

Optimal Transport Networks in Spatial Equilibrium*

Pablo D. Fajgelbaum¹ and Edouard Schaal²

¹UCLA and NBER

²CREI, Universitat Pompeu Fabra, Barcelona GSE and CEPR

December 2019

Abstract

We study optimal transport networks in spatial equilibrium. We develop a framework consisting of a neoclassical trade model with labor mobility in which locations are arranged on a graph. Goods must be shipped through linked locations, and transport costs depend on congestion and on the infrastructure in each link, giving rise to an optimal transport problem in general equilibrium. The optimal transport network is the solution to a social planner's problem of building infrastructure in each link. We provide conditions such that this problem is globally convex, guaranteeing its numerical tractability. We also study cases with increasing returns to transport technologies in which global convexity fails. We apply the framework to assess optimal investments and inefficiencies in the road networks of European countries.

JEL Classification: F11, O18, R13

*E-mail: pfajgelbaum@econ.ucla.edu, eschaal@crei.cat. We thank the editor, Liran Einav, and four anonymous referees. We thank Costas Arkolakis, Dave Donaldson, and Jonathan Eaton for their discussions. We thank Thomas Chaney, Arnaud Costinot, Don Davis, Manuel Garcia Santana, Cecile Gaubert, Gueorgui Kambourov, Carlos Llano, Ezra Oberfield, Stephen Redding, Esteban Rossi-Hansberg, and Jonathan Vogel for helpful comments. We also thank Dorian Henricot and Cristiano Mantovani for superb research assistance. Edouard Schaal acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563), the Barcelona GSE (grant SG-2017-07) and the European Research Council Starting Grant 804095.

1 Introduction

Trade costs are a ubiquitous force in international trade and economic geography, as they shape the spatial distributions of prices, real incomes, and trade flows. Transport infrastructure is a key determinant of trade costs.¹ Governments, international organizations, and private companies routinely invest large amounts of resources to improve transport networks within and across countries. How should these investments be allocated? Are observed transport networks suboptimal, and if so how important are these inefficiencies?

In this paper, we develop and apply a framework to study optimal transport networks in general equilibrium spatial models. We solve a global optimization over the space of networks, given any primitive fundamentals, in a general neoclassical framework. In contrast to the standard approach, here trade costs are an outcome rather than a primitive, endogenously responding to fundamentals such as resource endowments and geographic frictions through optimal investments in the transport network. We apply the framework to European road networks, where we assess the aggregate and regional impacts of optimal infrastructure growth, the inefficiencies of observed networks, and the optimal placement of roads as a function of observable regional characteristics.

The point of departure for the framework is a neoclassical economy with multiple goods, factors, and locations, nesting standard trade models (such as the Ricardian, Armington, and factor-endowment models) and allowing for either a fixed spatial distribution of the primary factors (as in international trade models) or for labor to be perfectly mobile (as in economic geography models). The key methodological innovation is that locations are arranged on a graph and goods can only be shipped through connected locations subject to transport costs that depend both on how much is shipped (e.g., because of congestion or decreasing returns to shipping technologies) and on how much is invested in infrastructure (e.g., the number of lanes or the quality of the road). We tackle the planner’s problem of simultaneously choosing the transport network (i.e., the set of infrastructure investments), the allocation of production and consumption, and the gross trade flows across the graph.

Solving this problem may be challenging because of dimensionality—the space of all networks is large—and interactions—an investment in one link asymmetrically impacts the returns to investments across the network. It is also complicated by the potential presence of increasing returns due to the complementarity between infrastructure investments and shipping. We exploit the fact that the planner’s subproblem of choosing gross trade flows is an optimal flow problem on a network, a well understood problem in the operations research and optimal transport literatures. A key insight from these literatures is that the optimal flows derive from a “potential field”—prices in our context—that can be efficiently solved numerically using duality techniques. We make assumptions such that the full planner’s problem, involving the general equilibrium allocation and the network investments alongside the optimal transport, inherits the tractability of optimal flow problems. Our assumptions, including a continuous mapping from infrastructure investments to trade costs and

¹See [Limao and Venables \(2001\)](#) and [Atkin and Donaldson \(2015\)](#). For a review of various determinants of trade costs see [Anderson and Van Wincoop \(2004\)](#).

curvature in the technology to transport goods, ensure that the full planner’s problem is convex and that the set of optimal infrastructure investments can be expressed as a function of equilibrium prices. As a result, we solve the full planner’s problem while avoiding a direct search in the space of networks. Instead, we optimize in the space of equilibrium prices applying the numerical methods typically used for optimal transport problems.

While strong enough congestion in transport guarantees the convexity of the planner’s problem and enables the use of a duality approach, our framework can also be used when congestion is weak or absent—a case that implies increasing returns in the transport technology. We numerically approximate the global solution in non-convex cases by combining the duality approach to obtain the optimal flows with global-search numerical methods that build upon standard simulated annealing techniques. Even though in non-convex cases we only find local optima, the ensuing networks display the qualitative features that one would expect in the presence of economies of scale, such as more concentration in fewer links and a larger amount of zeros.

The framework has enough flexibility to be matched to data on actual transport networks. The quantification relies on two steps. First, the model’s fundamentals can be calibrated such that the solution to the planner’s optimal allocation of consumption, production, and gross flows matches spatially disaggregated data on economic activity given an observed transport network. This step is enabled by the fact that, given the transport network, the welfare theorems hold assuming Pigouvian taxes to correct congestion externalities. Second, assuming a specific technology to build infrastructure makes it possible to undertake counterfactuals involving the optimal network.

We apply these steps in the context of European road networks. We calibrate the productivity and the endowment of non-traded goods such that the model reproduces the observed population and value added at a high spatial resolution separately for each of the 24 countries in our data. We construct a measure of the road infrastructure linking any two contiguous cells in the data and entertain different assumptions on labor mobility and on the returns to infrastructure, encompassing both convex and non-convex cases. We either assume that the observed road network is the outcome of the full planning problem—allowing us to back out these costs from the first-order conditions of the planner’s problem—or use existing estimates for how building costs vary with observable geographic features.

Our counterfactuals in the benchmark parametrization with convex costs imply that, across countries, the average welfare gain from an optimal 50% expansion in the observed road networks and the average welfare loss from road misallocation are on average 2% and range between 0.1% and 7%. The optimal expansion or reallocation of roads reduces regional inequalities in real consumption, reflecting that optimal infrastructure investments reduce dispersion in the marginal utility of consumption of traded commodities. We illustrate the alternative road investment plans implied by the different assumptions and counterfactuals by considering two of the largest economies in our data, France and Spain. We conclude with an exercise involving multiple countries in Western Europe, which highlights the importance of trade across borders and international coordination in infrastructure policy.

The rest of the paper proceeds as follows. Section 2 discusses the connection to the literature. Section 3 develops the framework, establishes its key properties, and discusses the numerical implementation. Section 4 presents simple illustrative examples. Section 5 applies the model to road networks in Europe. Section 6 concludes. We relegate proofs, additional derivations, details of the quantitative exercise, tables, and figures to the appendix.

2 Relation to the Literature

A quantitative literature in international trade and spatial economics studies the role of trade costs in rich geographic settings. [Eaton and Kortum \(2002\)](#) and [Anderson and Van Wincoop \(2003\)](#) developed quantitative versions of the Ricardian and Armington trade models, respectively, allowing counterfactuals with respect to trade costs in multi-country competitive equilibrium. A standard approach to study the gains from market integration is to fit these models to data on the geographic distribution of economic activity, and then ask what would happen if trade costs between specific locations were to change by some predetermined amount.² We develop a different approach to implementing counterfactuals. We first pinpoint the best set of infrastructure investments in a transport network, and then ask what would happen if trade costs were to change in the way implied by the efficient transport network.

Recent studies undertake counterfactuals with respect to the cost of shipping across specific links in models where traders choose least cost routes to ship goods.³ [Allen and Arkolakis \(2014\)](#) measure the aggregate effect of the U.S. highway system, [Donaldson and Hornbeck \(2016\)](#) calculate the historical impact of railroads on the U.S. economy, and [Redding \(2016\)](#) compares the impact of infrastructure changes in models with varying degrees of increasing returns. [Alder \(2019\)](#) simulates counterfactual transport networks in India, [Nagy \(2016\)](#) studies how the development of U.S. railways affected city formation, and [Sotelo \(2016\)](#) simulates the impact of highway investments on agricultural productivity in Peru. Other recent studies allowing for factor mobility and trade frictions within countries include [Bartelme \(2015\)](#), [Caliendo et al. \(2017\)](#) and [Ramondo et al. \(2016\)](#).

Some papers feature an optimization over transport networks. [Alder \(2019\)](#) applies a heuristic algorithm that adds or removes links in a specific order based on their contribution to net aggregate income.⁴ [Felbermayr and Tarasov \(2015\)](#) study optimal infrastructure investments by competing planners in an Armington model where locations are arranged on a line. [Allen and Arkolakis \(2019\)](#) compute the welfare gradient with respect to reductions in the cost of shipping across specific links

²[Costinot and Rodríguez-Clare \(2013\)](#) review the quantitative gravity literature on changes in trade costs focused on measuring gains from international trade. [Redding and Rossi-Hansberg \(2017\)](#) review a body of research using similar frameworks to study counterfactuals involving changes in infrastructure within countries. See [Donaldson \(2015\)](#) and [Redding and Turner \(2015\)](#) for reviews of empirical analyses of actual changes in transport infrastructure, as well as the literature review below for additional references.

³[Chaney \(2014a\)](#) studies endogenous networks of traders in contexts with imperfect information. For a review of recent literature on the role of various types of networks in international trade see [Chaney \(2014b\)](#).

⁴See Section D of the Supplementary Material for a comparison with the [Alder \(2019\)](#) algorithm.

in an Armington model with spillovers.⁵ Their approach is suitable to compute the first-order welfare impact of infrastructure investments around an initially observed allocation.

We solve instead a global optimization over the space of networks in a neoclassical framework. Both our model and the studies cited above include an optimal transport problem, defined as the trader’s problem of choosing least-cost routes across pairs of locations.⁶ In most of the studies cited above, the optimal transport problem does not include congestion and can therefore be solved independently from the general equilibrium. In our context, congestion in transport renders the infrastructure investment problem convex, enabling the search for the global optimum. The least-cost route optimization from the applications of the gravity trade models discussed before corresponds to the solution of our optimal transport problem in the special case without congestion.

