a given country before being allowed to lift all roads from their current location, freely shuffle them around, and then reorganise them in the most efficient way for mutual trade. She does not get to build completely new roads, but is only allowed to move infrastructure from one part of the country to another. In this exercise, she takes into account local incentives for trade between all sets of neighbours on a complex network graph, regional differences in trade costs caused by geographical and network characteristics, and heterogeneous costs to constructing new roads depending on the underlying terrain.

I then compare these optimal networks to the current system. I argue that the degree to which the optimum differs from the status quo can be interpreted as an intuitive measure for the inefficiency of a country's current road network. I show that potential welfare gains from reshuffling roads are substantial, improving overall welfare on the entire continent by about 1.15%. I also identify South Sudan as the country with the worst transport network in Africa. South Sudan, the world's youngest nation, stands to gain more than six per cent of total welfare solely through the better reorganisation of its road network.

On the regional level, this scenario creates winners and losers. The model identifies some areas as having *too* many roads and decides to put them to better use somewhere else. These areas were inefficiently overequipped with transportation infrastructure before the major reshuffling exercise. Other regions, however, did not have enough infrastructure given their relative position in the network and are now awarded additional roads by the social planner. I identify these areas as discriminated against by the current transportation network design. By comparing welfare levels before and after the hypothetical intervention, I create a novel dataset of local infrastructure discrimination for more than 10,000 cells covering the entire African continent.

Why are some regions systematically cut off from the benefits of efficient trade? I use a variety of empirical designs to analyse the substantial spatial variation present in my dataset. Firstly, I investigate the long-run effects of large infrastructure investments from

the colonial area. Similarly to Jedwab and Moradi (2016), I find a persistent impact of railway lines constructed by the colonial powers over a century ago. Plausibly exogenous variation in the number of kilometres crossing a given area significantly skews the current trade network towards a suboptimal state today. Even though many of the railway lines have fallen into disarray since independence, regions close to colonial railroads still have too much road infrastructure given their relative position in the network. I argue that the colonial era transport revolution coordinated infrastructure investments towards a spatial equilibrium that persists until today, even though it has since become inefficient. In contrast, proximity to railway lines that were planned, but by historical accident never built, does not predict any significant departure from the optimal spatial distribution, bolstering the causal nature of this relationship. I also provide suggestive evidence for local general equilibrium effects and find that areas bordering a railway might benefit at the expense of their immediate neighbours.

Secondly, I analyse how spatial inefficiencies in Africa's trade system are related to ethnic power dynamics. I find no evidence that ethnic regions that are politically discriminated against have less than optimal infrastructure stocks. I estimate precise null effects of ethnicities being excluded from the central government, historically involved in an ethnic war, or split by arbitrary colonial borders (Michalopoulos and Papaioannou, 2016) on current transport network inefficiency. However, I do find substantial evidence for regional favouritism – homelands of national leaders have significantly more infrastructure than is nationally efficient. Finally, I investigate the extent to which foreign aid projects have succeeded in alleviating the imbalances in Africa's transport networks. I present descriptive evidence demonstrating that World Bank funds have not gone towards the regions most in need of additional infrastructure. Instead, areas that are identified as having too many roads are associated with more Bank lending. The same patterns hold for development aid from China.

My study contributes to several strands of literature. In analysing the impact of transport revolutions, I add to the large body of work devoted to identifying the economic returns to improving infrastructure systems. A series of rigorous studies have gauged the welfare effects of the expansion of the US railway network in the 19th century (Donaldson

and Hornbeck, 2016; Swisher, 2017), colonial railway systems in India (Donaldson, 2018; Burgess and Donaldson, 2012), or highway systems in China (Faber, 2014; Baum-Snow et al., 2017). In contrast to these studies, I do not analyse the impact of historical transport revolutions, but rather measure how much a hypothetical first-best transport system would improve welfare. Methodologically, I harness recent advances bringing insights from the optimal transport literature into the economics discourse. Most directly, I apply the framework by Fajgelbaum and Schaal (2017) to construct the optimal trade network for every African country. They are the first to optimise over the space of networks in order to find the globally efficient transport system in an economics context, though the problem has long featured prominently in the mathematics literature (for a textbook treatment, see Bernot et al., 2009). To the best of my knowledge, my study is the first to employ their framework in a development context. Previous studies in economics relied on stepwise heuristics to eliminate suboptimal counterfactual networks like Alder (2017) in India or Burgess et al. (2015) in Kenya, but did not include a derivation of the globally optimal network design. Once constructed, my network features trade on a two-dimensional lattice geometry and is hence related to the theoretical work of Allen and Arkolakis (2014, 2016). In using satellite data to construct a spatial representation of the economy, my study also adds to a vast literature ignited by Henderson et al. (2012) and surveyed in Donaldson and Storeygard (2016). I also contribute to the literature employing regional trade models to explain subnational welfare disparities caused by internal geography in a development context. Coşar and Fajgelbaum (2016) and Atkin and Donaldson (2015) analyse how exposure to international trade hubs propagates through the local topography in China and Ethiopia / Nigeria, respectively. In constructing the efficient network to mitigate these dispersions over space, I also contribute to new explorations into conditions and characteristics of optimal spatial policies (Fajgelbaum and Gaubert, 2018). In my three strands of empirical inquiry, I first add to the literature examining long-run persistence of colonial transportation revolutions in Africa (Jedwab and Moradi, 2016; Jedwab et al., 2017). I also contribute to the literature examining how ethnic relations relate to comparative development in Africa (Michalopoulos and Papaioannou, 2013, 2014, 2016) and add to our understanding of how ethnic (De Luca et al., 2018) and regional favouritism

(Hodler and Raschky, 2014; Burgess et al., 2015) skew public goods spending towards an inefficient allocation. Lastly, I contribute to the literature on the distribution and effects of foreign aid (Clemens et al., 2012; Nunn and Qian, 2014; Dreher and Lohmann, 2015; Dreher et al., 2017).

This paper will proceed as follows: section 2 presents a network model of trade which allows for solving for the optimal transport network. In section 3, I calibrate the model with rich data and derive the efficient trade network design for every country in Africa. Section 4 describes and quantifies spatial imbalances in the continent’s current trade systems. Section 5 then presents empirical strategy and results for investigating three potential sources of spatial inefficiency: colonial infrastructure investments, ethnic relations, and foreign aid. Section 6 concludes.

2 A model of optimal transport networks

In this paper, I derive the unique optimal transport network for every country in Africa. To be able to maximise over the space of networks, I harnesses an altered version of a framework by Fajgelbaum and Schaal (2017). The model is introduced in the following paragraphs.

2.1 Geography

Following the set-up and notation of Fajgelbaum and Schaal (2017), I consider a set of locations $\mathcal{I} = \{1, \dots, I\}$. Each location $i \in \mathcal{I}$ inhabits a number of homogeneous consumers L_i . This number is treated as given and fixed for every location, such that consumers are not allowed to move between locations. Each consumer has an identical set of preferences characterised by

$$u = c^\alpha$$

where c denotes per capita consumption. Every consumer in location i consumes c_i and $C_i = L_i c_i$ denotes total consumption in location i .

There is a set of goods \mathcal{N} denoted by $n = \{1, \dots, N\}$. Total consumption in each

location is defined as the CES aggregation of these goods

$$C_i = \left(\sum_{n=1}^N (C_i^n)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

where σ denotes the standard elasticity of substitution and C_i^n denotes the consumption of good n in location i . Locations specialise in the production of goods such that each location only supplies one variety $n \in \mathcal{N}$. Let $Y_i = Y_i^n$ denote total production in location i .

2.2 Network topography

Locations \mathcal{I} represent nodes of an undirected network graph. Each location i is directly connected to a set of neighbours $N(i) \in \mathcal{I} \setminus \{i\}$. I consider locations to be arranged on a two-dimensional square lattice where each node is connected to the locations in its *Moore Neighbourhood* (i.e. its eight surrounding nodes to the north, north-east, east, and so on). Nodes at the border of the network graph might have fewer than eight neighbours. Let \mathcal{E} denote the set of edges connecting neighbouring nodes and note that $(\mathcal{I}, \mathcal{E})$ fully describes the underlying network topography.

All goods can be traded within the network. Let $Q_{i,k}^n$ denote the total flow of good n travelling between nodes i and $k \in N(i)$. While goods can only be traded between neighbouring nodes, nothing prevents them from travelling long distances through the network by passing multiple locations after each other. Sending goods from location i to location $k \in N(i)$ incurs trade costs, which are modelled in the canonical iceberg form. I follow Fajgelbaum and Schaal and model iceberg trade costs for trading good n between neighbouring locations i and k as

$$\tau_{i,k}^n(Q_{i,k}^n, I_{i,k}) = \delta_{i,k}^\tau \frac{(Q_{i,k}^n)^\beta}{I_{i,k}^\gamma} \quad (1)$$

where $I_{i,k}$ is defined as the level of *infrastructure* on the edge between nodes i and k . More infrastructure on a given link decreases the cost of trading between them. $\delta_{i,k}^\tau$ is a scaling parameter, which allows trade costs to be flexibly adjusted for any given origin-destination pair. Trade costs also depend on $Q_{i,k}^n$, the total flow of goods on the link. Higher existing

trade volumes on a given edge make sending an additional good more costly, a dynamic Fajgelbaum and Schaal refer to as *congestion externality*. Sending one additional unit of goods from i to k makes all other existing shipments on that link more expensive. The social planner realises this and takes congestion into account when determining optimal trade flows.

In equilibrium, each location cannot consume and export more than it produced and imported. More formally

$$C_i^n + \sum_{k \in N(i)} Q_{i,k}^n (1 + \tau_{i,k}^n(Q_{i,k}^n, I_{i,k})) \leq Y_i^n + \sum_{j \in N(i)} Q_{j,i}^n \quad (2)$$

must hold for every n and i . equation (2) is a variation of what Fajgelbaum and Schaal call the *Balanced Flows Constraint*.

I follow the contribution of the Fajgelbaum and Schaal (2017) framework and proceed to endogenize infrastructure provision $I_{i,k}$ in order to facilitate optimal trade flows. Analytically, this problem nests the static trade flow exercise outlined above. The social planner chooses an infrastructure network, and given the network proceeds to compute the optimal trade flows subject to the *Balanced Flows Constraint* (2). To make the problem more interesting, I follow Fajgelbaum and Schaal in introducing a constraint on infrastructure. This is specified in fairly straightforward manner as the *Network Building Constraint*

$$\sum_i \sum_{k \in N(i)} \delta_{i,k}^I I_{i,k} \leq K \quad (3)$$

where $\delta_{i,k}^I$ denotes the cost of building infrastructure on the edge between nodes i and k . Total spending on infrastructure is constrained by K , the sum originally spent on building the *existing* road network of a country. I observe the current road network of the economy, infer how much it must have cost to build it, and set K equal to this amount. The social planner's task of choosing an optimal network hence amounts to a reallocation exercise. She gathers all road building material available in the economy and gets to redistribute it in a more sensible way. Improving infrastructure between two nodes in order to foster local trade hence comes at the cost of having to take away infrastructure elsewhere. I argue that the degree to which the social planner has to rearrange existing edges serves

as a sensible measure of spatial inefficiency in the existing network.²

2.3 Planner's problem and equilibrium

In the nested problem, the social planner observes localities, endowments, population, and preferences and solves for trade flows between nodes that maximise overall welfare. She also solves for the optimal transport network which induces welfare-maximising trade flows in the nested problem while respecting the *Network Building Constraint* (3). The full planner's problem can hence be stated as

$$\begin{aligned}
& \max_{\left\{ \left\{ C_i^n, \{Q_{i,k}^n\}_{k \in N(i)} \right\}_n, \right.} \sum_i L_i u(c_i) \\
& \quad \left. c_i, \{I_{i,k}\}_{k \in N(i)} \right\}_i \\
\text{subject to} \quad & L_i c_i \leq \left(\sum_{n=1}^N (C_i^n)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \forall i \in \mathcal{I} & \text{CES Consumption} \\
& C_i^n + \sum_{k \in N(i)} Q_{i,k}^n (1 + \tau_{i,k}^n(Q_{i,k}^n, I_{i,k})) \\
& \quad \leq Y_i^n + \sum_{j \in N(i)} Q_{j,i}^n, \forall i \in \mathcal{I}, n \in \mathcal{N} & \text{Balanced Flows Constraint} \\
& \sum_i \sum_{k \in N(i)} \delta_{i,k}^I I_{i,k} \leq K & \text{Network Building Constraint} \\
& I_{i,k} = I_{k,i}, \forall i \in \mathcal{I}, k \in N(i) & \text{Infrastructure Symmetry} \\
& C_i^n, c_i, Q_{i,k}^n, I_{i,k} \geq 0, \forall i \in \mathcal{I}, n \in \mathcal{N}, k \in N(i). & \text{Non-Negativity}^3
\end{aligned}$$

My version of the planner's problem follows the baseline Fajgelbaum and Schaal (2017) model. However, my model has four important differences. First, in my model all goods are tradeable and no local amenities exist. Second, I do not allow workers to migrate between places and hence differences in marginal utility might still exist between nodes. Third, my model remains agnostic about the production function of each location and no analysis of the optimal use of input factors is undertaken. Fourth, I impose infrastructure symmetry. All these changes are undertaken with the later calibration and reshuffling

²I also impose infrastructure symmetry and restrict $I_{i,k} = I_{k,i} \forall i, k \in N(i)$.

³As will be discussed in chapter 3, I calibrate my version of the model with an even stronger lower bound to infrastructure $I_{i,k}$ than mere non-negativity. For reasons discussed below, I simulate the model while binding $I_{i,k} \geq 4$. For all other variables, merely non-negativity is required.

exercise in mind.

While optimising over the space of networks might appear daunting, Fajgelbaum and Schaal (2017) provide conditions under which deriving the unique spatial optimum is both ensured and feasible. Instead of solving for every single infrastructure link, I follow the authors and recast the problem in its dual representation as a set of first-order conditions from the subproblems, which only depend on Lagrange multipliers of each constraint. There are considerably fewer multipliers than primal control variables, namely one for every good in every node (interpretable as local prices). I am hence left to only find a price field from which under the convexity assumptions, all other properties follow.⁴ As spelled out more formally in the technical appendix (see section A), I obtain the optimal network by constructing the Lagrangian corresponding to the planner’s problem, deriving its first-order conditions, and recasting them as functions of the Lagrange parameters. Numerical optimisation now yields the solution to the dual problem and inserting the parameters back into the derived first-order conditions, I can immediately derive the optimal infrastructure network $I_{i,k}$, optimal trade flows $Q_{i,k}^n$ over this network, and ensuing consumption patterns C_i^n in each location.

3 Deriving optimal trade network designs

To calibrate a topography of economic activity and trade in all African countries, I construct a novel network representation covering the entire continent and enrich its nodes and edges with data from a variety of sources.

3.1 Network nodes

I first divide the entire continent into grid cells of 0.5 degrees latitude by 0.5 degrees longitude (roughly 55 by 55 kilometres at the equator). For all of Africa, this amounts to 10,167 cells. Using GIS, I locate the geometric centroid of each cell and overlap it with current political borders to assign countries to each centroid. I then aggregate spatial data on economic and geographic characteristics onto this grid cell level. Raster

⁴It is still a quite demanding task to solve the ensuing dual problem, even numerically. Invoking duality reduces the scale of the problem, but I am still left with optimising over $I \times N$ variables.

data on 2015 population totals come from the *Gridded Population of the World* dataset (GPW, Socioeconomic Data and Applications Center, 2016). This NASA-funded project gathers data from hundreds of local census bureaus and statistical agencies in order to construct a consistent high-resolution spatial dataset of the world’s population. When a datasource only reports population totals for large, higher-level administrative districts, the dataset smoothes population uniformly over the entire area.⁵ Africa is the continent with the coarsest resolution of administrative input data. However, the average coverage of $(57\text{KM})^2$ neatly matches the grid cell resolution of my study. I overlay the GPW data with my 10,000+ grid cells to obtain the total number of people living in each cell. On average, a cell is home to 110,000 people, with the median much to the left of that (25,000). The most populous cell contains Cairo and inhabits almost 18 million people. 212 cells are uninhabited. To proxy for heterogeneities in economic activity over space, I rely on the established practise of using satellite imagery of light intensity at night (Henderson et al., 2012). Pre-processed data on 2010 night luminosity come from Henderson et al. (2018) and are also aggregated onto my study’s 0.5×0.5 degree grid resolution.

3.2 Network edges

To quantify the degree to which network nodes are connected to each other, I make use of the open source routing service OPEN STREET MAP (OSM). For every centroid location, I scan OSM for the optimal route to each of their respective eight surrounding neighbours. Since I am interested solely in within-country transport networks, I perform the exercise for each country separately and do not elicit connections between locations of different countries. Hence, centroids located near the coastline or national borders often have less than eight immediate neighbours. For all of the resulting almost 75,000 routes, I gather distance travelled, average speed, and step-by-step coordinates of the travel path.⁶

The OSM routing algorithm is specified for cars and takes into account differential

⁵GPW does not employ any auxiliary data sources – like satellite data – to weight-adjust population totals over space (Doxsey-Whitfield et al., 2015).