As mentioned in the introduction, the planner’s subproblem of choosing how to ship goods given demand, supply and infrastructure formally defines an optimal transport problem. Optimal transport problems were studied early on by [Monge \(1781\)](#) and [Kantorovich \(1942\)](#).⁷ Because we analyze the optimal route problem instead of the direct assignment of sources to destinations, our approach is more closely related to optimal flow problems on a network as in Chapter 8 of [Galichon \(2016\)](#).⁸ Our problem differs from this literature in two important aspects. First, in our model, consumption and production are endogenous because they respond to standard general-equilibrium forces. Instead, the aforementioned optimal flows problems map sources with fixed supply to sinks with fixed demand.⁹ Second, our ultimate focus is on the optimal network investments in the presence of general-equilibrium forces, whereas this literature usually takes the transport costs between links as a primitive. In that regard, the problem that we study is akin to the optimal transport network problems in non-economic environments analyzed in [Bernot et al. \(2009\)](#).

Despite these differences, our model inherits key appealing properties of optimal transport problems. While the optimal transport literature shows that strong duality holds under weak conditions, it holds under some conditions in our model as a special case of convex duality. Hence, our way of embedding an optimal transport problem into a general neoclassical equilibrium model extended with a network design problem does not preclude the validity of key earlier insights from the optimal transport literature. The main benefit of duality, in our context, is a reduction of the search space and substantial gains in computation times.¹⁰

A large body of research estimates how actual changes in transport costs impact economic activity. For instance, [Fernald \(1999\)](#) estimates the impact of road expansion on productivity

⁵[Swisher IV \(2015\)](#) and [Trew \(2016\)](#) allow for endogenous transport costs in different historical contexts.

⁶Note that “optimal transport” refers to the optimal shipping of goods throughout the network. This is one of the subproblems embedded in our framework, alongside the optimal network design problem.

⁷See [Villani \(2003\)](#) for a textbook treatment of the subject.

⁸See also [Bertsekas \(1998\)](#) for a survey of algorithms and numerical methods for optimal flow and transport problems on a network.

⁹See [Beckmann \(1952\)](#) for an early continuous-space example of such an optimal transport problem in economics. See [Carlier \(2010\)](#) and [Ekeland \(2010\)](#) for lecture notes on optimal transport and its connection to economics.

¹⁰A network-design and planning literature in operations research studies related network-design problems in telecommunications and transport without embedding them in general-equilibrium spatial models. See [Ahuja et al. \(1989\)](#).

across U.S. industries; [Chandra and Thompson \(2000\)](#), [Baum-Snow \(2007\)](#) and [Duranton et al. \(2014\)](#) estimate the impact of the U.S. highways on various regional economic outcomes; [Donaldson \(2018\)](#) estimates the impact of access to railways in India; and [Faber \(2014\)](#) estimates the impact of connecting regions to the expressway system in China.¹¹ Our application measures the aggregate country-level welfare gains from optimally expanding current road networks in European countries. In the counterfactuals, we inspect the relationship between optimal infrastructure investment and population growth across regions.

As we apply the model to measure the potential losses from misallocation of roads, the paper is broadly related to studies of the aggregate effects of misallocation such as [Restuccia and Rogerson \(2008\)](#) and [Hsieh and Klenow \(2009\)](#). [Desmet and Rossi-Hansberg \(2013\)](#), [Brandt et al. \(2013\)](#), [Asturias et al. \(2016\)](#), [Fajgelbaum et al. \(2018\)](#), and [Hsieh and Moretti \(2019\)](#) among others focus on geographical misallocation arising from frictions or spatial policies.

3 Model

3.1 Environment

Preferences The economy consists of a discrete set of locations $\mathcal{J} = \{1, \dots, J\}$. We let L_j be the number of workers located in $j \in \mathcal{J}$, and L be the total number of workers. We entertain cases with and without labor mobility. Workers consume a bundle of traded goods and a non-traded good in fixed supply, such as land or housing. Utility of an individual worker who consumes c units of the traded goods bundle and h units of the non-traded good is

$$u = U(c, h), \tag{1}$$

where the utility function U is homothetic and concave in both of its arguments. In location j , per-capita consumption of traded goods is $c_j = C_j/L_j$, where C_j is the aggregate demand of the traded goods bundle in location j .

There is a discrete set of tradable sectors $n = 1, \dots, N$, combined into C_j through a homogeneous of degree 1 and concave aggregator (e.g., a CES aggregator),

$$C_j = D_j(D_j^1, \dots, D_j^N), \tag{2}$$

where D_j^n is sector n 's output used in location j .

Production The supply-side corresponds to a standard neoclassical economy. In addition to labor, there is a fixed supply $\mathbf{V}_j = (V_j^1, \dots, V_j^M)'$ of primary factors $m = 1, \dots, M$ in location j . These factors are immobile across regions but mobile across sectors. The production process may

¹¹See also [Coşar and Demir \(2016\)](#) and [Martincus et al. \(2017\)](#) for empirical studies of how infrastructure investments impact international shipments.

also use goods from other sectors as intermediate inputs. Output of sector n in location j is:

$$Y_j^n = F_j^n(L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n), \quad (3)$$

where L_j^n is the number of workers, $\mathbf{V}_j^n = (V_j^{1n}, \dots, V_j^{Mn})'$ is the quantity of other primary factors, and $\mathbf{X}_j^n = (X_j^{1n}, \dots, X_j^{Nn})$ is the quantity of each sector's output allocated to the production of sector n in location j . The production function F_j^n is either neoclassical (constant returns to scale, increasing and concave in all its arguments) or a constant (endowment economy).

Underlying Graph The locations \mathcal{J} are arranged on an undirected graph $(\mathcal{J}, \mathcal{E})$, where \mathcal{E} denotes the set of edges (i.e., unordered pairs of \mathcal{J}). For each location j there is a set $\mathcal{N}(j)$ of connected locations, or neighbors. Goods can only be shipped through connected locations; i.e., goods shipped from j can be sent to any $k \in \mathcal{N}(j)$, but to reach any $k' \notin \mathcal{N}(j)$ they must transit through a sequence of connected locations. The transport network design problem will consist of determining the level of infrastructure linking each pair of connected locations. A natural interpretation is that j is a geographic unit such as county, $\mathcal{N}(j)$ are its bordering counties, and shipments are done by land. More generally, the neighbors in the model do not need to be geographically contiguous, since it could be possible to ship directly between geographically distant locations by land, air or sea. The fully connected case in which every location may ship directly to every other location, $\mathcal{N}(j) = \mathcal{J}$ for all j , is one special case.

Transport Technology In the model, goods transit through several locations before reaching a point where they are consumed or used as intermediate input. We let Q_{jk}^n be the quantity of goods in sector n shipped from j to $k \in \mathcal{N}(j)$, regardless of where the good was produced. We adopt two alternative specifications for transport costs: iceberg costs without congestion across commodities, and a case with cross-goods congestion. For simplicity of exposition, we discuss the former case here. In section 3.7 we discuss congestion across goods, which will also be the benchmark specification in our applications. Appendix A presents the model in a general case encompassing both formulations.

Transporting each unit of good n from j to k requires τ_{jk}^n units of the good n itself, so that $1 + \tau_{jk}^n$ corresponds to the iceberg cost. This per-unit cost is specified as a function of the total quantity Q_{jk}^n of good n shipped along the link jk and of the level of infrastructure I_{jk} :

$$\tau_{jk}^n = \tau_{jk}(Q_{jk}^n, I_{jk}). \quad (4)$$

The per-unit cost of shipping is increasing in the quantity of commodities shipped:

$$\frac{\partial \tau_{jk}(Q, I)}{\partial Q} \geq 0. \quad (5)$$

This assumption allows for decreasing returns in the shipping sector. We refer to these decreasing

returns as congestion, with the understanding that this concept encapsulates several real-world forces whereby an increase in shipping activity leads to higher marginal transport costs. These forces may include higher travel times or road damage, as well as decreasing returns to scale in transportation due to land-intensive fixed factors such as warehousing or specialized inputs. In short, the more is shipped, the higher the per-unit shipping cost.¹²

We interpret I_{jk} as capturing features that lead to reductions in the cost of transporting goods. For example, when shipping over land, I_{jk} may correspond to whether a road linking j and k is paved, its number of lanes or the availability of roadside services. Hence, we assume:

$$\frac{\partial \tau_{jk}(Q, I)}{\partial I} \leq 0.$$

The transport technology $\tau_{jk}(\cdot)$ is allowed to vary by jk , capturing variation in shipping costs across links for the same quantity shipped and infrastructure. This variation may reflect geographic characteristics such as distance or ruggedness. The per-unit cost function $\tau_{jk}(Q, I)$ may also depend on the direction of the flow; e.g., if elevation is higher in j than k and it is cheaper drive downhill then $\tau_{jk}(Q, I) < \tau_{kj}(Q, I)$.

Flow Constraint In every location there may be tradable commodities being produced, as well as coming in or out. The balance of these flows requires that, for all locations $j = 1, \dots, J$ and commodities $n = 1, \dots, N$,

$$\underbrace{D_j^n + \sum_{n'} X_j^{nn'} + \sum_{k \in \mathcal{N}(j)} (1 + \tau_{jk}^n) Q_{jk}^n}_{\text{Consumption + Intermediate Use + Exports}} \leq \underbrace{Y_j^n + \sum_{i \in \mathcal{N}(j)} Q_{ij}^n}_{\text{Production + Imports}}. \quad (6)$$

The left-hand side of this inequality is location j 's consumption D_j^n of good n , intermediate-input use $X_j^{nn'}$ by each sector n' , exports to neighbors Q_{jk}^n and inputs to the transport sector $\tau_{jk}^n Q_{jk}^n$. These flows are bounded by the local production Y_j^n and imports from neighbors Q_{ij}^n . In standard minimum-cost flow problems this restriction is known as the conservation of flows constraint.

We let P_j^n be the multiplier of this constraint. This multiplier reflects society's valuation of a marginal unit of good n in location j . In the decentralized allocation, this multiplier will equal the price of good n in location j . Therefore, we refer to P_j^n as the price of good n in location j .