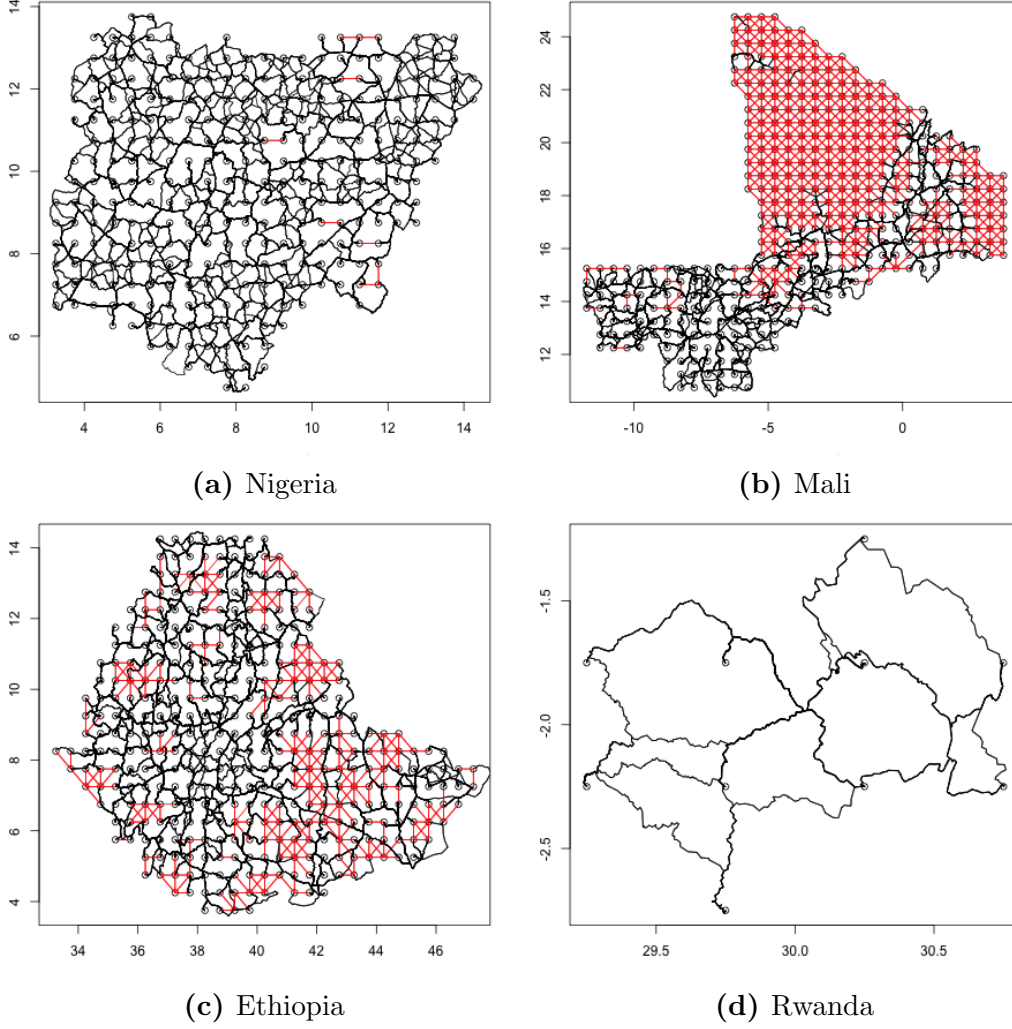
⁶Scans of OSM were conducted in November 2017. The service does not allow a retrospective scan over past road databases, so a time difference between lights (2010), population (2015), and roads (2017) can not be overcome. If anything, this renders my network inefficiency measure a lower bound to its true value since government officials will have had time to adjust their network to any spatial economic imbalances. If the 2017 network only inefficiently supports 2010-15 trade, chances are it did even worse in 2015.

speeds attainable on different types of roads. However, if either start or destination location do not directly fall onto a street, the optimal route jumps to the nearest road and goes from there. To take this into account, I add a walking distance to the travel path. Agents are assumed to walk in straight lines to the nearest street at a fixed speed of 4 km/h. They then take the car and drive the route with average speed as specified by OSM, before they potentially have to walk the last stretch again to their exact centroid destination. For some particularly remote areas, even the nearest street is very far away, such that the car routing provided by OSM is not sensible. To counter these cases, I also calculate for all 70,000+ connections the outside option of walking the entire link in a straight line at 4 km/h. I then identify cases in which walking directly is actually faster than using OSM's proposed route (plus the travel to and from roads). In these cases, I replace OSM's route with the walking distance and constant 4 km/h speed.

Figure I presents the resulting road networks for four countries. Figure Ia displays every optimal route for Nigeria, which appears overall fairly well connected. Commuters mostly seem to be able to drive relatively direct routes between locations, even though cases with substantial detours are also evident at second glance. Connections in which walking were the preferred alternative are displayed in red and fairly rare in Nigeria. Figure Ib presents the case of Mali, which paints a different picture: for many connections through the Sahara desert in the north-east of the country, walking straight lines in the sand is actually the fastest way to get from A to B. Ethiopia in Figure Ic displays only a few trails connecting the country's east to the west. Small Rwanda in Figure Id zooms in on the actual roads taken and displays the intricacies of the optimal routing provided by OSM.

Relying on the open source community of OSM does come with some drawbacks. The most pressing concern is that data on the position and quality of roads are user-generated and hence subject to reporting bias. Richer areas may appear to be equipped with more roads if local residents have the time and necessary access to a computer to enter their neighbourhoods into the database. As soon as inference is conducted on the relationship between streets and any covariate of development, the resulting estimates will be biased. While this is certainly troubling, I believe this bias to be much more important on finer

Figure I: Road networks for different countries as scanned off OSM



Road networks as scanned off Open Street Map (OSM). Black lines represent optimal routes from each grid cell centroid to each of its eight surrounding neighbours. These routes may include a portion walked by foot in order to get to the nearest street. Connections in which walking the entire distance is faster are printed as red straight lines. Axes denote degrees longitude (x) and latitude (y), respectively. Data scanned in November 2017.

resolutions than the operating one in this study. Start and destination of the elicited routes are on average more than 55 kilometres apart and travel will hence take place mostly on larger roads and national highways. It is unlikely that these major streets are systematically underreported in OSM, the primary open source routing platform on the internet. It is nevertheless important to keep this potential flaw of the data in mind when conducting inference later on.⁷

⁷In some rare cases (less than 0.1 per cent of all connections), the OSM algorithm cannot find any route between two neighbouring centroid locations. This is mostly due to an obvious geographic impossibility to connect two nodes. In Guinea-Bissau, for instance one location lies on the Bolama Islands just off the shore of mainland Guinea-Bissau. Its neighbouring locations are all on the mainland and hence unreachable by car. In other cases, both locations to be connected are in deep jungle or swampy regions. In all these cases, I treat the link as if the two locations were not neighbours in the first place. That implies I even forgo the backup possibility of walking the entire distance, assuming that agents cannot walk between islands or through the densest jungle.

I use the average attainable speed between locations according to the OSM algorithm as a proxy for the quality of current infrastructure on the edge between them. If two locations are linked by a faster connection, I assume this to be the result of higher infrastructure $I_{i,k}$ on this edge. I hence set

$$I_{i,k} = \text{Average Speed}_{i,k} \quad (5)$$

This measure is naturally bound from below at 4 km/h, as walking the air-line distance is always available as a backup. Empirically, average speeds range between 6 km/h (Mauritania, where most routes go through the desert and have to be covered by walking) and 33 km/h (Swaziland).

To parametrise iceberg trade costs defined in equation (1), I follow Fajgelbaum and Schaal and set $\beta = 1.245$ and $\gamma = 0.6225$. I calibrate $\delta_{i,k}^\tau$ by harnessing a recent contribution by Atkin and Donaldson (2015). They show that trade costs are significantly increasing in (log) distance between origin and destination, impeding much of mutually beneficial trade. Directly taking the average of the authors' two point estimates for Ethiopia and Nigeria, I calculate

$$\delta_{i,k}^\tau = 0.0466 \times \ln(\text{Distance}_{i,k}) \quad (6)$$

as the trade cost elasticity to distance travelled.⁸

$\delta_{i,k}^I$ from equation (3) denotes the relative constant cost of increasing the average speed on a given link by one. I follow Fajgelbaum and Schaal who in turn make use of a recent study by Collier et al. (2015), which estimates infrastructure building costs in developing countries. Readily applying their specification, I calculate

$$\ln(\delta_{i,k,c}^I) = \delta_c^I - 0.11 \times \mathbb{1}_{\text{Distance}_{i,k,c} > 50km} + 0.12 \times \ln(\text{Ruggedness}_{i,k,c}) + \ln(\text{Distance}_{i,k,c}) \quad (7)$$

as the constant cost of increasing infrastructure $I_{i,k}$ on the link between i and k in country c . $\text{Distance}_{i,k,c}$ denotes the road distance travelled between nodes and enters pos-

⁸Atkin and Donaldson (Table 2, page 44) estimate the coefficient as 0.0374 for Ethiopia and 0.0558 for Nigeria. My parameter is the simple average of these two point estimates.

itively, implying that longer roads are costlier to develop as every single road kilometre will have to be improved. Moreover, the building cost per kilometre falls discretely when the distance surpasses 50 kilometres, as embodied by the indicator $\mathbb{1}_{\text{Distance}_{i,k,c} > 50\text{km}}$. Note that every route in my sample is longer than 50 kilometres and the corresponding dummy term is hence always equal to 1. $\text{Ruggedness}_{i,k,c}$ denotes the average ruggedness between grid cells i and k and enters positively, highlighting the additional expenses accompanied with building on uneven terrain.⁹ δ_c^I is a country-specific scaling parameter. Its main purpose is to ensure that equation (3) is satisfied with $K = 1$. I first appraise the infrastructure network $I_{i,k}$ of all countries and then flexibly alter δ_c^I for each nation individually in order to comply with equation (3).

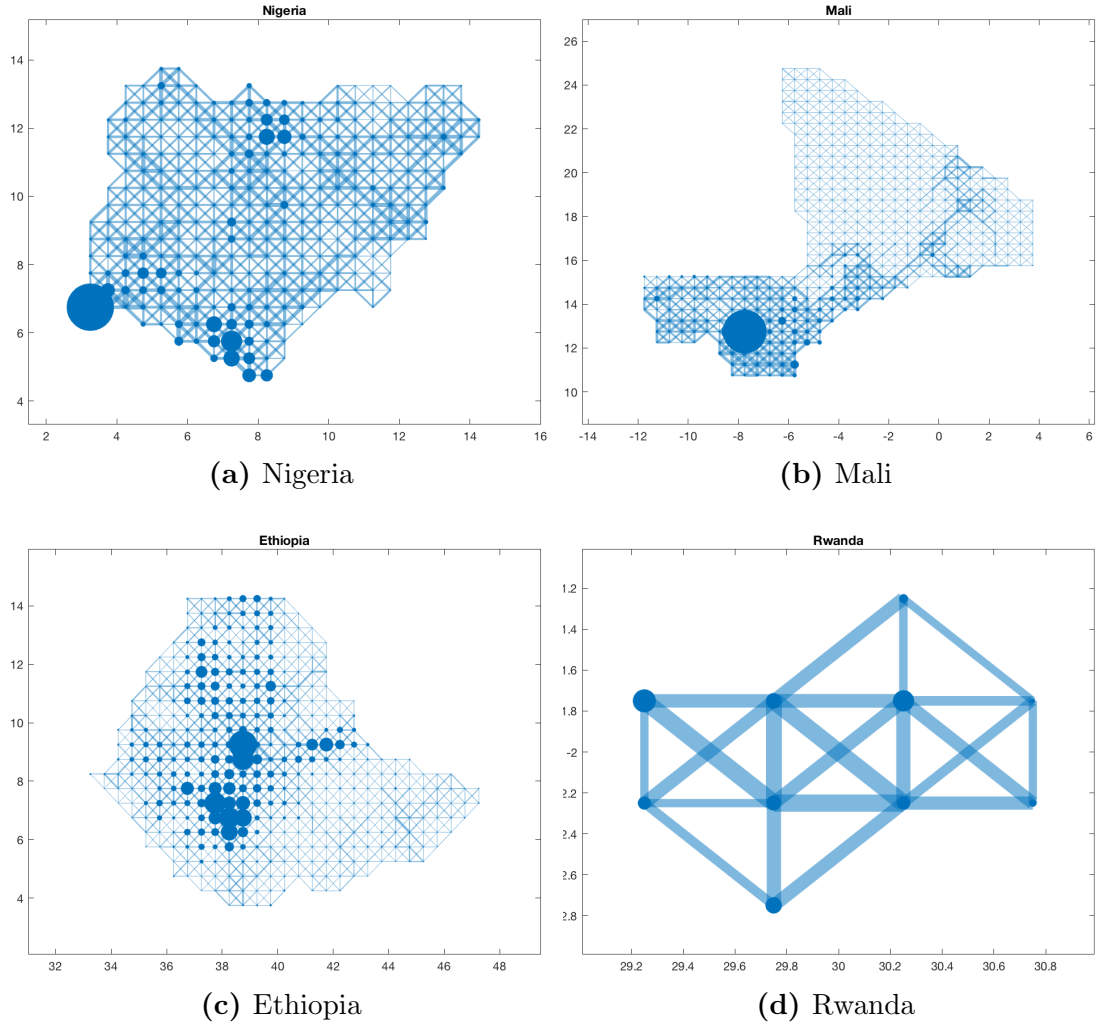
To build incentives for trade, I introduce $N = 2$ different goods: an agricultural and an urban good. To classify grid cells as urban or rural I use an iterative procedure, which seeks to match each country’s 2016 urbanisation rate as reported by The World Bank (2017). I start by assuming every location is a city and then gradually proceed to re-classify the least densely populated locations, until the ratio of people living in urban areas to total population equals that of the World Development Indicators.¹⁰ With this procedure, seven per cent of grid cells are classified as urban. These cells inhabit 40 per cent of the continent’s population, matching recent figures from Lall et al. (2017) fairly well.

After these steps, a discretised network representation exists for every African country. Nodes in the network are the spaced centroid locations of each grid cell. They combine the characteristics of the entire grid cell (population, output, etc.) in one point. Edges in the network are road connections between centroids. Each edge also carries a number of characteristics (average speed, trade costs, and infrastructure building costs). Figure II presents this discretised network representation for the four countries from above. Nodes are printed larger proportional to their population. Edges are drawn thicker proportional to the initial infrastructure investment.

⁹Data on local ruggedness come from Henderson et al. (2018) and is described in more detail with other geographical covariates below.

¹⁰For three countries, the WDI do not report urbanisation rates. In these cases, I match the overall urbanisation rate for the entire African continent of 42 per cent as reported by Lall et al. (2017).

Figure II: Discretised networks for different countries



Discretised representation of the infrastructure networks from Figure I. Nodes are drawn with radius proportional to total population, edges' thickness correspond to average attainable speed on the route connecting the nodes.

3.3 Trade network optimisation

For each country, I conduct two simulations. In both exercises, I calibrate the curvature parameter of the utility function at $\alpha = 0.4$ and the elasticity of substitution parameter at $\sigma = 4$. In the first simulation exercise, infrastructure $I_{i,k}$ is treated as fixed. This is to obtain a baseline estimate of the spatial variation of welfare in each country. Locations are still allowed to trade with each other, but only over the exogenous current road network. By construction, the resulting solution will have two properties. Firstly, total output over the entire country will remain untouched. Inputs are not defined and hence do not shift to more productive regions. Indeed, any welfare gains will be attained solely by shipping the right mix of goods to the right regions. Secondly, labor immobility will leave welfare

differences between regions as agents cannot simply move to more privileged cells. The social planner would like to overcome these differences, but is confronted with trade costs which might leave certain remote areas much worse off than well-connected ones.

Following this static exercise, I proceed to the main task of endogenizing the infrastructure matrix $I_{i,k}$. With the *Network Building Constraint* binding total infrastructure investment at the level of the current road network, the social planner is now free to reshuffle roads within the country in order to improve connections as she chooses. If she wants to improve the connection between two given locations, she will have to take away infrastructure from somewhere else in the country. This reallocation exercise does not seek to identify where to place the optimal next investment, but rather represents an utterly fictitious scenario in which every road can be lifted from the ground, reshuffled, and eventually located someplace else.¹¹ The procedure does not measure how many roads a country has, but rather how well they are placed. It does not look at whether the entire country is full of speedy roads, but rather whether those roads connect the right locations.

I conduct the reallocation scenario for every African country. Six small countries (Cape Verde, Comoros, The Gambia, Mauritius, São Tomé and Príncipe, and Reunion) are too small to form a sensible network as they only show up as a single location in the dataset and are henceforth no longer considered. Optimisations are performed via MATLAB's `fmincon` command. When conducting the simulations, I bind the social planner's set of permissible roads from below, at 4 km/h (such that $I_{i,k} \geq 4 \forall i \in \mathcal{I}, k \in N(i)$). This is motivated by the assumption at the beginning that walking straight lines at this speed is conceived as an outside option and always available to any commuter. The social planner should not be able to force commuters to travel slower than walking in order to build a faster road elsewhere.¹²

Figure III visualises this reallocation exercise for several countries. Subfigure IIIa dis-

¹¹Note that equation (3) only fixes $\sum_i \sum_{k \in N(i)} \delta_{i,k}^I I_{i,k} = K$. Hence, not the overall sum of infrastructure is fixed, but more precisely the overall cost of infrastructure. This still allows the social planner to take away one unit of infrastructure on a very expensive (high $\delta_{i,k}^I$) link and exchange it for much more than one unit on a cheaper (low $\delta_{i,k}^I$) link.

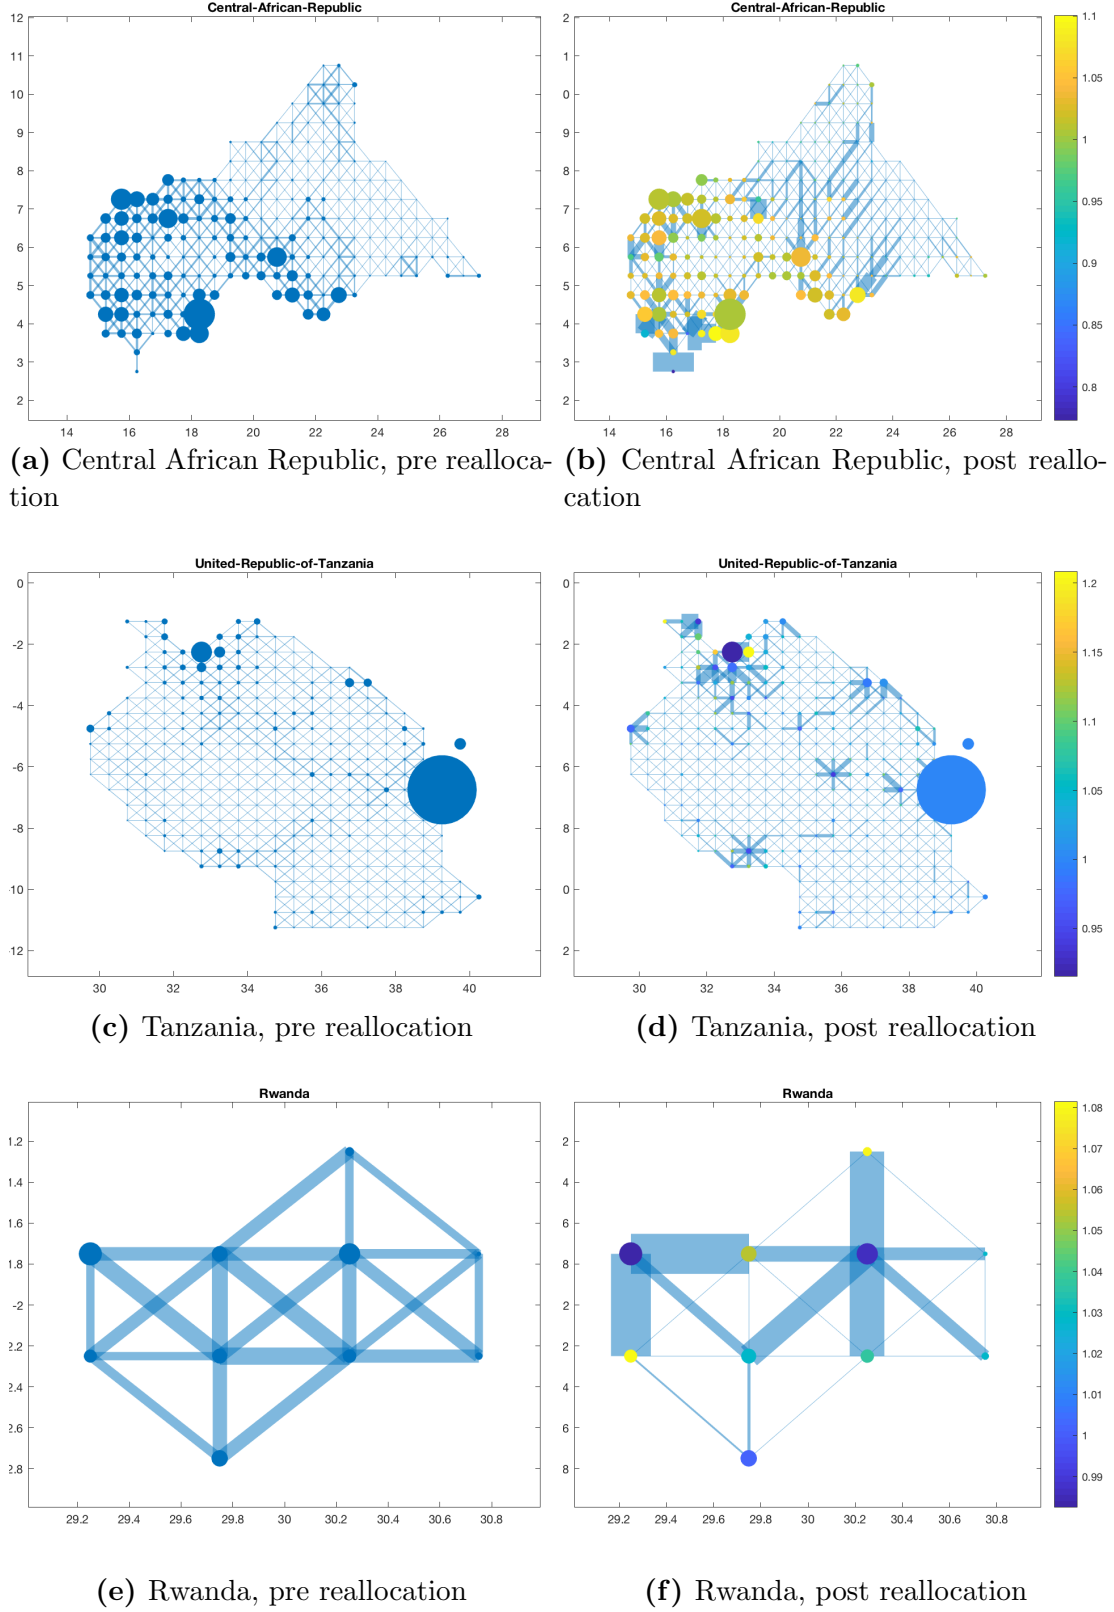
¹²Contrarily, I do not explicitly restrict possible investments from above (at least not in addition to the sum-restriction imposed by equation 3), as this could violate the strong convexity of the problem. Not bounding the problem in principle allows the social planner to combine every available infrastructure from all over the country into one supersonic speed highway on one particular edge. However, the model is calibrated in a way which makes this very unattractive to the planner anyway. After simulating reallocation in every African country, less than 0.8% of all 70,000+ built roads were suggested to be over 260 km/h. Still, one outlier of 2007 km/h (in Egypt) and one of 1755 km/h (in South Africa) remain.

plays the discretised network representation of the Central African Republic, comparable to Figures IIa – IIId. The edges to this network are printed almost evenly thick, implying that infrastructure is fairly evenly distributed across the country. Subfigure IIIb then displays the country after the network reshuffling exercise. Three patterns stand out. First, the social planner sees a clear need to connect the populous areas in the south-west of the country with each other. Some southern nodes are granted extensive, almost highway-like connections to their immediate neighbours. For that, the social planner is willing to salvage some of the apparently unnecessary infrastructure in the middle or north of the country. Second, there still is a benefit to having a few trails connecting the south-west with the north-east. Some clear north-south and east-west routes spanning multiple regions emerge. Thirdly, nodes are printed in a colour scale corresponding to individual welfare gains and losses for each location. As can be seen from first-glance, most southern regions stand to gain between five and ten per cent of total welfare from this scenario. Hardly any nodes seem to lose welfare, even though on second glance a few instances become apparent.

Tanzania in Figures IIIc – IIId displays a more decentralised optimal network solution. The reallocation scenario results in the main urban areas being better connected to their immediate surroundings, but no clear overarching network emerges. There also does not appear to be any necessity to better connect hinterland regions with the primal city Dar es Salaam in the east. Indeed, the largest city slightly loses welfare with the reallocation at the expense of multiple smaller population centres in the north.¹³ Small Rwanda in Figures IIIe – IIIf helps to illustrate some of the forces at hand in a less crowded graph. Starting from a fairly evenly distributed transport network, the reallocation dynamics lead to much more variation in infrastructure provision. Some links are deemed superfluous and hence reduced to the smallest admissible level, while others are scaled to multiple times their starting infrastructure stock. Furthermore, high welfare gains are reported by direct neighbours of big production centres and urban grid cells. These are unassuming grid cells with average population or output levels, merely equipped with the geographical

¹³On a side note, Tanzania also illustrates an interesting case where a tiny fraction of the country is fully detached from the rest of the network. Just north of Dar es Salaam, the island of Zanzibar constitutes a one-node subnetwork of its own. Not surprisingly, it remains completely unaffected by the reshuffling of roads on the mainland. Instances like these are relatively common in the dataset.

Figure III: Reallocation scenario for different countries



Results from optimally reshuffling roads in three African countries. In each network graph, every node represents a grid cell centroid location with radius proportional to the size of its local population. Edges are drawn thicker depending on their allotted infrastructure $I_{i,k}$ (i.e. average attainable speed). In the optimal networks on the right, nodes are coloured based on their relative welfare gains and losses. Note the slightly different color scales for each country.

blessing of being close to a bigger neighbour. This leads to the conclusion that while better infrastructure combats the welfare costs of geographical distance, proximity to hubs still matters. Even an optimally designed transport network is ultimately not able to fully overcome the *curse of distance*.¹⁴

4 A measure of spatial transport network inefficiency

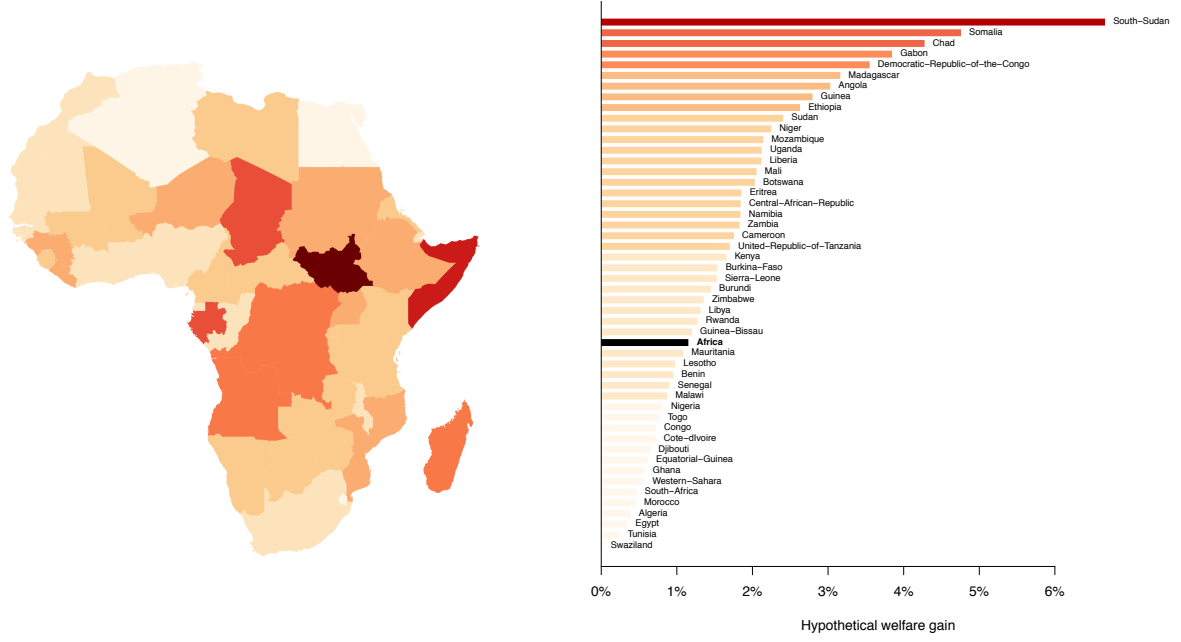
After successfully reshuffling a country’s transport network, overall welfare will necessarily (weakly) increase. It is the social planner’s objective to maximise overall welfare, and since the original network composition is always still available, the entire country cannot on aggregate be worse off than before. Note again that overall production (light output) will be unaffected by the entire exercise. Welfare gains are solely caused by enabling mutual benefits from trade through connecting the right locations. Nevertheless, they are substantial. The Central African Republic of Figure III, for instance, stands to gain 1.84% of overall welfare just by reshuffling its roads. Tanzania (1.7%) and Rwanda (1.27%) are slightly closer to their hypothetical optimum.

Figure IV displays all African countries and their hypothetical welfare gain. The three countries from above perform rather well in comparison. Some nations like South Africa (0.5% welfare gains) or Tunisia (0.2%) perform even better. Many countries are leaving much more on the table, like Somalia (4.8%) or Chad (4.3%). No African country, however, has a more ill advised transport network than South Sudan. Its citizens stand to gain almost 6.7% of overall welfare if just their roads were better placed. This might not come as a surprise, as the world’s newest country has largely inherited a road network that was not conceived to sustain an independent nation, but rather connect it to its former capital up north. For the entire continent, optimal reallocation of national road systems would improve overall welfare by 1.15%.

Forgone welfare gains can be conceived as an intuitive measure for overall network inefficiency. The closer hypothetical gains to zero (the lighter the country’s colour), the more efficient the current allocation of roads. Vice-versa, if a country stands to gain a lot from reshuffling, then the current network is deemed more inefficient. On a simple

¹⁴A term coined by Boulhol and de Serres (2010).

Figure IV: African countries by network inefficiency



Countries coloured according to their welfare gain under the optimal reallocation counterfactual. Scale from almost 7% (dark red) to < 1% (white) welfare gains. Gains are computed by comparing the population-weighted sum of utility levels over all grid cells in a country before and after the reallocation exercise.

cross-section, countries with less efficient networks are significantly correlated with more corruption ($p < 0.01$), less property rights ($p < 0.01$) and less 2010 log GDP ($p = 0.07$).¹⁵

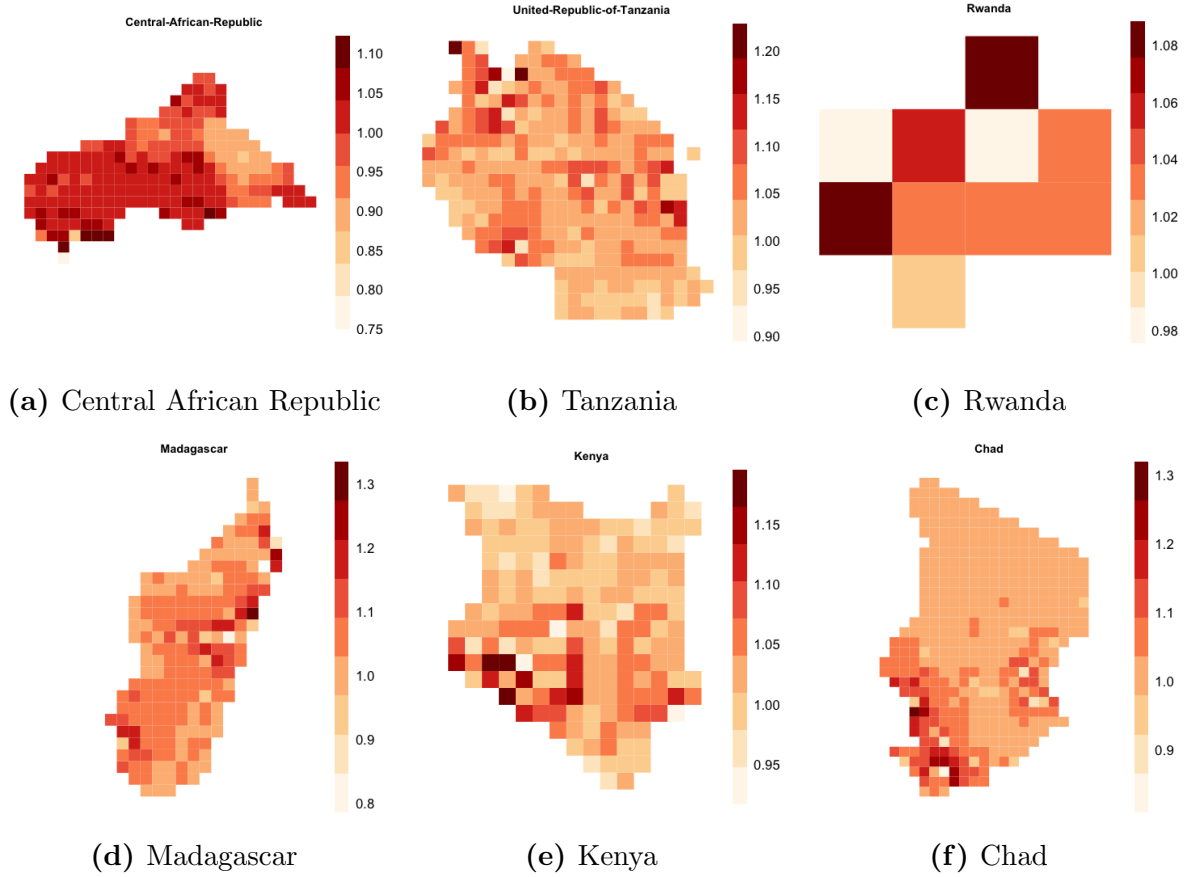
While each country only stands to gain overall welfare from the reallocation procedure, individual locations might very well lose in the process. Intuitively, some regions might be equipped with far too many good roads such that the social planner takes these roads away to use someplace else. Comparing each grid cell's welfare before and after the major reshuffling can help to identify regions which are currently over or underprovided for. More formally, I define

$$\Lambda_i = \frac{\text{Welfare under the optimal Infrastructure}_i}{\text{Welfare under the current Infrastructure}_i} \quad (8)$$

as the *Local Infrastructure Discrimination Index* for grid cell i . Areas with high Λ_i scores ($\Lambda_i > 1$) would be gaining under the optimal reallocation scenario and are hence under provided for in the network's current state. A score of $\Lambda_i < 1$ on the other hand,

¹⁵Data from The World Bank (2017). For corruption and property rights, data are only available for 35 countries and correlations are hence performed on this truncated sample. Interestingly, network efficiency is *not* statistically associated with earlier independence years ($p = 0.8$) or more artificial border designs ($p = 0.3$) as reported by Alesina et al. (2011).