Network Building Technology We define the transport network as the distribution of infrastructure $\{I_{jk}\}_{j \in \mathcal{J}, k \in \mathcal{N}(j)}$. The network-design problem will determine this distribution. We assume that building infrastructure requires a resource ("concrete" or "asphalt") in fixed aggregate supply

¹²For a review of the early literature on production-function estimates of returns to scale in the transport sector see [Winston \(1985\)](#). [Newbery \(1988\)](#) theoretically studies road damage externalities, whereby the road damage caused by one vehicle increases the operating costs of subsequent vehicles. [Maibach et al. \(2013\)](#) lists higher travel times, higher accident rate, and road damage as reasons why increased road use may impact transport costs. Other social costs include environmental damage and noise.

K , which can be freely shipped across locations and cannot be used for other purposes. This assumption represents a situation where an amount of resources has been sunk into network-building but must still be allocated across the network. When characterizing the planner's problem, it will lead to the intuitive property that the opportunity cost of building infrastructure in any location is only foregoing infrastructure elsewhere.

The cost of setting up infrastructure may vary across links jk . Specifically, building a level of infrastructure I_{jk} on the link jk requires an investment of $\delta_{jk}^I I_{jk}$ units of K . The network-building constraint therefore is:

$$\sum_j \sum_{k \in \mathcal{N}(j)} \delta_{jk}^I I_{jk} \leq K. \quad (7)$$

We allow the network-design problem to take place when some lower bound for infrastructure \underline{I}_{jk} is already in place. We also allow for an upper bound \bar{I}_{jk} to how much can be built, possibly representing geographic constraints on the capacity to build on a specific link. While the graph $(\mathcal{J}, \mathcal{E})$ is undirected, the infrastructure matrix $\{I_{jk}\}$ defines a weighted directed graph, as there is no need to impose symmetry in investments or costs between connected locations.

While both the transport technology $\tau_{jk}(Q, I)$ in (4) and the infrastructure building cost δ_{jk}^I may vary across links jk , each type of variation reflects different forces. Variation in $\tau_{jk}(Q, I)$ by jk captures how features of the terrain impact per-unit shipping costs given quantity shipped and infrastructure, whereas δ_{jk}^I captures the marginal cost of setting up infrastructure. In the planner's problem below, δ_{jk}^I will not impact the allocation other than through infrastructure I_{jk} .

3.2 Planner's Problem

We solve the problem of a utilitarian social planner who maximizes welfare under two extreme scenarios: labor is either immobile or freely mobile. In the former case, we let ω_j be the planner's weight attached to each worker located in region j . We define each problem in turn.

Definition 1. *The planner's problem with immobile labor is*

$$W = \max_{c_j, h_j, \{I_{jk}\}_{k \in \mathcal{N}(j)}, \{D_j^n, L_j^n, \mathbf{v}_j^n, \mathbf{x}_j^n, \{Q_{jk}^n\}_{k \in \mathcal{N}(j)}\}_n} \sum_j \omega_j L_j U(c_j, h_j)$$

subject to:

(i) *availability of traded commodities,*

$$c_j L_j \leq D_j (D_j^1, \dots, D_j^N) \text{ for all } j;$$

and availability of non-traded commodities,

$$h_j L_j \leq H_j \text{ for all } j;$$

(ii) the balanced-flows constraint,

$$D_j^n + \sum_{n'} X_j^{nn'} + \sum_{k \in \mathcal{N}(j)} (1 + \tau_{jk} (Q_{jk}^n, I_{jk})) Q_{jk}^n \leq F_j^n (L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n) + \sum_{i \in \mathcal{N}(j)} Q_{ij}^n \text{ for all } j, n;$$

(iii) the network-building constraint,

$$\sum_j \sum_{k \in \mathcal{N}(j)} \delta_{jk}^I I_{jk} \leq K,$$

subject to a pre-existing network,

$$0 \leq \underline{I}_{jk} \leq I_{jk} \leq \bar{I}_{jk} \leq \infty \text{ for all } j, k \in \mathcal{N}(j);$$

(iv) local labor-market clearing,

$$\sum_n L_j^n \leq L_j \text{ for all } j;$$

and local factor market clearing for the remaining factors,

$$\sum_n V_j^{mn} \leq V_j^m \text{ for all } j \text{ and } m; \text{ and}$$

(v) non-negativity constraints on consumption, flows, and factor use,

$$\begin{aligned} C_j^n, c_j, h_j &\geq 0 \text{ for all } j \in \mathcal{N}(j), n \\ Q_{jk}^n &\geq 0 \text{ for all } j, k \in \mathcal{N}(j), n \\ L_j^n, V_j^{mn} &\geq 0 \text{ for all } j, m, n. \end{aligned}$$

If labor is freely mobile then the problem is defined as follows.

Definition 2. *The planner's problem with labor mobility is*

$$W = \max_{u, c_j, h_j, \{I_{jk}\}_{k \in \mathcal{N}(j)}, L_j, \{D_j^n, L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n, \{Q_{jk}^n\}_{k \in \mathcal{N}(j)}\}_n} u$$

subject to restrictions (i)-(v) above; as well as:

(vi) free labor mobility,

$$L_j u \leq L_j U(c_j, h_j) \text{ for all } j; \text{ and}$$

(vii) aggregate labor-market clearing,

$$\sum_j L_j = L.$$

This formulation restricts the planner's problem to allocations satisfying utility equalization across locations, a condition that must hold in the competitive allocation. Since U is strictly increasing, restriction (vi) implies that the planner will allocate $u = U(c_j, h_j)$ across all populated locations, and $c_j = 0$ otherwise.

We stop for a moment to discuss the generality achieved in the previous definitions. The

case without labor mobility corresponds to international trade models. The production structure encompasses neoclassical trade models.¹³ When labor mobility is allowed, the model nests urban economics model with a single homogeneous tradeable good in the tradition of [Roback \(1982\)](#). Since we have assumed neoclassical production functions, this formulation does not encompass new economic geography models such as [Krugman \(1991\)](#) and [Helpman \(1998\)](#) nor quantitative extensions with increasing returns ([Allen and Arkolakis, 2014](#); [Redding, 2016](#)). In Section 3.7 we discuss how to implement some cases with increasing returns and provide some examples.

The planner’s problem from Definition 1 can be expressed as nesting three problems:

$$W = \max_{I_{jk}} \max_{Q_{jk}^n} \max_{\{c_j, h_j, D_j^n, L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n\}} \sum_j \omega_j L_j U(c_j, h_j)$$

subject to the constraints. The innermost maximization problem is a standard allocation problem of choosing consumption and factor use subject to the production possibility frontier and the availability of goods in each location. In what follows we refer to it as the “optimal allocation” subproblem. We now discuss some intuitive features of the solution to optimal flows subproblem over Q_{jk}^n and the network design problem over I_{jk} .

Optimal Flows The optimal flow problem that determines the gross flows Q_{jk}^n combines an optimal transport problem—how to map production sources to destinations—and a least-cost route problem with congestion. Under the assumption that domestic absorption D_j^n and production Y_j^n are taken as given, this problem is well known in the optimal transport literature (see, for instance, Chapter 8 of [Galichon \(2016\)](#) or Chapter 4 of [Santambrogio, 2015](#)) and in operations research ([Bertsekas, 1998](#)). A general lesson from these literatures is that these problems are well behaved and admit strong duality. In other words, while the least-cost route and the optimal coupling of sources to destinations may appear to be high-dimensional problems, the solution boils down to finding a “potential field”, meaning one Lagrange multiplier (or price) for each location-good pair, and then expressing the flows as a function of the difference between the multipliers across locations.

The optimal flow problem in our model shares these properties as a special case of convex duality. To understand the solution, remember that P_j^n is the multiplier of the flows constraint (ii), equal to the price of good n in location j in the market allocation according to Proposition 4 below. The first-order condition (A.1) from the planner’s problem in Appendix (A.1) gives the following equilibrium price differential for commodity n between j and $k \in \mathcal{N}(j)$:

$$\frac{P_k^n}{P_j^n} \leq 1 + \tau_{jk}^n + \frac{\partial \tau_{jk}^n}{\partial Q_{jk}^n} Q_{jk}^n, = \text{if } Q_{jk}^n > 0. \quad (8)$$

¹³The Armington model ([Anderson and Van Wincoop, 2003](#)) corresponds to $N = J$ (as many sectors as regions) and $F_j^n = 0$ for $n \neq j$, so that Y_j^n is region j ’s output in the differentiated commodity that (only) region j produces. The Ricardian model corresponds to labor as the only factor of production and linear technologies, $Y_j^n = z_j^n L_j^n$. The specific-factors and Heckscher-Ohlin models are also special cases.




Figure 1: Example of Optimal Flows as a Function of the Price Field

Notes: The picture shows an example of optimal flows in a 15×15 square network with uniform infrastructure across links and one good produced at the origin (blue circle) and consumed in 10 other locations (orange circles). The price in each location is indicated by the z-axis coordinate, and corresponds to a solution of the optimal flow problem given production, consumption and population. The density of flows is represented by the thickness of links and their direction is indicated by the arrows.

Condition (8) is a no-arbitrage condition: the price differential between a location and its neighbors must be less than or equal to the marginal transport cost. From the planner's perspective, this marginal cost takes into account the diminishing returns due to congestion. In the absence of congestion, $\partial\tau_{jk}^n/\partial Q_{jk} = 0$, the price differential would be bounded by the trade cost.

This expression has a number of intuitive properties. Given the network investment, it identifies the trade flow Q_{jk}^n as a function of the price differential, as long as the right-hand side can be inverted. The inversion is possible if the total transport cost $Q_{jk}^n\tau_{jk}^n$ is convex in the quantity shipped. In that case, the gross trade flow Q_{jk}^n is increasing in the price differential. Condition (8) also implies that goods in each sector flow in only one direction, although a link may have flows in opposite directions corresponding to different sectors. In addition, not all goods need to be shipped and some links may be unused despite having positive infrastructure. This may occur if the price gap is not large enough at zero trade to justify shipping. To help visualize the geometric properties of the problem, Figure 1 illustrates how a price field can implement the optimal flows given consumption and production in an example with a single traded commodity. In the example, the good is produced in the location at the origin (blue circle) and demanded in ten locations (red circles). This example uses the functional form for τ_{jk} in (10), which implies that there are some shipments in every link although they become negligible in regions far away from the points of production and consumption. The prices, represented on the z-axis, attain their lowest value at the point of production, and gradually increase with the distance to that point. The optimal flows follow the price gradient according to equation (8) under equality. The consumption locations are local peaks of the price field as long as they do not re-ship the good.