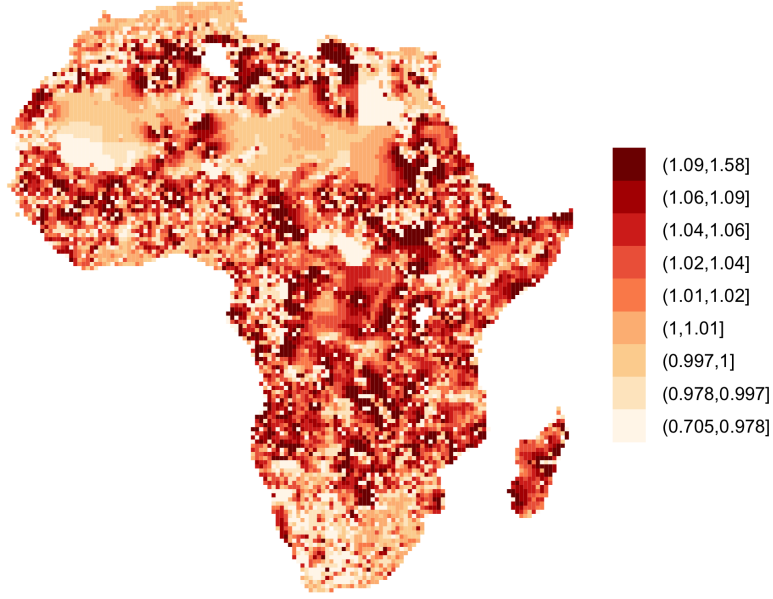
Figure V: Spatial distribution of Λ_i for sample countries



Six African countries by their Local Infrastructure Discrimination Index Λ_i on the grid cell level. Maps show each country as the 0.5×0.5 degree grid used for the network optimisation. For each map, darker shaded cells correspond to higher Λ_i levels and hence more infrastructure discrimination compared to the optimal network. To better visualise within-country variation, colour scale slightly changes from country to country.

implies that a region is *too* well off given its position in the network today and hence should be stripped off some of its infrastructure to increase overall welfare. Figure V displays the spatial distribution of Λ_i for six countries. The darker a grid cell's shade, the more it is disadvantaged by the inefficiencies of the current network. Figures Va – Vc display what could already be inferred from the colouring of the nodes in Figure III: The Central African Republic has not enough fast roads in the south-west and too many in the east, Tanzania shows no clear spatial pattern, and Rwanda only engages in minor reshuffling. In Figures Vd – Vf it, furthermore, becomes apparent that Madagascar's infrastructure network discriminates against the island's heartland (note how the coastal areas tend to be much lighter than the hinterland), Kenya would profit from connecting Nairobi to its surroundings, and Chad's south is discriminated against compared to the north.

Figure VI: Spatial distribution of Λ_i for entire sample



Africa as represented by 10,158 grid cells of 0.5×0.5 degrees. Cells coloured according to their Local Infrastructure Discrimination Index Λ_i . Darker cells would benefit from reallocating national infrastructure networks. Note that the hypothetical reallocation scenario is conducted on the country-level. Cells from different countries are hence not immediately comparable. Map's colouring follows an equal-interval rule such that every colour in the spectrum has an equal amount of members. This is to visualise the measure's variation but leads to unequal bracket-sizes for each colour.

In interpreting Λ_i , keep in mind that this is a measure of differences in *welfare*. In this, it need not be a direct mapping of changes in actual infrastructure provision. Indeed, the highly non-linear nature of the optimal reallocation scenario can lead to situations in which a certain region substantially profits from the optimal policy, even though it is not directly granted additional roads. Local changes in welfare can instead be caused also by fortuitous peculiarities of geography – maybe a neighbouring region emerges as a local trade hub, or the optimal network leads to improvements in the variety of goods, all without directly targeting each individual grid cell with additional roads. In my full dataset, changes in welfare Λ_i are positively correlated with changes in infrastructure, yet the opposite is also true for some outliers.¹⁶ In the remainder of this study, I will refer to high Λ_i values as implying that a region is being awarded additional infrastructure from the social planner, even though the two need not necessarily be always equivalent.

Figure VI displays the spatial variation of Λ_i over all 10,000+ grid cells of the entire continent. When interpreting this map, note that grid cells are undergoing the reshuffling scenario solely within their respective country. National borders hence play a role and

¹⁶Infrastructure changes computed as $\sum_{k \in N(i)} I_{i,k}^{\text{opt}} / \sum_{k \in N(i)} I_{i,k}^{\text{empirical}}$ correlate with Λ_i at $p < 0.01$

can at times even clearly be inferred from the printed map.¹⁷ Keeping this in mind, the map reveals substantial spatial variation in the index across the African continent. The luckiest region (in Namibia) stands to lose almost 30% of total welfare if the fictitious social planner intervened and reshuffled roads away from it. On the other extreme of the spectrum, the residents of one grid cell in Gabon are missing out on a welfare hike of more than 50%. On average, a grid cell gains 2.8 per cent of welfare, the median cell gains 1.7%. Figure A.1 in the appendix plots the distribution of Λ_i , which roughly follows a normal distribution. In Figure VI, abandoned regions are clearly displaying spatial correlation with large neighbouring swaths of land collectively missing out on infrastructure improvements in certain countries. This begs the conclusion that countries do not just overlook single grid cells but rather live with vast stretches of disadvantaged regions. The index is evidently representing more than just haphazard noise. In the following sections, I analyse patterns behind this heterogeneity of infrastructure discrimination over space.

5 The determinants of Africa’s spatial trade network imbalances

Why are some African roads not in the right place to promote beneficial trade? To investigate which areas have too much or too little infrastructure, I employ the *Local Infrastructure Discrimination Index* Λ_i as dependent variable in a standard OLS regression setting. In the base specification, I estimate

$$\Lambda_{i,c} = \beta v_{i,c} + \mathbf{X}_{i,c}\gamma + \delta_c + \epsilon_{i,c} \quad (9)$$

¹⁷There are two reasons why I conduct the simulation procedure within countries and not over the entire African continent. One is computational; the requirements for numerically solving the model increase quadratically in the number of locations \mathcal{I} . The largest country in Africa (Algeria) is made up of almost 900 locations and already strains computing power quite heavily. Simulating all of Africa’s 10,000+ locations at once is then almost unattainable with available technology. The second reason is interpretational; while lifting a country’s roads from the ground and flexibly reshuffling them across the nation is already a fictitious scenario, it still operates within a government transport authority’s locus of control. Regions disadvantaged by their own government can reasonably be considered discriminated against. This is less the case if one were to optimise over the entire continent. Without a central planning body for all of Africa, it is hard to interpret why a road in e.g. Tunisia should rather be moved into Namibia.

for a variety of different independent variables $v_{i,c}$ of grid cell i in country c . δ_c denotes country fixed effects, $\mathbf{X}'_{i,c}$ is a vector of controls, and β is the coefficient of interest. The dependent variable $\Lambda_{i,c}$ roughly follows a normal distribution and I hence do not transform it. As apparent in Figure VI, local infrastructure discrimination displays autocorrelation over space causing the error term $\epsilon_{i,c}$ to not be distributed independently. To account for this problem, I follow Bester et al. (2011) and construct a higher-level spatial grid of 3 degrees latitude by 3 degrees longitude and cluster standard errors within each of these higher-level grid cells. Errors are allowed to covary within each cluster, but not between them.¹⁸

Including country fixed effects is of particular importance, as Λ_i is constructed by optimising trade flows within each country separately. Larger, wealthier, or more urbanised countries have more flexibility in reallocating their transport network. A grid cell in Egypt is hence not directly comparable to one in Sierra Leone. Country fixed effects account for this underlying heterogeneity and make observations comparable internationally. The vector of controls $\mathbf{X}'_{i,c}$ captures observable characteristics of each grid cell that plausibly account for some of the variation in Λ_i . Henderson et al. (2018) show that a surprisingly parsimonious set of geographical and agricultural covariates explains a substantial part of the global variation in economic activity. Making use of their data, I include in \mathbf{X} each grid cell's average altitude, average temperature and precipitation, land suitability for agriculture, length of the annual growing period, and an index for the stability of malaria transmission. I also include mutually exclusive (and collectively exhaustive) dummy variables classifying each grid cell into one of twelve predominant vegetation regions (or *biomes*, see Henderson et al., 2018).¹⁹ To flexibly control for any broad geographic trend over the entire continent, I additionally add fourth-order polynomials of both latitude and longitude for each grid cell. To take into account that some regions have a natural advantage in conducting trade, I also include indicators for whether a grid

¹⁸This technique draws its power from constructing clusters in the most arbitrary manner possible without relying on potentially endogenous partitions like national borders or administrative units (see e.g. Michaels and Rauch, 2017). There are 332 such clusters. Since 3 degrees is evenly divisible by the observation-level grid cell size of 0.5 degrees, each cluster in principle fits 36 observations. The median cluster does indeed comprise 36 cells, but some border-regions fit fewer observations. On average, there are 31 observations in a cluster.

¹⁹Only eight of the twelve vegetation patterns are actually present on the African continent, the other indicators (biomes 4, 6, 8, and 11) are dropped from consideration.

cell’s centroid is within 25 kilometres of a natural harbour, big lake, or navigable river respectively (again using data from Henderson et al., 2018). Lastly, to account for potential gravitational trade forces from abroad, I create and include a dummy for whether a cell is at the border of a country’s network and hence has less than eight immediate neighbours.

I call the set of controls outlined so far “*Geographic Controls*”. They are in principle unaffected by human decisions about the design of trade networks and therefore plausibly exogenous. Another set of covariates, however, poses more difficulties. These are the variables that were already used to calibrate the optimal reallocation simulation from above, namely a cell’s population, light output, ruggedness, and classification into urban and rural.²⁰ I call these “*Simulation Controls*”. It is crucial to be aware that Λ_i is, among others, already a product of the intricate interplay between these factors. There is hence a danger for plain OLS to detect a spurious, mechanical relationship between them, potentially biasing results. On the other hand, not controlling for the spatial distribution of people and economic activity creates the risk of confounding estimates by means of omitted variable bias. To confront this dilemma, I always report estimates with and without the set of simulation controls. As I will demonstrate, results turn out to be largely similar between the two, hinting at the highly non-linear genesis of Λ_i . Table A.1 in the appendix prints basic correlations of my various control sets with the outcome Λ_i . On average, my model reallocates roads towards border cells, as well as colder and more malaria-prone areas, and takes infrastructure away from more illuminated grid cells. Ruggedness, population, the classification of urban and rural, or fourth-order geographic trends do not systematically predict variations in Λ_i .

My measure of network inefficiency only pertains to systems of goods trade. As discussed above, there are clearly other rational motivations for building roads, mainly facilitating the commute of people to large administrative hubs – most immediately, a nation’s capital. To ensure that my results are not driven by systematic ignorance of these not-for-trade roads, I re-estimate every model while excluding grid cells containing a country’s capital. Unless otherwise reported, this leaves results virtually unchanged.

²⁰Recall that population and output were components of the planner’s problem, ruggedness went into the cost of building new infrastructure $\delta_{i,k}^f$, and the urban/rural classification determined which good a cell produced.


In this paper, I investigate three potential sources of network inefficiency in Africa: colonial era infrastructure investments, ethnic power relations, and foreign aid. Each individual setting faces different challenges to identification but they are all based on the framework outlined above. The following sections present empirical strategies and results for all three strands of inquiry.

5.1 Colonial infrastructure investments

The colonial powers transformed the landscape of many African regions by devising large scale infrastructure projects. Starting in the late 19th century, numerous railway lines were built to facilitate the transport of goods and troops through the vast newly appropriated territories. Between 1890 and 1960, British, French, Belgian, German, Italian, and Portuguese administrations all undertook efforts to permeate their colonies with more or less sophisticated railway networks (Jedwab and Moradi, 2016). There were two main motivations for this: supporting the extractive economies and ensuring military domination (Jedwab et al., 2017). These colonial railroads have been found to have a persistent impact on the spatial organisation of economic activity today. Jedwab and Moradi (2016) show how urbanisation started to centre around railway tracks in the decades following their construction. Even as most railway lines have fallen into disarray and road traffic has replaced trains as the most important means of transportation, economic activity today still clusters in places close to the former rail lines.

Did the transport revolution coordinate the economy on an efficient spatial equilibrium? To investigate whether railways from the colonial period still have an impact on trade network inefficiency today, I overlay the 10,000+ grid cells of my data set with every railway line built by the colonial powers in Sub-Saharan Africa. Figure VII prints in red 237 lines built between 1890 and the various independence dates. Data on railroad positioning comes from Jedwab and Moradi (2016), with the exception of South Africa, for which I manually digitise a map from Herranz-Loncán and Fourie (2017). No comparable data are available for Madagascar, Egypt, and the Maghreb countries, which also saw some colonial railway construction. As discussed below, findings are robust to excluding grid cells from these countries.

Figure VII: Colonial railway network



Maps displaying the network of railway lines (red) and placebo railroads (blue). Data from Jedwab and Moradi (2016) and Herranz-Loncán and Fourie (2017). Railroads built by the colonial powers between 1890 and 1960 are printed in red. Lines that were initially planned but never actually built are printed in blue.

For every grid cell, I compute the total number of colonial railway kilometres crossing the cell. This serves as a tangible measure for the stock of physical transport capital invested into each region. The majority of cells (91%) are not crossed by a colonial railway and hence have zero railroad kilometres. Those that are intersected by a line usually have between 20 and 60 railroad kilometres, while some important rail crossings or transport hubs have up to 100 kilometres of railroads. This measure captures the intensive margin of colonial infrastructure investment. I also construct a measure of extensive margin railroad exposure by computing the distance from a cell's centroid to its closest rail line. Every railroad line comes with a classification of being constructed primarily for military purposes or mining purposes (or neither, or both), allowing for more nuanced further analysis. Lastly, to account for potential endogeneity concerns, I also compute the same statistics for a set of railway lines the colonisers planned, but never built. As Jedwab and Moradi (2016) explain, these projects were not realised only for a series of arguably random historical events like unforeseeable cuts to financing, the outbreak of wars, or sudden retirements of administration officials. If any effects are to be attested to the construction of railroads during the colonial era, no impact should be

found for such *placebo* railroads. These tracks are printed in blue in Figure VII.²¹

Table I displays results from OLS estimation of equation (9) with Λ_i on the left hand side and total rail kilometres as explanatory variable v_i . Column (1) displays the plain cross-sectional relationship without controls and reveals a statistically significant and negative association between the two variables. Grid cells with high colonial railroad investment have significantly lower infrastructure discrimination today. Recall that low values of Λ_i correspond to regions losing welfare if the social planner were to optimally re-allocate infrastructure. On a merely descriptive level, the negative estimate of column (1) hence implies that regions with colonial railroads crossing through them would on average see infrastructure (and welfare) redistributed to those areas without colonial investments. Columns (2) – (4) gradually extend the set of observable controls. The point estimate and statistical significance of the persistence of railroads is robust to including country fixed effects, geographical controls, and simulation controls as described above. In the richest specification of column (4), every ten kilometres of colonial railway construction are associated with grid cells losing 0.2 percentage points of welfare at the hand of areas without any investment.

Columns (5) – (8) repeat the exercise for the set of placebo railroads. If these are assumed to have undergone the same planning process as actual railways, the regressions reveal the distorting power of site selection. None of the estimates are significantly different from zero, suggesting that the link in columns (1) – (4) is a causal one. The construction of railways by the colonial powers cause affected regions to be too well off today. The results are virtually identical when using only the subsample of 34 countries for which data on colonial railroad placement are available, or excluding all grid cells containing a national capital (not reported).

The effects described in Table I are small, yet remarkable. Across the African continent, areas that received large infrastructure investments a century ago are still *too* well off given their position in the national trade network. In contrast, areas that were not crossed by tracks are inefficiently short on infrastructure today. To see that this is a non-trivial finding, note that, firstly, most of the colonial railway lines have been in

²¹Data for placebo lines also come from Jedwab and Moradi (2016) and Herranz-Loncán and Fourie (2017).

Table I: Colonial railroads and local infrastructure discrimination index

	Dependent variable: Local Infrastructure Discrimination Index Λ_i							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
KM of Colonial Railroads	-0.0002 (0.0001)	-0.0001 (0.0001)	-0.0002 (0.0001)	-0.0002 (0.0001)				
KM of Colonial Placebo Railroads					0.00004 (0.0003)	-0.0002 (0.0003)	-0.0002 (0.0003)	-0.0003 (0.0003)
Country FE		Yes	Yes	Yes		Yes	Yes	Yes
Geographic controls			Yes	Yes			Yes	Yes
Simulation controls				Yes				Yes
Observations	10,158	10,158	10,158	10,158	10,158	10,158	10,158	10,158
R ²	0.001	0.099	0.124	0.126	0.00000	0.098	0.122	0.124

Results of estimation of equation (9) on the sample of 0.5×0.5 degree grid cells for the entire African continent (excluding six small countries, see text). Dependent variable is the Local Infrastructure Discrimination Index Λ_i for each grid cell. Columns (1)-(4) estimate the effect of colonial infrastructure investments as measured by the total number of colonial railroad kilometres crossing a cell. Starting with a simple univariate cross-section in (1), column (2) adds 49 country fixed effects. Column (3) adds geographic controls, consisting of altitude, temperature, average land suitability, malaria prevalence, yearly growing days, average precipitation, indicators for the 12 predominant agricultural biomes, indicators for whether a cell is within 25 KM of a natural harbour, navigable river, or lake, the fourth-order polynomial of latitude and longitude, and an indicator of whether the grid cell lies on the border of a country's network. Simulation controls are added in column (4) and are comprised of population, night lights, ruggedness, and a dummy for whether a cell is classified as urban. These indicators went into the original infrastructure reallocation simulation and are hence not orthogonal to Λ . Columns (5)-(8) repeat the estimations with railroads that were planned, but never built ("placebo railroads"). Results are robust to using only the subsample of 33 countries with colonial infrastructure investment as reported by Jedwab and Moradi (2016), plus South Africa (not reported). Results are also robust to excluding all grid cells containing a country's capital (not reported). Heteroskedasticity-robust standard errors are clustered on the 3×3 degree level and are shown in parentheses.

disrepair for decades and thus do not immediately dictate trade flows today. Secondly, recall that the optimal network reallocation and construction of Λ_i was based on roads and cars, not rails and trains. The implication is hence *not* that colonial railway systems themselves are inadequate to efficiently sustain inter-regional trade today. Rather the transport revolution a century ago coordinated the entire economy into a certain spatial equilibrium, which persists even though it has become inefficient. African nations would benefit from moving to a better equilibrium, but are locked in the current state. The placement of colonial railroads set in motion a process of spatial sorting with people, output, and infrastructure clustering in locations which are suboptimal today. The social planner identifies this, seeks to overcome these misallocations, and move infrastructure away from regions once considered important by the colonisers. Jedwab and Moradi show that colonial investments helped the economy to coordinate on one of many spatial equilibria – my findings suggest that this is not the optimal one.²²

²²Note that the spatial equilibrium induced by colonial railroads could have still been optimal *at the time*. My argument solely concerns the persistent effects of investments a century ago on network efficiency *today*.

Which railways were responsible for coordinating African economies into a suboptimal spatial equilibrium? To further understand the forces behind this dynamic, I split the sample of colonial railways along their initial construction purpose and separately calculate the total number of mining and military railroad kilometres crossing a cell. Table II repeats the estimations from above for both subsets of railways. As can be inferred, the effect is exclusively driven by railways built for military purposes. The social planner seeks to take welfare away from regions which were crossed by lines built for strategic military domination. Railroads constructed to support the mining trade are not associated with areas too well off today. This distinction holds both when including the variables separately (columns 1 – 4) or jointly (columns 5 – 6). This finding offers an intuitive understanding of how the current spatial distribution is inefficient. All colonial infrastructure investments spurred urbanisation close-by, regardless of their construction purpose. Since mining lines arguably cross areas that are still important trade routes today, the social planner sees no need to reorganise infrastructure away from them. Areas surrounding rails built with military motives in mind, however, have since lost their strategic importance. As the nature of military domination, state authority, and conflict have changed since the 19th century, there is no immediate value to clustering economic activity and infrastructure close to former military lines anymore. It is those sunk investments that skew the spatial equilibrium towards an inefficient equilibrium.

Until how far away does the confounding effect of former train lines reach? Table III displays results from regressing Λ_i on a series of indicators denoting whether a grid cell's centroid is within certain distance intervals from its closest colonial rail line. Columns (1) – (3) jointly estimate the effects of being within $[0 - 10]$, $(10 - 20]$, $(20 - 30]$, $(30 - 40]$, or $(40+)$ kilometres from the closest passing railroad (with $(40+)$ being the omitted category). These analyses uncover general equilibrium effects. Estimates from columns (1) – (3) suggest that areas close to railroads are better off at the expense of their immediate neighbours. Grid cells with centroids less than 10 kilometres away from a passing railroad line are between 1.3 and 1.7 percentage points too well off compared to the omitted category. The estimate is significantly different from zero and robust to gradually introducing additional controls. The same dynamic also holds for cells within 20 kilometres of a pass-

Table II: Heterogeneous effects of colonial railroads

	<i>Dependent variable: Local Infrastructure Discrimination Index Λ_i</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
KM of Rails for Military Purposes	-0.0002 (0.0001)	-0.0002 (0.0001)			-0.0002 (0.0001)	-0.0002 (0.0001)
KM of Rails for Mining Purposes			-0.0001 (0.0001)	-0.0001 (0.0001)	-0.0001 (0.0001)	-0.0001 (0.0001)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Simulation controls		Yes		Yes		Yes
Observations	10,158	10,158	10,158	10,158	10,158	10,158
R ²	0.123	0.125	0.122	0.124	0.123	0.125

Replication of estimations of Table I in estimating effects of colonial railroads on the Local Infrastructure Discrimination Index Λ . Colonial rails are classified as built for military or mining purposes (or neither or both) by Jedwab and Moradi (2016). Geographic controls consist of altitude, temperature, average land suitability, malaria prevalence, yearly growing days, average precipitation, indicators for the 12 predominant agricultural biomes, indicators for whether a cell is within 25 KM of a natural harbour, navigable river, or lake, the fourth-order polynomial of latitude and longitude, and an indicator of whether the grid cell lies on the border of a country's network. Simulation controls are comprised of population, night lights, ruggedness, and a dummy for whether a cell is classified as urban. Results are robust to using only the subsample of 34 countries with colonial infrastructure investment. Results are also robust to excluding all grid cells containing a country's capital (not reported). Heteroskedasticity-robust standard errors are clustered on the 3×3 degree level and are shown in parentheses.

ing line, and becomes undetectable further out. Moreover, one can even detect an adverse effect for cells between 30 and 40 kilometres away. Those cells would be granted additional welfare from redistributive efforts by the social planner. Though the estimate becomes increasingly imprecise and even insignificantly different from zero as further controls are introduced, the point estimate does not move by much. This is suggestive evidence that the confounding effect of colonial infrastructure policies is locally contained. Areas blessed with a close-by railway line are still too well off today, which comes at the expense of their neighbouring regions just a few kilometres away. To gain more confidence in the causal nature of this dynamic, columns (4) – (6) again repeat the same exercise with placebo railroads. Of the twelve estimates produced, none is significantly different from zero.²³

²³Results are virtually identical when restricting the sample to only grid cells from the 34 countries with data on colonial railway investment. One single coefficient on placebo railroads turns marginally significant ($p < 0.1$). While this is certainly noteworthy, I believe this association to be merely spurious. In fact, by the law of large numbers, one would even statistically expect at least one of 16 placebo estimations to be within this significance band.

Table III: General equilibrium effects of colonial railroads

	<i>Dependent variable: Local Infrastructure Discrimination Index Λ_i</i>					
	Full Sample					
	(1)	(2)	(3)	(4)	(5)	(6)
< 10 KM to Colonial Railroad	−0.013 (0.003)	−0.015 (0.004)	−0.017 (0.004)			
10 – 20 KM to Colonial Railroad	−0.013 (0.003)	−0.016 (0.004)	−0.017 (0.004)			
20 – 30 KM to Colonial Railroad	−0.002 (0.004)	−0.004 (0.004)	−0.005 (0.004)			
30 – 40 KM to Colonial Railroad	0.010 (0.005)	0.008 (0.005)	0.007 (0.005)			
< 10 KM to Colonial Placebo Railroad				−0.005 (0.004)	−0.006 (0.004)	−0.006 (0.004)
10 – 20 KM to Colonial Placebo Railroad				−0.003 (0.005)	−0.004 (0.005)	−0.005 (0.005)
20 – 30 KM to Colonial Placebo Railroad				−0.001 (0.004)	−0.001 (0.004)	−0.001 (0.004)
30 – 40 KM to Colonial Placebo Railroad				0.007 (0.004)	0.006 (0.004)	0.005 (0.004)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls		Yes	Yes		Yes	Yes
Simulation controls			Yes			Yes
Observations	10,158	10,158	10,158	10,158	10,158	10,158
R ²	0.101	0.126	0.129	0.099	0.123	0.125

Effects of various distance-intervals on the Local Infrastructure Discrimination Index Λ . Explanatory covariates are dummy-variables indicating whether a cell's centroid is within X kilometres to its closest colonial railroad. Distance larger than 40 kilometres is the omitted category. Geographic controls consist of altitude, temperature, average land suitability, malaria prevalence, yearly growing days, average precipitation, indicators for the 12 predominant agricultural biomes, indicators for whether a cell is within 25 KM of a natural harbour, navigable river, or lake, the fourth-order polynomial of latitude and longitude, and an indicator of whether the grid cell lies on the border of a country's network. Simulation controls are comprised of population, night lights, ruggedness, and a dummy for whether a cell is classified as urban. Columns (1)-(3) examine the effect of actually built colonial railroads. Columns (4)-(6) repeat these calculations with railroads that were planned, but never built ("placebo railroads"). Results are robust to excluding all grid cells containing a country's capital (not reported). Heteroskedasticity-robust standard errors are clustered on the 3×3 degree level and are shown in parentheses.

5.2 Ethnic relations

An ominous legacy of the colonial era is the design of Africa's national borders. Drawn without much regard for local circumstances and with European rather than African interests in mind, many territories amalgamated various previously unrelated tribes and ethnicities under the umbrella of one nationality. After independence, most borders persisted and today African nations are among the most ethnically diverse in the world (Alesina et al., 2016). A well established literature has investigated the impacts of ethnic diversity and identity on international comparative development (Easterly and Levine, 1997; Alesina and La Ferrara, 2005; Gennaioli and Rainer, 2007; Alesina et al., 2016). More recently, a series of new studies has begun to shine light on the causes and effects

of ethnicity-level heterogeneity within African countries. Ethnic homelands and tribes have been shown to be better developed when they are represented in national leadership (Franck and Rainer, 2012), have more centralised deep rooted institutional systems (Michalopoulos and Papaioannou, 2013), or simply are not split in two by artificial national border designs (Michalopoulos and Papaioannou, 2016).

The political economy of public good provision along ethnic lines presents an interesting case for the spatial imbalances in my dataset. Are ethnic homelands with less political clout systematically discriminated against in the provision of trade infrastructure? To analyse spatial patterns in the Local Infrastructure Discrimination Index Λ over ethnic homelands, I follow Michalopoulos and Papaioannou (2013, 2014, 2016) and intersect an ethnolinguistic map of pre-colonial homelands from Murdock (1959) with current national borders.²⁴ I spatially aggregate my grid cell measure of network inefficiency Λ_i onto the ethnicity-country level by assigning each grid cell an ethnicity based on its centroid location and weighing grid cells by their respective population.²⁵ 280 ethnic homelands are too small to overlay with any grid cell centroid, leaving me with 932 observations. Figure VIII presents the spatial variation of the Local Infrastructure Discrimination Index Λ_h on the ethnicity-country level.²⁶

I investigate two patterns of infrastructure discrimination along ethnic lines, namely the preferential treatment of groups sharing the ethnic homeland with the national leader (*ethnic favouritism*) and the negative treatment of groups excluded from the government

²⁴Ethnicity data are available for every country but Western Sahara. Ethnicities present in more than one country count as multiple observations. Of the 835 inhabited homelands identified by Murdock, 314 are split in two or more parts by the current national borders, creating 1,212 ethnicity-country observations.

²⁵Recall that on the grid cell level Λ_i was defined as

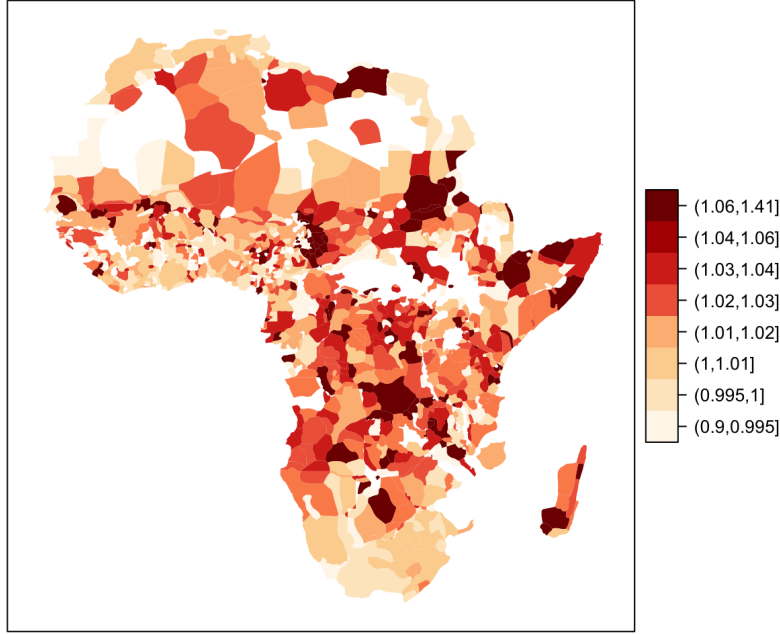
$$\Lambda_i = \frac{\text{Welfare under the optimal Infrastructure}_i}{\text{Welfare under the current Infrastructure}_i} = \frac{L_i u_i^{\text{optimal Infrastructure}}}{L_i u_i^{\text{current Infrastructure}}}$$

To aggregate this onto the ethnic homeland level, I sum over all grid cells i in homeland h

$$\Lambda_h = \frac{\sum_{i \in h} L_i u_i^{\text{optimal Infrastructure}}}{\sum_{i \in h} L_i u_i^{\text{current Infrastructure}}}$$

²⁶The homeland which would benefit most from national reshuffling of roads is the tiny Tunisian part of the Ghadames homeland. This area was identified as an oil basin in the 1990s, stands to gain more than 40% of welfare from reallocating national roads, and hence arguably presents an outlier (Echikh, 1998). Other ethnicities that are discriminated against mostly by the current network are the Kababish in Sudan (25%), the Bata in Cameroon (24%), or the Aushi in Congo-Kinshasa (23%). The most disproportionately advantaged ethnic homelands in Africa are those of the Mober in Nigeria (who stand to lose 8% of welfare if optimal networks were imposed) and the Sanga (7%) and Kreish (11%) of the Central African Republic.

Figure VIII: Λ_h over ethnic homelands



Spatial distribution of local infrastructure discrimination index Λ_h , aggregated over ethnic homelands. Unit of observation is pre-colonial homelands as initially defined in an ethnolinguistic map by Murdock (1959) intersected by current political borders (following Michalopoulos and Papaioannou, 2016). Λ_h is from grid cell level and weighted by population. Maps's colouring follows an equal-interval rule such that every colour in the spectrum has an equal amount of members. This is to visualise the measure's variation but leads to unequal bracket-sizes for each colour.

(ethnic discrimination). I begin with the latter. To measure ethnic discrimination, I rely on four measures of political relationships between ethnicities. Firstly, I make use of the Ethnic Power Relations (EPR) database by Vogt et al. (2015), which globally identifies “politically relevant ethnic groups and their access to state power” (Vogt et al., 2015, p. 1328) over the past seven decades. Not every ethnic homeland inhabits a group that is “politically relevant”, significantly truncating the sample by 46%.²⁷ EPR reports a yearly time series of political discrimination for every group in the sample. In particular, a group is coded as discriminated against by the central government if there is “active, intentional, and targeted discrimination by the state against group members in the domain of public politics” (Vogt et al., 2015, p. 1331). I follow Michalopoulos and Papaioannou (2016) and analyse effects of a dummy variable taking on the value one if a group has experienced discrimination in at least one year between 1960 and 2010 and use this measure to investigate whether infrastructure discrimination covaries with political discrimination. Secondly, I broaden the definition of ethnic discrimination

²⁷Merging EPR observations with ethnic homelands is non-trivial. Thankfully, I am able to rely on the conversion established by Michalopoulos and Papaioannou (2016).

and more generally look at groups which are excluded from the central government. As defined by EPR, this classification entails all groups that are discriminated against (from above), plus groups that are defined as either powerless or self-excluded (Vogt et al., 2015, p. 1331).²⁸ Thirdly, I analyse the effects of an EPR indicator denoting whether an ethnicity was part of a civil war with an explicitly ethnic dimension at some point between 1960 and 2010. The construction of this indicator is identical to the dummies described above and is obtained from Michalopoulos and Papaioannou (2016). As a more indirect measure for ethnic discrimination, I finally analyse whether ethnicities which were split in two or more parts by the arbitrary border drawings of the colonial powers get less than optimal infrastructure investment. Michalopoulos and Papaioannou (2016) show that ethnic homelands of which ten or more per cent are in more than one country are significantly more prone to violence, have less political clout, and report less overall well-being.