The least-cost route optimization present in the applications of gravity trade models discussed in the literature review corresponds to the solution to this optimal transport problem assuming no congestion. In that case, the optimal transport problem can be solved independently from the rest of the model. In our case, determining the least-cost routes requires information about the flows, the supply, and the demand for each good, which are endogenously solved as part of the allocation. Therefore, the optimal transport problem must be solved jointly with the optimal allocation problem.

Optimal Network Consider now the outer problem of choosing the transport network I_{jk} for all $j \in \mathcal{J}$ and $k \in \mathcal{N}(j)$. Letting P_K be the multiplier of the network-building constraint (iii) (in other words, the shadow price of asphalt), and as long as the (possibly infinite) upper bound \bar{I}_{jk} is not binding, the planner’s choice for I_{jk} implies

$$\underbrace{P_K \delta_{jk}^I}_{\text{Marginal Building Cost}} \geq \underbrace{\sum_n P_j^n Q_{jk}^n \left(-\frac{\partial \tau_{jk}^n}{\partial I_{jk}} \right)}_{\text{Marginal Gain from Infrastructure}}, \quad (9)$$

with equality if there is actual investment, $I_{jk} > \underline{I}_{jk}$. This condition compares the marginal cost and benefits from investing on the link jk . The left-hand side is the opportunity cost of building an extra unit of infrastructure along jk , equal to the marginal value of the scarce resource K in the economy (the multiplier P_K of the network building constraint (7)) times the rate δ_{jk}^I at which that resource translates to infrastructure. The gain from the additional infrastructure, on the right hand side of (9), is the reduction in per-unit shipping costs, $-\partial \tau_{jk} / \partial I_{jk}$, applied to the value of the goods used as inputs in the transport technology.¹⁴

Importantly, the network investment problem inherits the properties that make the optimal transport problem tractable. Substituting the solution for Q_{jk}^n as a function of the price differentials P_k^n / P_j^n into (9) implies that the optimal infrastructure I_{jk} between locations j and k is only a function of prices in each location. Hence, rather than searching in the very large space of all networks, this condition allows us to solve for the optimal investment link by link given the smaller set of all prices. Similar properties can be attained in a case with cross-goods congestion, as described in Appendix Section 3.6.

3.3 Properties

Convexity We establish conditions for the convexity of the planner’s problem, which guarantee its numerical tractability.

¹⁴Various papers measure the first-order impact of changes in bilateral trade costs on world welfare (Atkeson and Burstein, 2010; Burstein and Cravino, 2015; Lai et al. 2015; Allen et al., 2019) or in trade costs in specific links of a transport network on country-level welfare (Allen and Arkolakis, 2019) around an observed equilibrium. The right-hand side of (9) could be used for a similar purpose, given a specific set of changes in trade costs.

Proposition 1. (*Convexity of the Planner’s Problem*) (i) Given the network $\{I_{jk}\}$, the joint optimal transport and allocation problem in the fixed (resp. mobile) labor case is a convex (resp. quasiconvex) optimization problem if $Q\tau_{jk}(Q, I_{jk})$ is convex in Q for all j and $k \in \mathcal{N}(j)$; and (ii) if in addition $Q\tau_{jk}(Q, I)$ is convex in both Q and I for all j and $k \in \mathcal{N}(j)$, then the full planner’s problem including the network design problem from Definition 1 (resp. Definition 2) is a convex (resp. quasiconvex) optimization problem. In either the joint transport and allocation problem, or the full planner’s problem, strong duality holds when labor is fixed.

The first result establishes that the joint optimal allocation and optimal transport subproblems, taking the infrastructure network $\{I_{jk}\}$ as given, define a convex problem for which strong duality holds under the mild requirement that the transport technology $Q\tau_{jk}(Q, I_{jk})$ is (weakly) convex in Q . This property ensures that our specific way of introducing an optimal-transport problem into a general neoclassical economy is tractable. Specifically, it guarantees the existence of Lagrange multipliers that implement the optimal allocation and transport subproblems and ensures the sufficiency of the Karush-Kuhn-Tucker (KKT) conditions, in turn allowing us to apply a duality approach to solve the model numerically—an approach which, as discussed in Section 3.6, substantially reduces computation times. Even if the full problem, including the network design, is not convex due to increasing returns to the network building technology (i.e., if part (ii) of the proposition fails but part (i) holds), a large subset of the full problem can still be solved using these efficient numerical methods.

The second result establishes the convexity of the full planner’s problem, including the network design, under the stronger requirement that the transport cost function $Q\tau_{jk}(Q, I_{jk})$ is jointly convex in Q and I .¹⁵ This condition restricts how congestion in shipping and the returns to infrastructure enter in the transport technology in each link through $\tau_{jk}(Q, I)$. In the absence of congestion (i.e., if $\partial\tau_{jk}/\partial Q = 0$), convexity fails unless τ_{jk} is a constant. The intuition for this convexity requirement is that the model features two complementarity forces between infrastructure investments and commodity shipments: the higher the investment in a link, the lower the transport costs and the higher the flows. In turn, higher shipments lead to more congestion and to more incentives to develop its infrastructure. The global convexity of the transport cost function ensures that the congestion forces eventually dominate and that the solution to the investment problem is interior and stable. Section E of the Supplementary Material develops this point more formally.

Log-Linear Parametrization of Transport Costs A convenient parametrization of (4) is the constant-elasticity transport technology,

¹⁵The proof of Proposition 1 is immediate: given the neoclassical assumptions, the objective function is concave and the constraints are convex, except possibly for the balanced-flows constraint; convexity of the transport cost $Q\tau_{jk}(Q, I_{jk})$ ensures convexity of that constraint as well. In the case with labor mobility, the planner’s problem can only be recast as a quasiconvex optimization problem, but the Arrow-Enthoven theorem for the sufficiency of the Karush-Kuhn-Tucker conditions under quasiconvexity, requiring that the gradient of the objective function is different from zero at the optimal point, is satisfied (Arrow and Enthoven, 1961).

$$\tau_{jk}(Q, I) = \delta_{jk}^\tau \frac{Q^\beta}{I^\gamma} \text{ with } \beta \geq 0, \gamma \geq 0. \quad (10)$$

If $\beta > 0$, this formulation implies congestion in shipping: the more is shipped, the higher the per-unit shipping cost; when $\beta = 0$, the marginal cost of shipping is invariant to the quantity shipped, as in the standard iceberg formulation. In turn, γ captures the elasticity of the per-unit cost to infrastructure. The scalar δ_{jk}^τ captures the geographic frictions that affect per-unit transport costs given the quantity shipped Q and the infrastructure I .

When the transport technology is given by (10), many of the preceding results admit intuitive closed-form formulations. First, the restriction that $Q\tau_{jk}(Q, I)$ is convex in both arguments from Proposition 1 holds if and only if $\beta \geq \gamma$. This inequality captures a form of diminishing returns to the overall transport technology: the elasticity of per-unit transport costs to investment in infrastructure is smaller than its elasticity with respect to shipments.

Second, from the no-arbitrage condition (8), we obtain the following solution for total flows from j to k as function of prices:

$$Q_{jk}^n = \left[\frac{1}{1 + \beta} \frac{I_{jk}^\gamma}{\delta_{jk}^\tau} \max \left\{ \frac{P_k^n}{P_j^n} - 1, 0 \right\} \right]^{\frac{1}{\beta}}. \quad (11)$$

This solution naturally implies that better infrastructure is associated with higher flows given prices and geographic trade frictions. It also shows that the total flows fall with congestion β and increase with the average price differentials. Third, using the log-linear transport technology (10), the optimal level of infrastructure is

$$I_{jk} = \min \left[\max \left(I_{jk}^*, \underline{I}_{jk} \right), \bar{I}_{jk} \right], \quad (12)$$

where I_{jk}^* is the optimal infrastructure (9) arising from the unconstrained optimal-network problem ($\underline{I}_{jk} = 0$ and $\bar{I}_{jk} = \infty$),

$$I_{jk}^* = \left[\frac{\gamma}{P_K} \frac{\delta_{jk}^\tau}{\delta_{jk}^I} \left(\sum_n P_j^n (Q_{jk}^n)^{1+\beta} \right) \right]^{\frac{1}{1+\gamma}}. \quad (13)$$

Given the prices at origin, the optimal infrastructure increases with gross flows. Given these flows, infrastructure also increases with prices at origin, as a higher sourcing cost implies a higher marginal saving from investing. Conditioning on these outcomes, infrastructure increases with δ_{jk}^τ , reflecting that the optimal investments offset geographic trade frictions, and decreases with δ_{jk}^I , reflecting that the investment is smaller where it is more costly to build. Because it satisfies the Inada condition, the log-linear specification (10) implies that the solution to the planner's problem features a positive investment whenever the price of any good varies between neighboring locations, $P_j^n \neq P_k^n$ for any n .

Combining (11) with (13), we reach an explicit characterization of the optimal infrastructure

in each link as a function of prices, elasticities and geographic frictions:

$$I_{jk}^* = \left[\frac{\gamma}{P_K \delta_{jk}^I (\delta_{jk}^\tau)^{\frac{1}{\beta}}} \left(\frac{1}{1+\beta} \sum_{n: P_k^n > P_j^n} P_j^n \left(\frac{P_k^n}{P_j^n} - 1 \right)^{\frac{1+\beta}{\beta}} \right) \right]^{\frac{\beta}{\beta-\gamma}}. \quad (14)$$

where the multiplier P_K is such that the network-building constraint (7) is satisfied.

Proposition 2. (*Optimal Network in Log-Linear Case*) *When the transport technology is given by (10), the full planner's problem is a convex (resp. quasiconvex) optimization problem if $\beta \geq \gamma$. The optimal infrastructure is given by (12).*

Under a general formulation of the transport technology $\tau_{jk}(Q, I)$ and in the absence of a pre-existing network ($\underline{I}_{jk} = 0$) the solution to the full planner's problem may feature no infrastructure (and therefore no trade) in some links, even if prices vary between the nodes connected by those links. However, when the transport technology takes the loglinear form (10), this possibility arises if and only if there are no incentives to trade ($P_j^n = P_k^n$ for all n) due to the Inada condition on I_{jk} in the transport technology (10) and the property that the marginal shipping costs are zero when no shipping is done as long as $\beta \geq 0$, respectively.