To analyse patterns of infrastructure discrimination on the ethnic homeland level, I estimate a slightly different version of equation (9)

$$\Lambda_{h,c} = \beta v_{h,c} + \mathbf{X}_{h,c}\gamma + \delta_c + \epsilon_{h,c} \quad (10)$$

where $\Lambda_{h,c}$ is the Local Infrastructure Discrimination Index for homeland h in country c , $\mathbf{X}'_{h,c}$ and δ_c again denote controls and country fixed effects respectively, $v_{h,c}$ are the explanatory covariates discussed above, and β is the coefficient of interest. The number of ethnicity observations (about 900) is significantly smaller than the number of grid cells. In order to avoid overfitting, I slightly truncate the set of controls $\mathbf{X}'_{h,c}$. I replace the latitude and longitude polynomials, the classification into urban or rural, as well as dummies indicating proximity to a natural harbour, river, lake, and national border with two continuous measures of distance to the nearest border and distance to the coast. As homelands are much more irregularly shaped than grid cells, I also include the natural logarithm of each homelands' area (as in Michalopoulos and Papaioannou, 2016). Apart from these adjustments, $\mathbf{X}'_{h,c}$ entails all geographical and simulation controls of the mod-

²⁸I again rely on the transformation by Michalopoulos and Papaioannou (2016) who code an indicator as one if the ethnic group has experienced exclusion from the government at any point between 1960 and 2010.

els on the grid cell level. As many ethnicities appear more than once, the error term will reasonably be autocorrelated beyond the country-level. $\epsilon_{h,c}$ is hence plausibly not independent even across countries. To account for this, I follow Michalopoulos and Papaioannou (2016) and double-cluster standard errors at both the country level as well as the ethnic family level using the mechanism proposed by Cameron et al. (2011).

Table IV reports results of estimating equation (10) for each of the four indicators of ethnic discrimination as independent variable. As before, the table prints estimates both with and without the set of simulation controls (consisting of population, night lights, and ruggedness). No estimate is significantly different from zero, implying that the null hypothesis of no linear relation between v_h and Λ_h cannot be rejected. There is no evidence suggesting that inefficiencies of the national trade network systematically covary with ethnicities that are historically discriminated against (columns 1–2), excluded from the government (3–4), involved in an ethnic war (5–6), or split by arbitrary colonial borders (7–8). After the social planner has reshuffled a country’s roads, historically victimised groups are not better off than they were before. These (null) findings do not change when (a) excluding homelands containing a nation’s capital, (b) including ethnic family fixed effects (c) controlling for pre-colonial differences in societal structure between ethnicities, namely complexity of hierarchies, the existence of compact settlement structures, and existence of a class system (Michalopoulos and Papaioannou, 2013), (d) using a more lenient definition of split homelands whereby only 5% of the territory has to be in more than one country (instead of 10%), (e) clustering standard errors along merely (either) one dimension, and (f) excluding all geographic controls. All estimates hardly move and remain statistically indistinguishable from zero (not reported).

To ensure that these null results are not merely an artefact of lack of statistical power in a small sample, Table IV additionally reports the *Minimum Detectable Effect Size* (MDE) under each estimate in brackets (Haushofer and Shapiro, 2016). This is the smallest estimate that would have still been detectable with 80% power at the 5% significance level and can be computed $MDE = 2.8 \times SE(\hat{\beta})$ (see Haushofer and Shapiro, 2016). MDEs in Table IV are all around 0.02. This implies that I was powered to detect at least a 2% increase (or reduction) in the Local Infrastructure Discrimination Index Λ_h – slightly

Table IV: Null effect of ethnic discrimination

	Dependent variable: Local Infrastructure Discrimination Index Λ_h							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ethnicity discriminated against 1960–2010 (SE) [MDE]	−0.001 (0.008) [0.021]	−0.001 (0.007) [0.021]						
Ethnicity excluded from the central government 1960–2010			−0.006 (0.005) [0.014]	−0.005 (0.005) [0.014]				
Ethnicity involved in an ethnic war 1960–2010					0.002 (0.008) [0.022]	0.002 (0.008) [0.022]		
Ethnicity split by colonial borders							−0.002 (0.004) [0.011]	−0.002 (0.004) [0.012]
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Simulation controls		Yes		Yes		Yes		Yes
Observations	496	496	496	496	496	496	932	932
R ²	0.156	0.166	0.158	0.168	0.156	0.167	0.164	0.167

Statistically insignificant effects of various indicators of ethnic discrimination on the Local Infrastructure Discrimination Index Λ_h . The sample comprises ethnic homelands, projected on current national borders. Independent variable in columns (1)–(2) is a dummy variable indicating if an ethnicity has experienced discrimination from the government at some point between 1960–2010. In columns (3)–(4), independent variable is a dummy indicating if an ethnicity has been excluded from the central government at some point between 1960–2010. (5)–(6) analyse impacts of ethnicities being involved in a major ethnic war at some point between 1960–2010. All data gathered by Vogt et al. (2015) and obtained in transformed form from Michalopoulos and Papaioannou (2016). Explanatory variable in (7) and (8) is an indicator of ethnicities being split by national borders, defined as having at least 10 per cent of their homeland in more than one country (from Michalopoulos and Papaioannou, 2016). Data for regressions in (1)–(6) only exist for politically relevant ethnic groups, truncating the sample by 46%. All observations exclude Western Sahara, for which no ethnic homeland data exist. Geographic controls consist of altitude, temperature, average land suitability, malaria prevalence, yearly growing days, average precipitation, indicators for the 12 predominant agricultural biomes, distances to the nearest coast and border, and the natural logarithm of the homeland area. Simulation controls are comprised of population, night lights, and ruggedness. Results are robust to excluding homelands containing a country’s capital (not reported). Heteroskedasticity-robust standard errors are double-clustered on the country level and the ethnic-family level and are reported in parentheses. All columns also report minimum detectable effect sizes (MDEs) in brackets. This is the smallest effect that would have still been detectable with 80% power at 5% significance (Haushofer and Shapiro, 2016).

larger effect sizes than being very close to a colonial railway (see Table III). Considering the importance of ethnic power relations on subregional development in Africa, I regard this as a reasonably fine resolution. The reported null effects of Table are hence not an artefact of underpowered tests, but rather bolster the contention that active ethnic victimisation and trade network inefficiency are not systematically linked.

I also analyse ethnic favouritism. Are ethnicities systematically better off when the country’s leader was born in their homeland? Existing studies have shown that the rise to power of a new national leader leads to temporarily more consumption and output in the leader’s birth region (Hodler and Raschky, 2014) and ethnic homeland (De Luca et al.,

2018). During the leader’s time in office, birth region and ethnic homelands also benefit from more foreign aid and infrastructure investment being channeled their way (Dreher et al., 2016; Burgess et al., 2015).

To investigate whether ethnic favouritism accounts for imbalances in trade infrastructure provision over space, I make use of a dataset of African national leaders provided by Dreher et al. (2016). The data entail information about the birth region and time in office of 117 heads of state holding power in 44 African countries dating back to 1969.²⁹ Using Open Street Map, I obtain coordinates for birthplaces and spatially merge them with my dataset on ethnic homelands. I then use this information to calculate for each ethnic group the total number of years someone born in the respective homeland has held high office. This allows me to analyse whether a homeland’s over provision with transport infrastructure covaries with personal ties to national power.

Table V investigates effects of ethnic and regional favouritism. Panel A, columns (1) through (4) estimate equation (10) on the full sample of ethnic homelands intersected with current national borders. In (1) and (2), the explanatory variable is the total number of years someone born in the ethnic homeland was holding national power. The covariate enters with a significant and negative coefficient, implying that for each year one of their members was in power, an ethnic homeland is about 0.04 percentage points too well off given their relative position in the country’s trade network. If the social planner were to intervene, she would strip the homeland with ethnic ties to power from some infrastructure and reallocate it towards areas with no such ties. The effect size is small – with the median leader staying in power for nine years, ethnic favouritism distorts network efficiency in the order of magnitude of less than half a percentage point. Nevertheless, it provides further evidence for many of the recent findings on ethnic favouritism (De Luca et al., 2018).

After having analysed intensive margin effects of the number of years in power, columns (3) and (4) extend the inquiry to the extensive margin. The explanatory variable is now a dummy equalling one if an ethnic homeland has ever had someone represent them as head of state (regardless of how long). The coefficient enters significantly and negatively,

²⁹No data on national leaders are reported for Algeria, Western Sahara, South Sudan, Somalia, and Djibouti. Even for countries with data, coverage is not comprehensive as the dataset excludes leaders born abroad or with unknown birthplaces. 93.6% of homelands never sent anyone to the highest office in the country. For those that did, tenures last from merely one year (the Zerma in Niger) to 42 years (the Duma in Gabon, homeland of long-term head of state Omar Bongo).

Table V: Ethnic favoritism

	<i>Dependent variable: Local Infrastructure Discrimination Index Λ_h</i>					
	Full Sample				Excluding Capitals	
	(1)	(2)	(3)	(4)	(5)	(6)
Years in Power	-0.0004 (0.0002)	-0.0004 (0.0002)			-0.0003 (0.0002)	
In Power Dummy			-0.007 (0.004)	-0.008 (0.004)		-0.007 (0.004)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Simulation controls		Yes		Yes	Yes	Yes
Observations	932	932	932	932	895	895
R ²	0.165	0.169	0.165	0.169	0.177	0.177

Persistent impacts of holding power on Local Infrastructure Discrimination Index Λ in leaders' homelands and birthplaces. Estimates equation (10) on the sample of ethnic homelands split by current national borders. Independent variable in columns (1)–(2) is the number of years since 1969 someone born in the homeland was the country's leader. In (3)–(4), the independent variable is a dummy indicating whether the homeland ever held power. Columns (5)–(6) replicate the regressions while excluding observations containing a country's capital. Data for leaders' birthplaces from Dreher et al. (2016). All observations exclude Western Sahara, for which no ethnic homeland data exist. Geographic controls consist of altitude, temperature, average land suitability, malaria prevalence, yearly growing days, average precipitation, indicators for the 12 predominant agricultural biomes. For ethnic homelands also distances to the nearest coast and border, and the natural logarithm of the homeland area. For grid cells also indicators for whether a cell is within 25 KM of a natural harbour, navigable river, or lake, the fourth-order polynomial of latitude and longitude, and an indicator of whether the grid cell lies on the border of a country's network. Simulation controls are comprised of population, night lights, ruggedness, and whether a cell is classified urban (for grid cells). Heteroskedasticity-robust standard errors are (double-)clustered on the country level (and the ethnic-family level for homelands) and are reported in parentheses.

implying that ethnic favouritism is not constrained to a few long-term national leaders.

The regressions of Table V are not perfectly identified. While I do control for an extensive set of auxiliary variables, unobservable differences between homelands might still simultaneously impact both their infrastructure provision, as well as their chances of sending some of their own to be national leader. This bias is particularly evident for national capitals: leaders disproportionally are born in the capital, and capitals also have significantly better infrastructure provision. To try to account for this confound, I re-estimate columns (2) and (4) while excluding all homelands which include a nation's capital. Results are printed in columns (5) – (6) and are not qualitatively different from the full sample estimates, yet significance becomes slightly weaker. The effect thus prevails even in homelands not geographically connected to the places of power.³⁰

³⁰Results are robust to replicating these regressions on the grid cell level (using equation 9) rather than the homeland level (not reported). Together, these findings lend support to both ethnic and regional

5.3 Foreign aid

Africa is the primary target of international aid. In 2017, no other world region was awarded more development disbursements from The World Bank – indeed, African countries received more aid than Europe, Central Asia, Latin America, and the Caribbean combined (The World Bank, 2017). Of almost 12 billion US dollars worth of lending commitments, the biggest share was awarded to projects aimed at improving transportation infrastructure.³¹ The World Bank is not alone – in the past decade, non-traditional players have entered and disrupted the international development aid system (Dreher and Fuchs, 2015). Most notably, China has emerged as a significant donor nation, funding development projects in at least 50 African countries since the turn of the millennium (Strange et al., 2017). Yet despite the vast amount of resources involved, foreign aid has not yet been unequivocally proven to be linked with positive economic outcomes in recipient countries (Burnside and Dollar, 2000; Easterly et al., 2004; Rajan and Subramanian, 2008; Clemens et al., 2012; Clemens and Kremer, 2016; Nunn and Qian, 2014).

To investigate whether international development aid is quantitatively associated to my measure of trade network inefficiency, I make use of two recently established datasets of geo-referenced aid flows to Africa. Firstly, AidData (2017) in cooperation with The World Bank, tracks over 5,600 lending lines from The World Bank to African nations and reports precise coordinates of over 60,000 projects financed through these funds, totalling more than 300 billion US dollars. The sample comprises all projects approved between 1996–2014. As Strandow et al. (2011) describe, attributing projects to locations relies on a double-blind coding procedure of various World Bank documents. Secondly, I explore patterns from a similar database on Chinese aid projects by Strange et al. (2017). The motivations behind China’s involvement in Africa are opaque and data on aid flows are much less transparent than from traditional donors, rendering precise, geo-referenced attribution much more cumbersome. Strange et al. resort to reports from numerous local and international media outlets to track official and unofficial financing lines to over 1,500 projects worth 73 billion US dollars in the period 2000–2011. As Strange et al. point

favouritism. My analysis is not able to conclusively distinguish between the two.

³¹The transport sector made up 18% of total IBRD and IDA lending to African nations, followed by water and sanitation (14%), energy and extractives (14%), and public administration (12%) (The World Bank, 2017).

out, media reports are often based on initial press releases and do not necessarily follow up on the eventual disbursement of every promised dollar. In that, the dataset is likely to capture Chinese funding *commitments* rather than actual *disbursements*. Insofar as donors usually commit to more than they eventually deliver, these figures present an upper bound of realised development assistance. Furthermore, while AidData (2017) claim their dataset on World Bank projects to be exhaustive, the dataset on Chinese aid will naturally miss some unofficial flows, as significant parts of Chinese involvement remain untracked.

For the purpose of this study, I exclude aid projects with no clear-cut geographical target like unconditional lending lines to the central government or assistance for political parties. I also exclude flows with unknown or only vague information on eventual project location.³² I also ignore projects which were still under construction or otherwise not fully completed by the end of 2017 (when I was scanning the Open Street Map database for road data). Together, these steps truncate the World Bank sample by 35% and the China sample by 52%. In Figure IX, I map the spatial distribution of aid projects from both remaining samples. Each circle represents a project site, with radius proportional to (log) US dollar disbursement value.³³

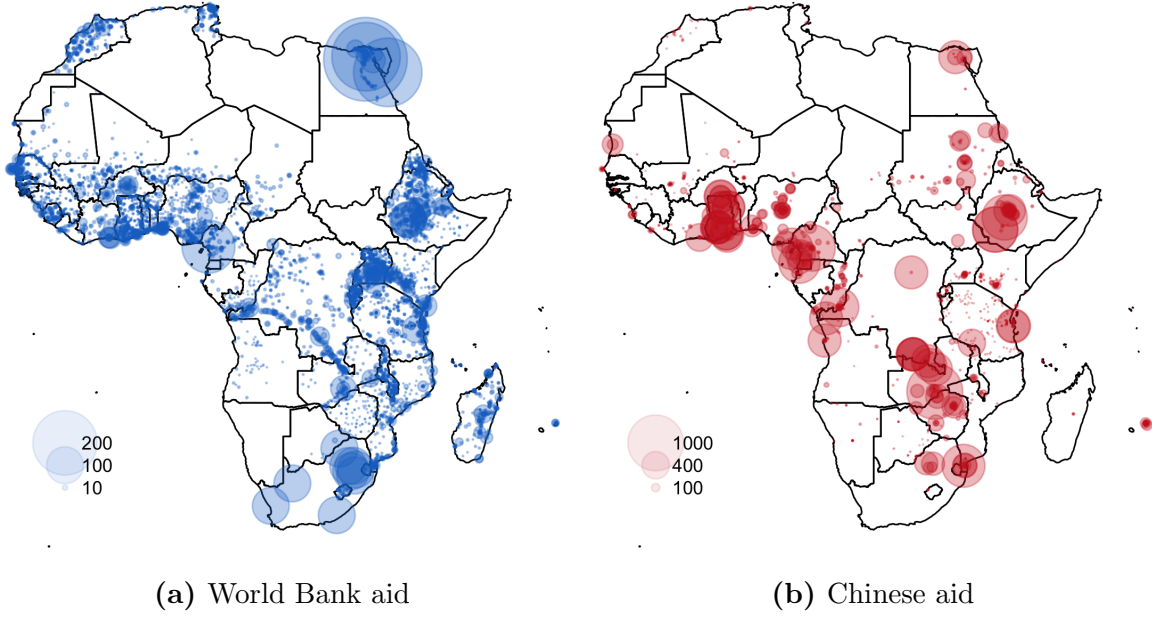
In order to analyse spatial relationships between foreign aid and my measure of local network inefficiency, I aggregate the total value of aid disbursements from the remaining 10,786 World Bank projects and 1,420 Chinese projects onto the grid cell level. Of the 10,158 grid cells of my sample, more than 21% have received some form of assistance from either source. On average, these cells receive aid volumes of more than 30 million US dollars. The area receiving the most total World Bank funding is the grid cell containing Uganda’s capital Kampala. The biggest beneficiary of Chinese development assistance is a grid cell in the south of Congo-Kinshasa, where Chinese funds of almost 5 billion US dollars helped construct a vast copper mining infrastructure.

Do donor institutions identify places most in need of additional infrastructure? I employ various indicators of aid provision in the standard grid cell level framework based on equation (9). I rely on two measures to quantify the prevalence of foreign aid: the

³²Specifically, I exclude all projects with a precision code of more than 3 – this corresponds to projects only identified at province-level or above. The remaining entries are geo-coded either exactly (61%), within a 25 kilometre radius (4%), or with municipality-level precision (35%) (Strandow et al., 2011).

³³All disbursements are adjusted to 2011 US dollars. For projects with multiple sites, I assume total disbursement value to be split evenly between sites.

Figure IX: Spatial distribution of development aid projects to African nations



Foreign aid projects funded by The World Bank (IXa) and China (IXb). Each dot represents one project site with radius proportional to the logarithm of total disbursements flowing to each site. World Bank data comprise all projects approved between 1996–2014. Chinese data include tracked projects between 2000–2011. Map only depicts projects coded with sufficient precision to not be excluded (see text). If a project has multiple sites, total disbursements are assumed evenly distributed between locations. Data from AidData (2017) and Strange et al. (2017). Legend denotes disbursement values in million 2011 US dollars. Note that the legends have different scales.

total value of aid disbursements to a grid cell in 2011 US dollars and the number of distinct project sites within a given cell. This distinction is useful to capture effects on both intensive and extensive margins. I also put additional emphasis on infrastructure by separately analysing variation in funds going only to infrastructure projects in the transportation sector.

Table VI reports results. Panel A investigates the spatial distribution of World Bank assistance. The estimates reveal seemingly opposing objectives between the Bank and the social planner. Negative estimates in columns (1) and (2) reveal that grid cells receiving more World Bank assistance score lower on the Local Infrastructure Discrimination Index Λ_i . These are areas the social planner identifies as overly privileged. Indeed, every additional million US dollar flowing into an area is associated with the grid cell being about 0.04 percentage points too well off, even after adjusting for geography, population, and economic activity, and only analysing within-country variation (column 2). Were the social planner to intervene and reallocate infrastructure, she would systematically take roads and welfare away from cells which received more World Bank funding. Columns

(3) and (4) zoom in on this dynamic and only analyse funds going to 2,949 transport infrastructure projects. Results are qualitatively similar, yet much stronger. The average transport infrastructure project size of around 3 million US dollars goes to grid cells which stand to lose 0.3 percentage points of welfare under the reallocation exercise. Similar effects hold on the extensive margin reported in columns (5)–(8). More project sites are associated with regions judged too well off by the social planner, regardless of whether the focus lies on assistance in general (5–6), or more specifically transportation infrastructure projects (7–8).

Panel B replicates the estimations with data from the Chinese aid sample. Results are very similar, yet two differences in magnitude stand out. Firstly, the association of Chinese money with network inefficiency is smaller, but also significantly different from zero (columns 1–2). Chinese assistance systematically flows into privileged cells, with point estimates of the association ranging between a quarter and a third of the World Bank results in Panel A. Again, the relationship is more pronounced for aid flows towards projects in the transport sector (columns 3–4). Secondly, on the extensive margin, more Chinese projects are similarly associated with higher trade network imbalances. For each new development site financed by China in a certain cell, the social planner intervenes and allocates 0.4 percentage points of welfare away from the cell (columns 5–6). For transport-sector projects only, the estimate grows to 1.4 percentage points – almost a quarter standard deviation in Λ_i per project (7–8). These results are substantially larger than the respective World Bank estimates, which does not come as a surprise as Chinese projects in the database are on average much more voluminous (and capture commitments rather than potentially lower disbursement sums).³⁴

These relationships should by no means be interpreted as clear causal effects. Since the placement of aid projects is not random, numerous other channels could account for the patterns depicted in Table VI. And though the evidence in Table VI may seem damning,

³⁴Results from Panel A are fully robust to following Dreher and Lohmann (2015) and taking the natural logarithm of aid disbursements (plus a small number to include places without aid) as dependent variable, or removing the upper 1st-percentile of aid receiving cells. Results also stay unchanged when excluding grid cells containing a nation’s capital or excluding grid cells containing a national leader’s birthplace. Findings from Panel B are also unchanged by all of these additional specifications, with the exception of the relationship between transport sector aid volumes and network inefficiency when excluding capitals. This estimate remains negative, yet narrowly loses its distinction from zero at the 10% significance level ($p = 0.106$).

Table VI: Foreign aid projects

	<i>Dependent variable: Local Infrastructure Discrimination Index Λ_i</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: Worldbank Projects</i>								
Total disbursements in million 2011 US dollars	-0.0003 (0.0001)	-0.0004 (0.0001)						
Transport-sector disbursements in million 2011 US dollars			-0.001 (0.0002)	-0.001 (0.0002)				
Number of projects					-0.002 (0.0004)	-0.003 (0.0004)		
Number of transport projects							-0.003 (0.001)	-0.004 (0.001)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Simulation controls		Yes		Yes		Yes		Yes
Observations	10,158	10,158	10,158	10,158	10,158	10,158	10,158	10,158
R ²	0.125	0.128	0.125	0.127	0.127	0.131	0.126	0.129
<i>Panel B: Chinese Development Projects</i>								
Total commitments in million 2011 US dollars	-0.0001 (0.00004)	-0.0001 (0.00004)						
Transport-sector commitments in million 2011 US dollars			-0.0003 (0.0001)	-0.0003 (0.0001)				
Number of projects					-0.003 (0.001)	-0.004 (0.001)		
Number of transport projects							-0.013 (0.004)	-0.014 (0.005)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Simulation controls		Yes		Yes		Yes		Yes
Observations	10,158	10,158	10,158	10,158	10,158	10,158	10,158	10,158
R ²	0.123	0.125	0.123	0.125	0.124	0.126	0.123	0.125

Grid cell level estimations of equation (9) with local infrastructure discrimination Λ_i as dependent variable and different measures of foreign aid flows into grid cells as explanatory covariates. Panel A investigates World Bank assistance. Columns (1)–(2) analyse total disbursement value from World Bank projects approved from 1996–2014 in 2011 US dollars, which were completed by 2017. (3)–(4) only use a subset of projects in the transport sector. (5)–(8) use the same data but focus on the number of distinct project sites within each grid cell. Panel B repeats the same estimations, but with data on Chinese aid projects between 2000–2011. Geographic controls consist of altitude, temperature, average land suitability, malaria prevalence, yearly growing days, average precipitation, indicators for the 12 predominant agricultural biomes, indicators for whether a cell is within 25 KM of a natural harbour, navigable river, or lake, the fourth-order polynomial of latitude and longitude, and an indicator of whether the grid cell lies on the border of a country’s network. Simulation controls are comprised of population, night lights, ruggedness, and a dummy for whether a cell is classified as urban. Data from AidData (2017) and Strange et al. (2017). Chinese aid data are more likely to reflect commitments rather than actual disbursements. Heteroskedasticity-robust standard errors are clustered on the 3×3 degree level and are shown in parentheses.

it might nevertheless be the result of very legitimate transport investment decisions. One vindication for the observed pattern could come at hand of the distinctively discrete nature of transformative infrastructure projects. To modernise and economise an overstressed road system, The World Bank (or China, for that matter) is often faced with the choice

between devising a large and necessary project like a new highway or modern bridge, which by means of their discrete nature always slightly overcorrect previous imbalances. This in turn leads the social planner to again readjust their impact. The donor’s investment strategies might also be exonerated by way of assuming increasing returns to scale. If The World Bank believes in an environment with multiple equilibria, where small initial investments set in motion a dynamic of spillover externalities, labour migration, and follow-up investments, it is often the right decision to fund projects in places that will not immediately harness their full capabilities (Krugman, 1991a,b, 1996; Fujita et al., 1999; Duranton and Venables, 2017). These investments will necessarily appear inefficient in promoting optimal trade *today*, yet spur transformative development *tomorrow* (see Michaels et al., 2018). Embedding the reallocation exercise in a New Economic Geography framework of increasing returns and labour mobility might be a valuable extension to better evaluate specific place-based policies.

6 Conclusion

In this study, I have identified spatial inefficiencies in Africa’s trade network. I first constructed a comprehensive economic topography of the entire continent, bringing together data from a variety of sources like satellites, census bureaus, and open source online routing services. I then presented a simple two-sector endowment network trade model and simulated the flow of goods through the internal geography formed by 10,000 African regions and almost 75,000 network connections. Harnessing the recent theoretical contribution by Fajgelbaum and Schaal (2017), I proceeded to endogenize the transport network in order to derive the unique optimally reorganised road network for every country in Africa.

In the second part of this study, I compared each country’s current network to its hypothetically optimal one. I ranked countries by overall network efficiency and presented a fine-resolution spatial dataset quantifying which sub-national areas are disadvantaged by the status quo. I empirically investigated patterns of trade network imbalances over space and linked inefficiencies to persistent lock-in effects caused by colonial infrastructure investments and differential treatment on the basis of ethnic favouritism. I found no effect

of ethnic discrimination. I also documented how development assistance by The World Bank and China has not targeted the regions identified as most in need of additional investment.

In contributing a comprehensive spatial measure on the differential provision of a primary public good covering an entire continent, my study provides the quantitative foundation for many more research questions pertaining to inequality over space. Future research designs could employ my dataset to analyse regional roots of conflict, political activism, social mobility, or subjective overall wellbeing. Another interesting avenue for inquiry could be to investigate whether infrastructure inefficiency spatially covaries with the provision of other public goods like education, health, or security.

Identifying spatial inefficiencies and understanding their historical, cultural, and political roots can be the first step in outlining effective place-based policies. Equipped with an unparalleled availability of spatial data and computing power, policymakers in Africa and around the world should feel empowered to combat local imbalances and design powerful interventions to better connect millions.

References

- AidData (2017).** World Bank Geocoded Research Release, Version 1.4.2. <http://aiddata.org/data/world-bank-geocoded-research-release-level-1-v1-4-2>. Williamsburg, VA and Washington, DC
- Alder, Simon (2017).** Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development. Working Paper, University of North Carolina at Chapel Hill
- Alesina, Alberto, William Easterly, and Janina Matuszeski (2011).** Artificial States. *Journal of the European Economic Association* 9(2), pp. 246–277
- Alesina, Alberto and Eliana La Ferrara (2005).** Ethnic Diversity and Economic Performance. *Journal of Economic Literature* 43(3), pp. 762–800
- Alesina, Alberto, Stelios Michalopoulos, and Elias Papaioannou (2016).** Ethnic Inequality. *Journal of Political Economy* 124(2), pp. 428–488
- Allen, Treb and Costas Arkolakis (2014).** Trade and the Topography of the Spatial Economy. *The Quarterly Journal of Economics* 129(3), pp. 1085–1140
- (2016). The Welfare Effects of Transportation Infrastructure Improvements. Working Paper, Dartmouth and Yale
- Atkin, David and Dave Donaldson (2015).** Who’s Getting Globalized? The Size and Implications of Intra-National Trade Costs. Working Paper, National Bureau of Economic Research
- Baum-Snow, Nathaniel, Loren Brandt, J. Vernon Henderson, Matthew A. Turner, and Qinghua Zhang (2017).** Roads, Railroads, and Decentralization of Chinese Cities. *The Review of Economics and Statistics* 99(3), pp. 435–448

- Bernot, Marc, Vicent Caselles, and Jean-Michel Morel (2009).** *Optimal Transportation Networks*, volume 1955 of *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, Berlin, Heidelberg
- Bester, C. Alan, Timothy G. Conley, and Christian B. Hansen (2011).** Inference with Dependent Data Using Cluster Covariance Estimators. *Journal of Econometrics* 165(2), pp. 137–151
- Boulhol, H. and A. de Serres (2010).** Have Developed Countries Escaped the Curse of Distance? *Journal of Economic Geography* 10(1), pp. 113–139
- Burgess, Robin and Dave Donaldson (2012).** Railroads and the Demise of Famine in Colonial India. Working Paper, LSE and NBER
- Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel (2015).** The Value of Democracy: Evidence from Road Building in Kenya. *The American Economic Review* 105(6), pp. 1817–1851
- Burnside, Craig and David Dollar (2000).** Aid, Policies, and Growth. *The American Economic Review* 90(4), pp. 847–868
- Cameron, A. Colin, Jonah Gelbach, and Douglas Miller (2011).** Robust Inference with Multi-Way Clustering. *Journal of Business & Economic Statistics* 29(2), pp. 238–249
- Clemens, Michael A. and Michael Kremer (2016).** The New Role for the World Bank. *Journal of Economic Perspectives* 30(1), pp. 53–76
- Clemens, Michael A., Steven Radelet, Rikhil R. Bhavnani, and Samuel Bazzi (2012).** Counting Chickens When They Hatch: Timing and the Effects of Aid on Growth. *The Economic Journal* 122(561), pp. 590–617
- Collier, Paul, Martina Kirchberger, and Mans Söderbom (2015).** The Cost of Road Infrastructure in Low-and Middle-Income Countries. *The World Bank Economic Review* 30(3), pp. 522–548
- Coşar, A. Kerem and Pablo D. Fajgelbaum (2016).** Internal Geography, International Trade, and Regional Specialization. *American Economic Journal: Microeconomics* 8(1), pp. 24–56
- De Luca, Giacomo, Roland Hodler, Paul A. Raschky, and Michele Valsecchi (2018).** Ethnic Favoritism: An Axiom of Politics? *Journal of Development Economics* 132, pp. 115–129
- Donaldson, Dave (2018).** Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *The American Economic Review* 108(4), pp. 899–934
- Donaldson, Dave and Richard Hornbeck (2016).** Railroads and American Economic Growth: A “Market Access” Approach. *The Quarterly Journal of Economics* 131(2), pp. 799–858
- Donaldson, Dave and Adam Storeygard (2016).** The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives* 30(4), pp. 171–198
- Doxsey-Whitfield, Erin, Kytt MacManus, Susana B. Adamo, Linda Pistolesi, John Squires, Olena Borkovska, and Sandra R. Baptista (2015).** Taking Advantage of the Improved Availability of Census Data: A First Look at the Gridded Population of the World, Version 4. *Papers in Applied Geography* 1(3), pp. 226–234
- Dreher, Axel and Andreas Fuchs (2015).** Rogue Aid? An Empirical Analysis of China’s Aid Allocation. *Canadian Journal of Economics/Revue canadienne d’économique* 48(3), pp. 988–1023
- Dreher, Axel, Andreas Fuchs, Roland Hodler, Bradley Parks, Paul Raschky, and Michael Tierney (2016).** Aid on Demand: African Leaders and the Geography of China’s Foreign Assistance. Working Paper, AidData
- Dreher, Axel, Andreas Fuchs, Bradley Parks, Austin Strange, and Michael Tierney (2017).** Aid, China, and Growth: Evidence from a New Global Development Finance Dataset. AidData Working Paper 46, AidData, Williamsburg, VA
- Dreher, Axel and Steffen Lohmann (2015).** Aid and Growth at the Regional Level. *Oxford Review of Economic Policy* 31(3-4), pp. 420–446