Other Convex Transport Technologies We provide two additional tractable transport technologies and the conditions that satisfy their convexity.

1. Exponential: $Q\tau_{jk}(Q, I) = \delta_{jk}^\tau \max(\exp(\beta Q - \gamma I) - 1, 0)$, convex for all $\beta \geq 0, \gamma \geq 0$;
2. CES: $Q\tau_{jk}(Q, I) = \delta_{jk}^\tau \max(Q^\beta - \zeta I^\gamma, 0)^{\frac{1}{\gamma}}$, $\zeta \geq 0$, convex for $0 \leq \gamma \leq 1$ and $\beta \geq \gamma$.

Non-Convexity: the Case of Increasing Returns to Transport When the condition guaranteeing global convexity in Proposition 1 fails, the constraint set in the planner's problem is not convex and the sufficiency of the first-order conditions is not guaranteed. We may nonetheless implement these cases numerically, as we discuss in Section 3.6. Focusing on the log-linear specification (10) introduced above, such nonconvexities arise when the transport technology features economies of scale, $\gamma > \beta$. We now show in a simple special case how the qualitative properties of the optimal network are affected by such economies of scale. In particular, increasing returns to investment in infrastructure create an incentive for the planner to concentrate flows on few links. As a result, the optimal network may take the form of a *tree*, a property already highlighted in other applications of optimal transport such as formation of blood vessels, irrigation or electric power supply systems (Banavar et al., 2000; Bernot et al., 2009).

Proposition 3. *In the absence of a pre-existing network (i.e., $\underline{I}_{jk} = 0, \bar{I}_{jk} = \infty$), if the transport technology is given by (10) and satisfies $\gamma > \beta$, and if there is a unique commodity produced in a single location, the optimal transport network is a tree.*

A tree is a connected graph without loops. Intuitively, under the conditions of the proposition, it is always better to remove alternative paths linking pairs of nodes and concentrate infrastructure investments in fewer links. As a result, in the optimal network a single path connects any two locations, a defining characteristic of a tree. This property is guaranteed to hold when there is only one source for one commodity. In the general case, it may still be optimal to maintain loops but the incentives to concentrate flows on fewer but larger routes remain. In Section 4 we present examples with multiple goods and multiple production locations where, if $\gamma > \beta$, the optimal network is sparser and concentrated on fewer links relative to cases with $\gamma \leq \beta$.

3.4 Decentralized Allocation Given the Network

We establish that the planner's optimal allocation ($\max_{c_j, h_j, D_j^n, L_j^n, \mathbf{v}_j^n, \mathbf{x}_j^n}$) and optimal transport ($\max_{Q_{jk}^n}$) subproblems given the network $\{I_{jk}\}$ correspond to a decentralized competitive equilibrium. For this decentralization, we do not take a stand on whether the network is the result of a planner's optimization.

Given the network, the decentralized economy corresponds to the perfectly competitive equilibrium of a standard neoclassical economy where consumers maximize utility given their budget, producers maximize profits subject to their production possibilities, and goods and factor markets clear. The only less standard feature is the existence of a transport sector with congestion. We assume free entry of atomistic traders into the business of purchasing goods in any sector at origin o and delivering at destination d for all $(o, d) \in \mathcal{J}^2$. The traders are price-takers and use a constant-returns to scale shipping technology. Each trader has a cost equal to $\tau_{jk}^n q_{jk}^n$ of delivering q_{jk}^n units of good n from j to $k \in \mathcal{N}(j)$ and takes the iceberg trade cost τ_{jk}^n as given, although this trade cost is determined endogenously through (10) as function of the aggregate quantity shipped.

As long as there is congestion in shipping, the traders will engage in an inefficient amount of shipping. We assume that the market allocation features policies that correct this externality. Specifically, the shipments of commodity n over link jk are subject to ad-valorem taxes $\varepsilon_{Q, jkn}^T \tau_{jk}^n$ on their value at j . Consider then a trader purchasing good n at location o and delivering it to location d . This company maximizes profits by optimizing over the route $r = (j_0, \dots, j_\rho) \in \mathcal{R}_{od}$, where j_0, \dots, j_ρ is a sequence of nodes from o to d and \mathcal{R}_{od} is the set of all such routes. The optimal route r_{od}^n maximizes the per-unit profits:

$$\pi_{od}^n = \max_{r=(j_0, \dots, j_\rho) \in \mathcal{R}_{od}} p_d^n - \underbrace{p_o^n}_{\text{Sourcing Costs}} - \underbrace{\sum_{k=0}^{\rho-1} p_{j_k}^n \tau_{j_k j_{k+1}}^n}_{\text{Transport costs}} - \underbrace{\sum_{k=0}^{\rho-1} p_{j_k}^n t_{j_k j_{k+1}}^n}_{\text{Taxes}}, \quad (15)$$

where p_j^n is the price of good n in location j in the market allocation. A shipper from o to d purchases each unit at price p_o^n and obtains the price p_d^n . In addition, shippers must pay the transport costs $p_{j_k}^n \tau_{j_k j_{k+1}}^n$ as well as the "toll" $p_{j_k}^n t_{j_k j_{k+1}}^n$ on each segment. In the absence of tolls, the shipping cost from o to d would equal the total iceberg cost, and the solution would correspond

to a standard least-cost route optimization.

To define the competitive equilibrium, we must also allocate the returns to factors other than labor. Under no labor mobility we assume that, in addition to the wage, each worker in location j receives a transfer t_j such that $\sum_{j=1}^J t_j L_j = \Pi$, where Π is an aggregate portfolio including ownership of fixed factors and government transfers. Hence, workers are rebated all tax revenues and own the primary factors and non-traded goods in the economy. This formulation allows for trade imbalances, which are needed to implement the planner's allocation.

Since they are standard, we relegate the definitions of the competitive allocation with and without labor mobility to Definition 3 in the appendix. Using that definition, we establish that the welfare theorems given the transport network hold.

Proposition 4. *(First and Second Welfare Theorems) If the tax on shipments of product n from j to k is*

$$t_{jk}^n = \varepsilon_{Q,jkn}^\tau \tau_{jk}^n,$$

where $\varepsilon_{Q,jkn}^\tau = \partial \log \tau_{jk}^n / \partial \log Q_{jk}^n$, then:

(i) *if labor is immobile, the competitive allocation coincides with the planner's problem under specific planner's weights ω_j . Conversely, the planner's allocation can be implemented by a market allocation with specific transfers t_j ; and*

(ii) *if labor is mobile, the competitive allocation coincides with the planner's problem if and only if all workers own an equal share of fixed factors and tax revenue regardless of their location, i.e., $t_j = \frac{\Pi}{L}$.*

In either case, the price of good n in location j , p_j^n , equals the multiplier on the balanced-flows constraint in the planner's allocation, P_j^n .

These results are useful for bringing our model to the data. Under the assumption that the observed allocation corresponds to the decentralized equilibrium, the first welfare theorem enables us to calibrate the model using the planner's solution to the optimal allocation and optimal transport subproblems given the network. In Section 3.7, we discuss how to calibrate the model assuming that the observed market allocation does not feature policies correcting the externality.. We note that the optimal allocation can be equivalently implemented by per-unit toll $\theta_{jk}^n = p_{jk}^n \varepsilon_{Q,jkn}^\tau \tau_{jk}^n$.

3.5 Decentralization of Network Investments

We now discuss a market structure that efficiently decentralizes the infrastructure investments. Consider a decentralized allocation as in Definition 3, including tolls. Suppose that, in addition, I_{jk} is endogenously determined by a link-specific builder who is granted the right to build infrastructure and receives in exchange a per-unit toll θ_{jk}^n . Builders can purchase the "asphalt" K at a price p_K , the stock of K may be initially owned by the government or by private individuals, and the price p_K adjusts such that the market for K clears. The builders will solve the problem

$$\max_{I_{jk}} \sum_n \theta_{jk}^n Q_{jk}^n(I_{jk}) - p_K \delta_{jk}^I I_{jk},$$

where $Q_{jk}^n(I_{jk})$ is the quantity of good n consistent with zero profits of shipping companies on the link jk given infrastructure I_{jk} and prices. The builders internalize that by adding infrastructure they can increase the flow of goods through their link, but we assume that they do not internalize general-equilibrium impacts on commodity prices. Now, the specific transfers t_j exhaust a portfolio Π which, in addition to fixed factors and government transfers, also includes ownership of K and net profits of builders.

We then obtain the following result.

Proposition 5. *If the global convexity condition of Proposition 1 is satisfied and the toll θ_{jk}^n is consistent with the optimal Pigouvian tax ($\theta_{jk}^n = P_j^n \varepsilon_{Q,jkn}^\tau \tau_{jk}^n$), then the decentralized infrastructure choice implements the optimal network investment.*

This result echoes the self-financing theorem of [Mohring and Harwitz \(1962\)](#) who showed that revenues from optimal congestion taxes are sufficient to cover capital costs of roads when the transport cost can be expressed as a function of the ratio Q/I . Our result is not restricted to the case in which the transport cost function is homogeneous of degree 0. The global convexity condition of Proposition 1 ensures the sufficiency of the first-order conditions in implementing the optimal allocation. In this general case, however, lump-sum transfers to builders may be needed to ensure participation.

3.6 Numerical Implementation

In this section we broadly discuss our numerical implementation and relegate details to Section D of the Supplementary Material.¹⁶

Convex Cases Under the conditions of Proposition 1, the full planner’s problem is a convex optimization problem and the KKT conditions are both necessary and sufficient. The system of first-order conditions is, however, a large system of non-linear equations with many unknowns. Gradient-descent based algorithms make large-scale convex optimization problems like ours numerically tractable, meaning that these algorithms are guaranteed to converge to the unique global optimum ([Boyd and Vandenberghe, 2004](#)).¹⁷

Our problem can be tackled numerically using two equally valid approaches. The first one is to feed the numerical solver the *primal* problem, meaning the full planner’s problem exactly as written in Definition 1. Specifically, letting \mathcal{L} be the Lagrangian of the planner’s problem as a function of the controls $\mathbf{x} = (c_j, h_j, D_j^n, L_j^n, \mathbf{V}_j^n, Q_{jk}^n, \dots)$ and the multipliers $\boldsymbol{\lambda} = (P_j^n, \dots)$, the primal consists of solving the saddle-point problem

$$\sup_{\mathbf{x}} \inf_{\boldsymbol{\lambda} \geq 0} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}).$$

¹⁶A Matlab toolbox implementing our model with detailed documentation and examples is available on the authors’ websites.