- Duranton, Gilles and Anthony J Venables (2017).** Place-Based Policies for Development. Working Paper, The World Bank
- Easterly, William and Ross Levine (1997).** Africa's Growth Tragedy: Policies and Ethnic Divisions. *The Quarterly Journal of Economics* 112(4), pp. 1203–1250
- Easterly, William, Ross Levine, and David Roodman (2004).** Aid, Policies, and Growth: Comment. *The American Economic Review* 94(3), pp. 774–780
- Echikh, K. (1998).** Geology and Hydrocarbon Occurrences in the Ghadames Basin, Algeria, Tunisia, Libya. *Geological Society, London, Special Publications* 132(1), pp. 109–129
- Faber, Benjamin (2014).** Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System. *The Review of Economic Studies* 81(3), pp. 1046–1070
- Fajgelbaum, Pablo D. and Cecile Gaubert (2018).** Optimal Spatial Policies, Geography and Sorting. Working Paper, UC Berkeley
- Fajgelbaum, Pablo D. and Edouard Schaal (2017).** Optimal Transport Networks in Spatial Equilibrium. Working Paper, National Bureau of Economic Research
- Foster, Vivien and Cecilia Briceño-Garmendia (2010).** *Africa's Infrastructure: A Time for Transformation*. Africa Development Forum. World Bank Publications, Washington DC
- Franck, Raphaël and Ilia Rainer (2012).** Does the Leader's Ethnicity Matter? Ethnic Favoritism, Education, and Health in Sub-Saharan Africa. *The American Political Science Review* 106(02), pp. 294–325
- Fujita, Masahisa, Paul R. Krugman, and Anthony Venables (1999).** *The Spatial Economy: Cities, Regions, and International Trade*. MIT Press, Cambridge, Mass. and London
- Gennaioli, Nicola and Ilia Rainer (2007).** The Modern Impact of Precolonial Centralization in Africa. *Journal of Economic Growth* 12(3), pp. 185–234
- Haushofer, Johannes and Jeremy Shapiro (2016).** The Short-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya. *The Quarterly Journal of Economics* 131(4), pp. 1973–2042
- Henderson, J. Vernon, Tim Squires, Adam Storeygard, and David N Weil (2018).** The Global Spatial Distribution of Economic Activity: Nature, History, and the Role of Trade. *The Quarterly Journal of Economics* 133(1), pp. 357–406
- Henderson, J. Vernon, Adam Storeygard, and David N Weil (2012).** Measuring Economic Growth from Outer Space. *The American Economic Review* 102(2), pp. 994–1028
- Herranz-Loncán, Alfonso and Johan Fourie (2017).** “For the Public Benefit”? Railways in the British Cape Colony. *European Review of Economic History* pp. 1–28
- Hodler, Roland and Paul A. Raschky (2014).** Regional Favoritism. *The Quarterly Journal of Economics* 129(2), pp. 995–1033
- Jedwab, Remi, Edward Kerby, and Alexander Moradi (2017).** History, Path Dependence and Development: Evidence from Colonial Railways, Settlers and Cities in Kenya. *The Economic Journal* 127(603), pp. 1467–1494
- Jedwab, Remi and Alexander Moradi (2016).** The Permanent Effects of Transportation Revolutions in Poor Countries: Evidence from Africa. *The Review of Economics and Statistics* 98(2), pp. 268–284
- Krugman, Paul (1991a).** History Versus Expectations. *The Quarterly Journal of Economics* 106(2), pp. 651–667
- (1991b). Increasing Returns and Economic Geography. *Journal of Political Economy* 99(3), pp. 483–499
- (1996). Urban Concentration: The Role of Increasing Returns and Transport Costs. *International Regional Science Review* 19(1), pp. 1–26
- Lall, Somik Vinay, J. Vernon Henderson, and Anthony J. Venables (2017).** *Africa's Cities: Opening Doors to the World*. World Bank Publications

- Limao, Nuno and Anthony J. Venables (2001).** Infrastructure, Geographical Disadvantage, Transport Costs, and Trade. *The World Bank Economic Review* 15(3), pp. 451–479
- Michaels, Guy, Dzhamilya Nigmatulina, Ferdinand Rauch, Tanner Regan, Neeraj Baruah, and Amanda Dahlstrand-Rudin (2018).** Planning Ahead for Better Neighborhoods: Long Run Evidence from Tanzania. Working Paper, LSE
- Michaels, Guy and Ferdinand Rauch (2017).** Resetting the Urban Network: 117–2012. *The Economic Journal* pp. 1–35
- Michalopoulos, Stelios and Elias Papaioannou (2013).** Pre-Colonial Ethnic Institutions and Contemporary African Development. *Econometrica* 81(1), pp. 113–152
- (2014). National Institutions and Subnational Development in Africa. *The Quarterly Journal of Economics* 129(1), pp. 151–213
- (2016). The Long-Run Effects of the Scramble for Africa. *The American Economic Review* 106(7), pp. 1802–1848
- Murdock, George Peter (1959).** *Africa: Its Peoples and Their Culture History*. McGraw-Hill, New York; London
- Nunn, Nathan and Nancy Qian (2014).** US Food Aid and Civil Conflict. *The American Economic Review* 104(6), pp. 1630–1666
- Rajan, Raghuram G and Arvind Subramanian (2008).** Aid and Growth: What Does the Cross-Country Evidence Really Show? *The Review of Economics and Statistics* 90(4), pp. 643–665
- Socioeconomic Data and Applications Center (2016).** Gridded Population of the World (GPW), v4 — SEDAC. <http://sedac.ciesin.columbia.edu/data/collection/gpw-v4>
- Strandow, Daniel, Michael Findley, Daniel Nielson, and Josh Powell (2011).** The UCDP and AidData Codebook on Georeferencing Aid Version 1.1. Technical report, Department of Peace and Conflict Research, Uppsala University, Uppsala
- Strange, Austin M., Axel Dreher, Andreas Fuchs, Bradley Parks, and Michael J. Tierney (2017).** Tracking Underreported Financial Flows: China’s Development Finance and the Aid–Conflict Nexus Revisited. *Journal of Conflict Resolution* 61(5), pp. 935–963
- Swisher, I. V. (2017).** Reassessing Railroads and Growth: Accounting for Transport Network Endogeneity. Working Paper, Faculty of Economics, University of Cambridge
- Teravaninthorn, Supee and Gaël Raballand (2009).** *Transport Prices and Costs in Africa: A Review of the Main International Corridors*. Directions in Development; Infrastructure. World Bank Publications, Washington DC
- The Economist (2015).** All Aboard. *The Economist* February 26, 2015
- The World Bank (2017).** The World Bank Annual Report 2017. Technical report, The World Bank, Washington DC
- Vogt, Manuel, Nils-Christian Bormann, Seraina Rüegger, Lars-Erik Cederman, Philipp Hunziker, and Luc Girardin (2015).** Integrating Data on Ethnicity, Geography, and Conflict: The Ethnic Power Relations Data Set Family. *Journal of Conflict Resolution* 59(7), pp. 1327–1342

Appendix

A Numerically solving the planner's problem

The full planner's problem on page 9 consists of a very large number of choice variables and hence requires vast computation efforts when solved directly. Fortunately, Fajgelbaum and Schaal (2017) provide guidance on how to transform this *primal* problem into its much simpler *dual* representation. The following section illustrates how to use their derivation to numerically solve my version of the model.

To show how a unique global optimum exists, first note that every constraint of the social planner's problem is convex but potentially for the *Balanced Flows Constraint*. However, the introduction of congestion causes even the *Balanced Flows Constraint* to be convex if $\beta > \gamma$. To see this, note that every part of the lengthy constraint is linear, but for the interaction term $Q_{i,k}^n \tau_{i,k}^n(Q_{i,k}^n, I_{i,k})$ representing total trade costs. Since $\tau_{i,k}^n$ was parameterised as in (1), this expands to

$$Q_{i,k}^n \tau_{i,k}^n(Q_{i,k}^n, I_{i,k}) = \delta_{i,k}^\tau \frac{(Q_{i,k}^n)^{1+\beta}}{I_{i,k}^\gamma} \quad (\text{A.1})$$

which is convex if $\beta > \gamma$. Under this condition, the social planner's problem is to maximise a concave objective over a convex set of constraints, guaranteeing that any local optimum is indeed a global maximum.^{A.1} $\beta > \gamma$ describes a notion of congestion dominance: increased infrastructure expenditure might alleviate the powers of congestion, but it can never overpower it. It precludes corner solutions in which all available concrete is spent on one link, all but washing away trade costs and leading to overwhelming transport flows on this one edge. If $\beta > \gamma$, geography always wins.

Consider first the full Lagrangian of the primal planner's problem

$$\begin{aligned} \mathcal{L} = & \sum_i L_i u(c_i) - \sum_i \lambda_i^C \left[L_i c_i - \left(\sum_{n=1}^N (C_i^n)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right] \\ & - \sum_i \sum_n \lambda_{i,n}^P \left[C_i^n + \sum_{k \in N(i)} Q_{i,k}^n (1 + \tau_{i,k}^n(Q_{i,k}^n, I_{i,k})) - Y_i^n - \sum_{j \in N(i)} Q_{j,i}^n \right] \\ & - \lambda^I \left[\sum_i \sum_{k \in N(i)} \delta_{i,k}^i I_{i,k} - K \right] - \sum_i \sum_{k \in N(i)} \zeta_{i,k}^S \left[I_{i,k} - I_{k,i} \right] \\ & + \sum_i \sum_{k \in N(i)} \sum_n \zeta_{i,k,n}^Q Q_{i,k}^n + \sum_i \sum_n \zeta_{i,n}^C C_i^n + \sum_i \sum_n \zeta_i^c c_i - \sum_i \sum_{k \in N(i)} \zeta_{i,k}^I \left[4 - I_{i,k} \right] \end{aligned} \quad (\text{A.2})$$

This is a function of the choice variables $(C_i^n, Q_{i,k}^n, c_i, I_{i,k})$ in all dimensions $\langle i, k, n \rangle$ and the Lagrange multipliers $(\lambda^C, \lambda^P, \lambda^I, \zeta^Q, \zeta^C, \zeta^c, \zeta^I)$ also in $\langle i, k, n \rangle$. Standard optimisa-

^{A.1}This is Fajgelbaum and Schaal Proposition 1.

tion yields first-order conditions which can be collapsed to the following set of equations

$$\begin{aligned}
c_i &= \left(\frac{1}{\alpha} \left(\sum_{n'} (\lambda_{i,n'}^P)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \right)^{\frac{1}{\alpha-1}} \\
C_i^m &= \left[\frac{\lambda_{i,n}^P}{\left(\sum_{n'} (\lambda_{i,n'}^P)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}} \right]^{-\sigma} L_i c_i \\
Q_{i,k}^n &= \left[\frac{1}{1+\beta} \frac{I_{i,k}^\gamma}{\delta_{i,k}^\tau} \max \left\{ \frac{\lambda_{k,n}^P}{\lambda_{i,n}^P} - 1, 0 \right\} \right]^{\frac{1}{\beta}} \\
I_{i,k} &= \max \left\{ \left[\frac{\kappa}{\lambda^I (\delta_{i,k}^I + \delta_{k,i}^I)} \left(\sum_n \max \left\{ (\delta_{i,k}^\tau)^{-\frac{1}{\beta}} \lambda_{i,n}^P \left(\frac{\lambda_{k,n}^P}{\lambda_{i,n}^P} - 1 \right)^{\frac{1+\beta}{\beta}}, 0 \right\} \right. \right. \right. \\
&\quad \left. \left. \left. + \sum_n \max \left\{ (\delta_{k,i}^\tau)^{-\frac{1}{\beta}} \lambda_{k,n}^P \left(\frac{\lambda_{i,n}^P}{\lambda_{k,n}^P} - 1 \right)^{\frac{1+\beta}{\beta}}, 0 \right\} \right) \right]^{\frac{\beta}{\beta-\gamma}}, 4 \right\}
\end{aligned} \tag{A.3}$$

These directly follow the more general framework outlined in the technical appendix of Fajgelbaum and Schaal applied to my version of the model. In the final equation denoting optimal infrastructure supply, $\kappa = \gamma(1+\beta)^{-\frac{1+\beta}{\beta}}$, and the multiplier λ^I is such that adherence to the *Network Building Constraint* is ensured. Note that there is a typo in the original authors' paper which prints one of the exponents as $(\delta_{i,k}^\tau)^{\frac{1}{\beta}}$ when it should be $(\delta_{i,k}^\tau)^{-\frac{1}{\beta}}$. Through these algebraic manipulations, I have expressed all choice variables as functions of merely the Lagrange parameters $\boldsymbol{\lambda}^P$ over dimensions $\langle i, k, n \rangle$. I can hence recast the entire Lagrangian in much simpler form as

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\lambda}, x(\boldsymbol{\lambda})) &= \sum_i L_i u(c_i(\boldsymbol{\lambda})) \\
&\quad - \sum_i \sum_n \lambda_{i,n}^P \left[C_i^m(\boldsymbol{\lambda}) + \sum_{k \in N(i)} Q_{i,k}^n(\boldsymbol{\lambda}) (1 + \tau_{i,k}^n (Q_{i,k}^n(\boldsymbol{\lambda})^n, I_{i,k}(\boldsymbol{\lambda}))) - Y_i^n - \sum_{j \in N(i)} Q_{j,i}^n(\boldsymbol{\lambda}) \right]
\end{aligned} \tag{A.4}$$

where $x(\boldsymbol{\lambda})$ denote the choice variables as functions of the Lagrange parameters as derived above. Fajgelbaum and Schaal note that thanks to complementary slackness, all other constraints can be readily dropped from consideration and only the *Balanced Flows Constraint* remains part of the problem.

As Fajgelbaum and Schaal further explain, the dual of this problem can now be conceived as the minimisation of


$$\min_{\boldsymbol{\lambda} \geq 0} \mathcal{L}(\boldsymbol{\lambda}, x(\boldsymbol{\lambda}))$$

which is an optimisation problem over merely $\|\boldsymbol{\lambda}^P\| = I \times N$ variables. Fajgelbaum and Schaal interpret $\boldsymbol{\lambda}^P$ as a field of prices varying over goods and locations. I am left only to minimise equation (A.4) to obtain the price-field $\boldsymbol{\lambda}^P$. I implement constrained optimisations within the `fmincon` environment in MATLAB and achieve fairly fast con-

vergence. Solving for smaller networks (like Rwanda or Djibouti) is a matter of seconds, yet the largest countries (Algeria, Angola, DRC, and Sudan) each take about a day of computation time (on a five-year old device, nonetheless). Plugging the derived λ^P parameters into the various FOCs in (A.3) yields the optimal transport network $I_{i,k}$, trade flows between locations $Q_{i,k}^n$, and consumption patterns C_i^n and c_i .

B Additional figures and tables

Figure A.1: Histogram of Λ_i



Histogram displaying the frequency distribution of Λ_i over all 10,158 grid cells. Plot also shows the respective PDF of a normal distribution with mean and standard deviation matching Λ_i .

Table A.1: Correlations of Λ_i with the various control sets

	<i>Dependent variable: Λ_i</i>		
	Geo	+ FE	+ Simulation
Altitude	0.00001 (0.00001)	-0.00001 (0.00001)	-0.00001 (0.00001)
Average Yearly Temperature	0.001 (0.001)	-0.003 (0.001)	-0.003 (0.001)
Average Land Suitability	0.008 (0.007)	0.007 (0.007)	0.007 (0.007)
Malaria Transmission Index	0.001 (0.0002)	0.0005 (0.0003)	0.0005 (0.0003)
Biome 1	0.018 (0.020)	0.022 (0.011)	0.022 (0.011)
Biomes 2 & 3	-0.027 (0.022)	-0.020 (0.014)	-0.020 (0.014)
Biome 5	-0.049 (0.024)	-0.036 (0.016)	-0.036 (0.016)
Biomes 7 & 9	0.017 (0.020)	0.025 (0.011)	0.025 (0.011)
Biome 10	0.015 (0.020)	0.022 (0.011)	0.022 (0.011)
Biome 12	-0.028 (0.022)	-0.026 (0.013)	-0.025 (0.014)
Biome 13	0.005 (0.020)	0.009 (0.010)	0.008 (0.011)
Biome 14	-0.005 (0.035)	-0.002 (0.033)	-0.002 (0.032)
< 25KM to Natural Harbour	-0.019 (0.012)	-0.017 (0.012)	-0.014 (0.013)
< 25KM to Navigable River	-0.019 (0.007)	-0.005 (0.005)	-0.001 (0.005)
< 25KM to Lake	-0.001 (0.009)	0.007 (0.009)	0.006 (0.008)
Yearly Growing Days	-0.00001 (0.00003)	-0.0001 (0.00004)	-0.0001 (0.00004)
Average Precipitation	0.00002 (0.0001)	0.00005 (0.0001)	0.00004 (0.0001)
Border Cell	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
Longitude	0.001 (0.0003)	-0.001 (0.001)	-0.001 (0.001)
Longitude ²	0.00003 (0.00002)	-0.00002 (0.00003)	-0.00002 (0.00003)
Longitude ³	-0.00000 (0.00000)	0.00000 (0.00000)	0.00000 (0.00000)
Longitude ⁴	0.00000 (0.00000)	-0.00000 (0.00000)	-0.00000 (0.00000)
Latitude	-0.001 (0.0003)	-0.001 (0.001)	-0.001 (0.001)
Latitude ²	0.00001 (0.00002)	0.00001 (0.00003)	0.00001 (0.00003)
Latitude ³	0.00000 (0.00000)	0.00000 (0.00000)	0.00000 (0.00000)
Latitude ⁴	0.000 (0.00000)	-0.000 (0.00000)	-0.000 (0.00000)
Terrain Ruggedness			-0.00000 (0.00000)
Average Night Lights			-0.001 (0.0003)
Total Population			0.000 (0.000)
Urban Grid Cell			0.006 (0.003)
Country FE		Yes	Yes
Observations	10,158	10,158	10,158
R ²	0.062	0.122	0.124