¹⁷We use the open-source large-scale optimization package IPOPT (<https://projects.coin-or.org/Ipopt>) which is based on an interior point method and is able to handle thousands of variables as long the problem is sufficiently sparse. The software converges in polynomial time ([Nesterov and Nemirovskii, 1994](#)).

The second approach, preferred in the optimal transport literature, is to solve instead the *dual* problem obtained by inverting the order of optimization, i.e.,

$$\inf_{\lambda \geq \mathbf{0}} \sup_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda).$$

In our context, the convexity of the full planner’s problem without labor mobility ensures that the dual coincides with the primal (Proposition 1), i.e., strong duality holds. The advantage of the dual is that we can use the first-order conditions from the optimal transport and the optimal investment problems, (8) and (13), as well as those from the neoclassical allocation problem, to express the control variables as functions of the multipliers, $\mathbf{x}(\lambda)$. The remaining minimization problem, $\inf_{\lambda \geq \mathbf{0}} \mathcal{L}(\mathbf{x}(\lambda), \lambda)$, is a convex minimization problem over fewer variables, subject to non-negativity constraints only.

Non-Convex Cases When the condition stated in Proposition 1 fails, the full planner’s problem is no longer globally convex, and the method described above is not guaranteed to find the global optimum. To solve for such non-convex cases, we exploit the property, stated at the beginning of Proposition 1, that the joint neoclassical allocation and optimal transport problem nested within the planner’s problem is convex as long as $Q\tau_{jk}(Q, I_{jk})$ is convex in Q . This condition is weaker and holds under the log-linear specification as long as $\beta \geq 0$, including the standard case without congestion ($\beta = 0$). We combine the primal and dual approaches to solve for the joint neoclassical allocation and optimal transport problems with an iterative procedure over the infrastructure investments. Specifically, starting from a guess on the network investment I_{jk} , we solve for the optimum over $c_j, h_j, D_j^n, L_j^n, \mathbf{V}_j^n$ and Q_{jk}^n , and then use the optimal network investment condition (9) to obtain a new guess over I_{jk} , and then repeat until convergence. We then refine the solution using a simulated annealing method that perturbs the local optimum and gradually reaches better solutions. Appendix D provides details.

3.7 Alternative Assumptions

Congestion Across Goods We have assumed that congestion only applies within good types. A natural assumption is that congestion takes place across goods. A simple way to incorporate this feature while preserving the convexity of the problem is to assume that the per-unit cost τ_{jk}^n is denominated in units of the bundle of traded goods aggregated through $D_j(\cdot)$ rather than in units of the good itself. We assume that transporting each unit of good n from j to $k \in \mathcal{N}(n)$ requires

$$\tau_{jk}^n = m^n \tau_{jk}(Q_{jk}, I_{jk}) \tag{16}$$

units of the traded goods bundle, where m^n measures product-specific characteristics such as weight or volume, and $Q_{jk} = \sum_{n=1}^N m^n Q_{jk}^n$ is the total weight or volume transported from j to k . Then, the number of units of the traded goods bundle D_j used to transport goods from j to its neighbors

is

$$T_j = \sum_{k \in \mathcal{N}(j)} \tau_{jk} (Q_{jk}, I_{jk}) Q_{jk}.$$

After properly adjusting the resource constraints in the definition of the planner’s problem, the convexity of the full planner’s problem is preserved under the same conditions stated in Proposition 1.

Appendices A.1 and A.2 present the definition and first-order conditions of the planner’s problem in a general case encompassing iceberg costs with own-good congestion, as well as this alternative formulation with congestion across goods. As it is more realistic, we adopt the case with congestion across goods as the benchmark in our application.

Externalities and Inefficiencies in the Market Allocation In Section 3.4, we assumed that the decentralized allocation is efficient. However, in some cases it may be desirable to consider an inefficient market allocation. For example, a standard formulation with agglomeration spillovers is to assume that the production technology is $Y_j^n = F_j^n (L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n; L_j)$, where the spillover from the total number of workers L_j on output Y_j^n is not internalized in the market allocation. Another common formulation is to assume externalities in the consumption of amenities entering through utility, $U(c_j, h_j; L_j)$. Similarly, without the Pigouvian taxes t_{jk}^n correcting the congestion externality in shipping, the market allocation is inefficient. In these cases, it is in principle still possible to calibrate the model and undertake counterfactuals using a “fictitious” planner who ignores the dependence of Y_j^n on L_j or of τ_{jk}^n on Q_{jk}^n . For example, in the case of production spillovers, the fictitious planner’s problem is defined as in Definition 2 under the assumption that the vector of aggregate population levels $\bar{\mathbf{L}} = \{\bar{L}_j\}$ in $Y_j^n = F_j^n (L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n; \bar{L}_j)$ is taken as given.¹⁸ As long as $F_j^n(\cdot)$ is neoclassical given \bar{L}_j , the statement in part (i) of Proposition 1, establishing convexity of the planner’s problem given the network, remains the same given \bar{L}_j . However, this approach requires solving an additional loop imposing that the vector of population $\mathbf{L} = \{L_j\}$ that solves the fictitious planner’s problem coincides with the aggregate distribution $\bar{\mathbf{L}}$ taken as given by the planner. Every distribution of population $\bar{\mathbf{L}}$ satisfying this fixed point problem corresponds to a market allocation and vice-versa.¹⁹ An important caveat, however, is that we can no longer establish the general convexity of the problem corresponding to part (ii) of Proposition 1. Section F of the Supplementary Material explains how to implement these cases along with a method to derive the optimal infrastructure gradient and conduct local optimization.

¹⁸The fictitious planner problem is defined exactly as in Definition 2 with $U(c_j, h_j; \bar{L}_j)$ in the case of consumption externalities, taking \bar{L}_j as given, or with $\tau_{jk} (Q_{jk}^n, I_{jk})$ in the case of congestion externalities, taking the shipments $\bar{\mathbf{Q}} = \{Q_{jk}^n\}_{j,k,n}$ as given.

¹⁹Whether such a fixed point exists depends on the specifics of the environment. It is beyond the scope of this paper to determine the conditions under which that is the case, but we note that, given the network $\{I_{kl}\}$, our environment can accommodate the specific parametric assumptions that guarantee existence or uniqueness of an inefficient decentralized allocation found in the previous literature. E.g., see Allen and Arkolakis (2014) for conditions that lead to existence and uniqueness in an Armington model with labor mobility and size spillovers.

4 Illustrative Examples

In this section we implement examples that illustrate the basic economic forces captured by the framework and its potential uses. All the figures can be found in Section C of the Supplementary Material. We start with an endowment economy without labor mobility and only one traded and one non-traded good in a symmetric graph. Then, we progressively move to more complex cases with multiple locations in asymmetric spaces, multiple sectors, labor mobility, and heterogeneous building costs due to geographic features. Throughout the examples, we illustrate the contrast between the globally optimal networks in convex cases, where the congestion forces dominate the returns to network building, and the approximate optimal networks in cases where global convexity of the planner’s problem fails. In all the examples, preferences are CRRA over a Cobb-Douglas bundle of traded and non-traded goods, $U = (c^\alpha h^{1-\alpha})^{1-\rho} / (1-\rho)$ with $\alpha = \frac{1}{2}$ and $\rho = 2$. There is a single factor of production, labor, and all technologies are linear. We adopt the constant-elasticity functional forms (10) for the transport and network-building technologies.

4.1 One Good on a Regular Geometry

Comparative Statics over K in a Symmetric Network To start we impose $\beta = \gamma = 1$, which lies at the boundary of the parameter space guaranteeing global convexity. We assume a single good, no labor mobility and no geographic frictions, $\delta_{jk}^\tau = \delta_{jk}^I = \text{Distance}_{jk}$.

Figure A.3 presents a network with 9×9 locations uniformly distributed in a square, each connected to 8 neighbors. All fundamentals except for productivity are symmetric: $(L_j, H_j) = (1, 1)$. Labor productivity is $z_j = 1$ at the center and 10 times smaller elsewhere.

Figure A.4 shows the globally optimal network when $K = 1$ (panel (a)) and when $K = 100$ (panel (b)). The upper-left figure in each panel displays the optimal infrastructure network I_{jk} corresponding to (13). The optimal network investments radiate from the center, and so do shipments. The bottom figures in each panel display the multipliers of the flows constraint (6)—the prices in the market allocation—and consumption. Because tradable goods are scarcer in the outskirts, marginal utility is higher and so are prices. As the aggregate investment grows from $K = 1$ to $K = 100$, the network grows into the outskirts and the differences in the marginal utility shrink. Panels (c) and (d) display the spatial distribution of prices and consumption. As the network grows, relative prices and consumption converge, and spatial inequalities are reduced.

Randomly Located Cities and Non-Convex Cases We now explore more complex networks and non-convex cases. Figure A.5 shows 20 “cities” randomly located in a space where each location has six neighbors. Population is $L_j = 1$ in each city and 0 otherwise. Productivity is again ten times larger at the center. The top panel shows the infrastructure and commodity flows in the optimal network. The optimal network radiates from the center to reach all destinations. Due to congestion, some destinations are reached through multiple routes. However, to reach some faraway locations such as the one in the northwest, only one route is built.

The middle panel inspects the same spatial configuration but assumes $\gamma = 2$. Now, the sufficient condition for global convexity from Proposition 1 fails. We see a qualitative change in the shape of the network. Due to increasing returns to network building, fewer roads are built but each has higher capacity. In particular, there is now only one route linking any two destinations, consistent with the no-loops result in Proposition 3.

Because in the non-convex network we can only guarantee convergence to a local optimum, we refine the solution by applying the numerical approach discussed in Supplementary Material Section D involving simulated annealing. The bottom panel compares the non-convex network before and after the annealing refinement. The refined network economizes on the number of links, leading to a welfare increase but preserving the no-loops property.

4.2 Many Sectors, Labor Mobility, and Non-Convexity

We now further introduce multiple traded goods and labor mobility. We allow for 11 traded commodities, one “agricultural” good (good 1) that may be produced everywhere outside of “cities” ($z_j^1 = 1$ in all “countryside” locations) and ten “industrial” goods, each produced in one random city only ($z_j^n = 1$ in only one city j and $z_j^n = 0$ otherwise). These goods are combined via a constant elasticity of substitution aggregator with elasticity of substitution $\sigma = 2$. Labor continues to be the sole factor of production, but is now mobile. The supply of the non-traded good is uniform, $H_j = 1$ for all j .

Figure A.6 shows the convex case ($\beta = \gamma = 1$). The first panel shows the optimal network. In the figure, each circle’s size denotes the population share. The remaining figures show the shipments of each good, with the circle sizes representing the shares in total production for the corresponding good. Figure A.7 shows the optimal network with annealing in the nonconvex case when $\gamma = 2$.

In these examples, we observe complex shipping patterns. There are bilateral flows over each link, now involving several commodities. Overall, the optimal network in the first panel reflects the spatial distribution of comparative advantages. Since industrial goods are relatively scarce, wages and population are higher in the cities that produce them. Due to the need to ship industrial goods to the entire economy and to bring agricultural goods to the more populated cities, the transport network has better infrastructure around the producers of industrial products. As Panel (a) of each figure illustrates, the optimal network links the industrial cities through wider routes branching out into the countryside. The agricultural good, being produced in many locations, travels short distances and each industrial city is surrounded by its agricultural hinterland.

The comparison between Figures A.6 and A.7 confirms the intuition that, in the presence of economies of scale in transportation, the optimal network becomes more skewed towards fewer but wider “highways”. Note, however, that the tree property from Proposition 3 no longer holds because there are multiple goods.

4.3 Geographical Features

We now show how the framework can accommodate geographical features like mountains and rivers. To highlight the role of these frictions we revert to a case with a single good and no factor mobility. Panel (a) of Figure A.8 shows 20 cities randomly allocated in a space where each location is connected to 8 other locations. Population equals 1 in all cities and productivity is the same everywhere (equal to 0.1) except in the central city, displayed in red, where it is 10 times larger. Each city’s size in the figure varies in proportion to consumption.

As implied by condition (13), the optimal infrastructure in a given link depends on the link-specific building cost δ_{jk}^I . In panel (a) we show the optimal network under the assumption that the cost of building infrastructure is proportional to the Euclidean distance:

$$\delta_{jk}^I = \delta_0 \text{Distance}_{jk}^{\delta_1}. \quad (17)$$

As in our first set of examples, the optimal network radiates from the highest-productivity city to alleviate differences in marginal utility.

In panel (b), we add a “mountain” by adding an elevation dimension to each link and re-configuring the building cost as

$$\delta_{jk}^I = \delta_0 \text{Distance}_{jk}^{\delta_1} \left(1 + |\Delta \text{Elevation}|_{jk}\right)^{\delta_2}. \quad (18)$$

Because it is more costly to build through the mountain, the optimal network circles around it to reach the cities in the northeast. Because more resources are invested in that region, the network shrinks elsewhere.

In the subsequent figures, we either increase or decrease the cost of building the network in specific links. Specifically, we allow for the more general specification:

$$\delta_{jk}^I = \delta_0 \text{Distance}_{jk}^{\delta_1} \left(1 + |\Delta \text{Elevation}|_{jk}\right)^{\delta_2} \delta_3^{\text{CrossingRiver}_{jk}} \delta_4^{\text{AlongRiver}_{jk}}. \quad (19)$$

In panel (c) we include a river and assume that $\delta_3 = \delta_4 = \infty$, so that investing in infrastructure either across or along the river is prohibitively costly. The optimal network linking cities across the river can only be built through the patch of dry land. In that natural crossing there is a “bottleneck” and a large infrastructure investment takes place.

In panel (d) we assume instead that no dry patch exists but that building bridges is feasible, $1 < \delta_3 < \infty$. Now, the planner builds two bridges, directly connecting the pairs of cities across the river. Panel (e) further allows for transport capacity along the river ($\delta_4 < \infty$). The planner retains the bridges, but now faraway locations in the southeast are reached by water instead of via ground transport.

Finally, panel (f) shows the non-convex case, $\gamma = 2 > \beta$, implemented through the combination of first-order conditions and simulated annealing described in Section 3.6. Now, a unique route links any two cities and fewer roads are built.

5 Road Network Expansion and Misallocation in Europe

We apply the framework for quantitative analysis of road networks in European countries. We start by describing the steps to represent data on economic activity and road networks in terms of the graph of our model. Then we choose the fundamentals to match the observed distribution of economic activity within each country. We conclude by implementing counterfactuals involving the optimal transport network. We implement the calibrations and counterfactuals country by country. In a final exercise, we also implement the analysis simultaneously for a connected set of countries in continental Europe.

5.1 Data

Sources We combine geocoded data on road networks, population, and income across European countries. The road network data is from EuroGlobalMap (EGM) by EuroGeographics.²⁰ The dataset combines shapefiles on the road network from each European country’s mapping and cadastral agencies, and it includes all major highways and roads connecting populated areas.²¹ For example, the French road network is represented by 38668 segments of active roads connecting 159258 geographic points with a total length of about 130000 km.

We perform the calibration and counterfactual analysis separately for each of the 24 countries included in EGM for which data on number of lanes is available. This set includes rich and poor countries, as well as geographically large and small. Table A.1 in Appendix B reports the list of countries with summary statistics about the size and average features of their road networks, the number of cells, and features of their discretized road networks.

An appealing feature of this dataset is that each segment of a road network has information about objective measures of road quality including type of road use (national, primary, secondary, or local), number of lanes, and whether it is paved or includes a median. National roads encompass each country’s highway system.²² On average across the countries in our data, national roads represent 10% of the road network, feature twice as many lanes per kilometer as other types of roads, are always paved (while 94% of the non-national networks are), and are more likely to include a median (87% relative to 4% of non-national networks). Since the roads labeled as primary, secondary and tertiary have very similar characteristics along these dimensions, we bundle them into a single “non-national roads” category.

We use population data from NASA-SEDAC’s Gridded Population of the World (GPW) v.4 and value added from Yale’s G-Econ 4.0. Both datasets correspond to the year 2005. The GPW population data is reported for 30 arc-second cells (approximately 1 kilometer) and the G-Econ

²⁰This product includes Intellectual Property from the European National Mapping and Cadastral Authorities and is licensed on behalf of these by EuroGeographics. Original product is available for free at www.eurogeographics.org. Terms of the license are available at <https://eurogeographics.org/services/open-data/topographic-data/>.

²¹We use GlobalMap data v8, corresponding to the year 2016, which is the earliest year we have access to. We only use road segments that are reported as “operational”.

²²E.g., roads labeled as national in the data include the *Autobahn* highway system in Germany, *autovias* and *autopistas* in Spain, and the *autoroute* system in France.

value-added data is reported for 1 arc-degree cells (approximately 100 km). To implement the model we must take a stand on the geographic units corresponding to each node. To strike a balance between the high spatial resolution of the EGM and GPW datasets and the coarser resolution of the G-Econ dataset, in most countries we use 0.5 arc-degree cells (approximately 50 km by 50 km) as benchmark. We adopt squares as geographic units so that the boundaries of the geographic units in the G-Econ dataset coincide with ours. We allocate population to each 0.5-degree cell by aggregating the smaller cells in GPW and we apportion income from the G-Econ cells according to the GPW-based population measure. In a few countries we use smaller or larger cells in order to allow for either a significant number of cells or avoid having a very large number.²³ We denote the population and value added observed in each cell j of each country by L_j^{obs} and GDP_j^{obs} .

The resulting number of cells in each country is reported in Table A.1. In the 19 countries where the NUTS subdivision of geographic units is available, the number of cells is larger than the number of level-2 NUTS regions. In most countries it is also larger than the number of level-3 NUTS.²⁴

Underlying Graph Using these data we construct empirical counterparts to the underlying geography $(\mathcal{J}, \mathcal{E})$ corresponding to the locations and links in the graph of our model, as well as an observed measure of infrastructure I_{jk}^{obs} for each link.

To define the set of nodes \mathcal{J} in each country, we use the GPW data to locate the population centroid of each cell. The population centroids are usually very close to a node on the road network. We relocate each population centroid to the closest point on a national road crossing through the cell, or on other types of roads if no national roads cross through the cell.²⁵ We define the observed population and income of each node $j \in \mathcal{J}$ to be equal to the total income GDP_j^{obs} and the population L_j^{obs} of the cell that contains it.

In turn, we define the set of edges \mathcal{E} as all the links between nodes in contiguous cells. This step defines a set of up to 8 neighbors $\mathcal{N}(j)$ for each node $j \in \mathcal{J}$: the 4 nodes in horizontal or vertical neighbors and the 4 nodes along the diagonals.

Discretized Road Network To construct a measure of infrastructure corresponding to I_{jk} in our model, we first aggregate the observed attributes of the road network over the actual roads linking each $j \in \mathcal{J}$ and $k \in \mathcal{N}(j)$. We use information on whether each segment s on the actual

²³We use the default 0.5 arc-degree cells in 17 countries (Austria, Belgium, Czech Republic, Denmark, Georgia, Hungary, Ireland, Latvia, Lithuania, Macedonia, Moldova, Netherlands, Northern Ireland, Portugal, Slovakia, Slovenia and Switzerland). Whenever assuming 0.5 arc-degree cells would lead to more than 200 cells, we use 1 arc-degree cells (Finland, France, Germany, Italy and Spain); and whenever doing so would lead to less than 20 cells, we use 0.25 arc-degree cells (Luxembourg and Cyprus).

²⁴NUTS (Nomenclature of Territorial Units for Statistics) is a standard developed by the European Union to divide the territory. NUTS 2 correspond to “basic regions for the application of regional policies”, and NUTS 3 correspond to “small regions for specific diagnoses” (<https://ec.europa.eu/eurostat/web/nuts/background>). Excluding overseas territories, Spain has 15 level-2 NUTS (autonomous communities) and 47 level-3 NUTS (provinces), whereas our partition has 61 cells. In France there are 21 level-2 NUTS (regions) and 94 level-3 NUTS (departments), whereas our partition has 74 cells.

²⁵On average across countries, the relocation across all cells within a country is 5.3 km.

road network belongs to a national road and its number of lanes. We define the average number of lanes and average road type for the link between j and k as follows:

$$\begin{aligned} lanes_{jk} &= \sum_{s \in S} \omega_{jk}(s) lanes(s), \\ nat_{jk} &= \sum_{s \in S} \omega_{jk}(s) nat(s), \end{aligned}$$

where $lanes(s)$ is the number of lanes on each segment s on the actual road network S , $nat(s)$ indicates whether segment s belongs to a national road, and $\omega_{jk}(s)$ is the weight attached to the infrastructure of each segment when computing the level of infrastructure from j to k . The weights $\omega_{jk}(s)$ should be larger on segments of the road network that are more likely to be used when shipping from j to k , and equal to zero for all $s \in S$ if no direct route exists linking j and k . We define $\omega_{jk}(s)$ based on the fraction of the cheapest path $\mathcal{P}(j, k)$ from j to k corresponding to that segment:

$$\omega_{jk}(s) = \begin{cases} \frac{length(s)}{\sum_{s' \in \mathcal{P}(j, k)} length(s')} & s \in \mathcal{P}(j, k) \\ 0 & s \notin \mathcal{P}(j, k) \end{cases}$$

where $length(s)$ is the length of segment s and $\mathcal{P}(j, k)$ is the cheapest path from j to k on the actual road network.²⁶ We follow these steps as long as the cheapest path does not stray from the cells containing j and k .²⁷ When that happens, we assume that no direct path from j to k exists in the actual road network, $\mathcal{P}(j, k) = \emptyset$, in which case $\omega_{jk}(s) = 0$ for all segments $s \in S$.

We define the observed measure of infrastructure I_{jk}^{obs} for each $j \in \mathcal{J}$ and $k \in \mathcal{N}(j)$ by letting $I_{jk}^{obs} = lanes_{jk}$ for national roads and $I_{jk}^{obs} = lanes_{jk} \times \kappa$ for non-national roads, where $\kappa < 1$ captures the smaller cost of non-national roads. We set $\kappa = 1/5$, which corresponds to the cost of road construction and maintenance per kilometer on trunk roads relative to federal motorways in Germany in 2007, as reported by Doll et al. (2008). We impose $I_{jk}^{obs} = I_{kj}^{obs}$, implying that infrastructure applies equally in either direction. In sum, we construct the observed infrastructure I_{jk}^{obs} as the average number of national road lanes over the path from j to k on the actual road network, if a direct path exists.²⁸

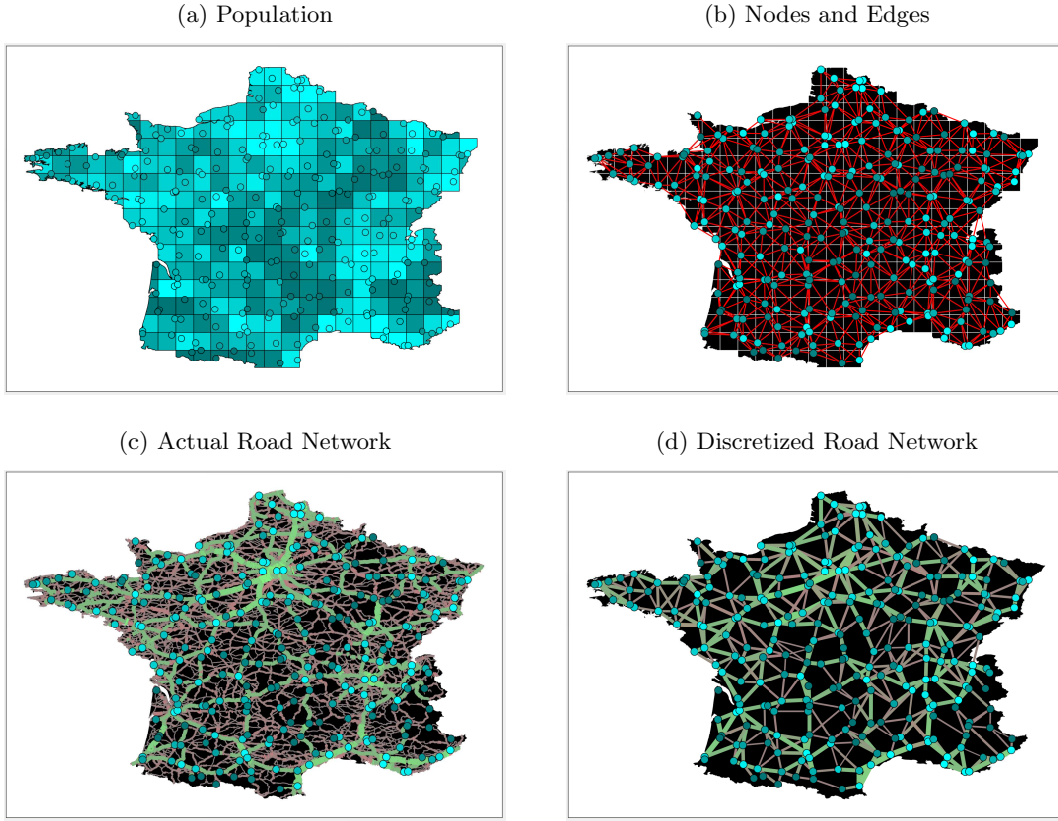
Examples: France and Spain Figures 2 and 3 represent each of the steps described above for two large countries in our data, France and Spain. Panel (a) of each panel shows the discretized map and associated population. Brighter cells are more populated, corresponding to higher deciles

²⁶This step does not use the model. In this step, for each pair of nodes $j \in \mathcal{J}$ and $k \in \mathcal{N}(j)$ we ask: what are the average characteristics (number of lanes and type of road) of the existing route connecting these two locations in the real world? To do this, we must choose some route connecting the pair of locations in the real world. We use the cheapest-route criterion as a way to choose this route. The cheapest path is constructed by weighting each segment s by its road user cost based on data from Combes and Lafourcade (2005) and other sources. See Appendix B for details.

²⁷We classify a path from j to k as straying from the cells containing j and k if more than 50% of the path steps over cells that do not contain j or k .

²⁸Across connected nodes in the discretized network, there is a correlation of 0.67 between I_{jk}^{obs} and the speed on the quickest path according to GoogleMaps.

Figure 2: Discretization of the French Road Network

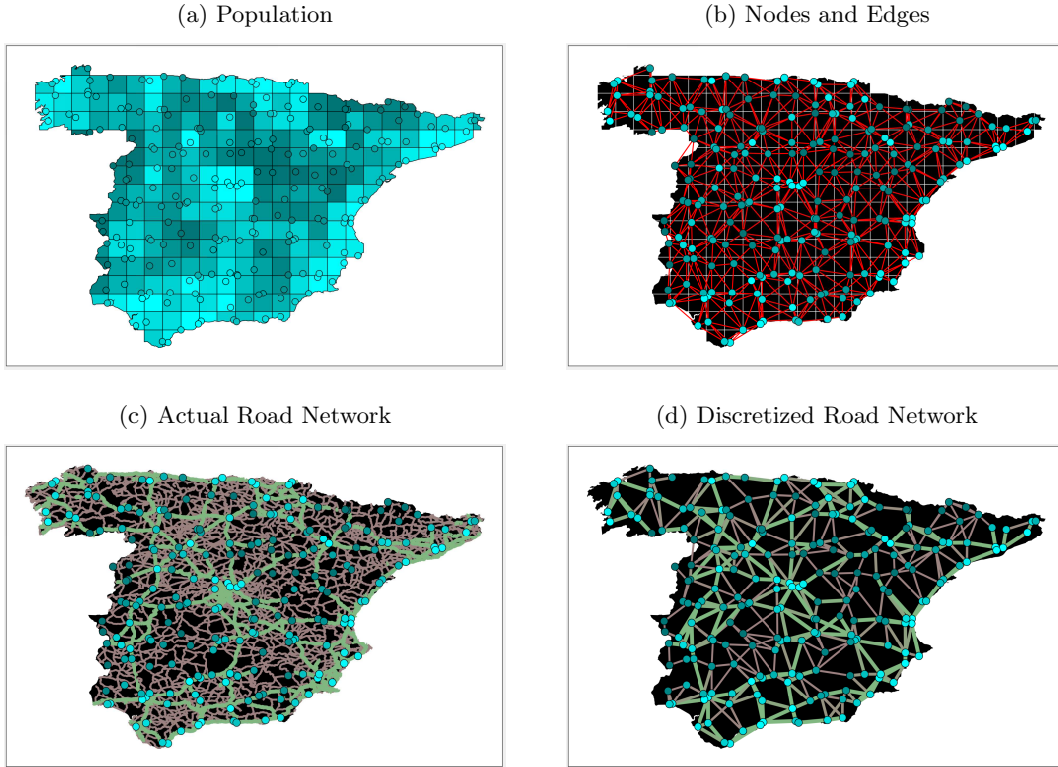


Notes: Panel (a) shows total population from GPW aggregated into 0.5 arc-degree (approximately 50 km) cells. Panel (b) shows the nodes \mathcal{J} corresponding to the population centroids of each cell in Panel (a), reallocated to their closest point on the actual road network, and the edges \mathcal{E} corresponding to all the vertical and diagonal links between cells. Panel (c) shows the centroids and the actual road network. Green segments correspond to national roads, red segments are all other roads, and the width of each segment is proportional to the number of lanes. Panel (d) shows the same centroids and the edges as the baseline graph in Panel (b), where each edge is weighted proportionally to the average number of lanes on the cheapest path between each pair of nodes on the road network. The color shade ranges from red to green according to the fraction of the shortest path traveled on a national road.

of the population distribution across cells. The (b) panels display the cells, the centroids (light blue circles) and the edges (red segments) of the underlying graph. The (c) panels show the centroids and the full road network. Green segments correspond to national roads and red segments correspond to other roads. The width of each road is proportional to its number of lanes.

Finally, the (d) panels show the infrastructure in the discretized road network. Each of the edges from the (b) panels is now assigned a width depending on the average number of lanes, $lanes_{jk}$, and a color ranging from red to green depending on the likelihood of using a national road, nat_{jk} . The width and color scale are the same as in panel (c). When no direct link from j to k is identified by our procedure, no edge is shown. The resulting discretized networks on the baseline grids clearly mirror the actual road networks for both countries, but they are now expressed in terms of the nodes and edges of our model and therefore allow us to quantify it.

Figure 3: Discretization of the Spanish Road Network



Notes: Panel (a) shows total population from GPW aggregated into 0.5 arc-degree (approximately 50 km) cells. Panel (b) shows the nodes \mathcal{J} corresponding to the population centroids of each cell in Panel (a), reallocated to their closest point on the actual road network, and the edges corresponding to all the vertical and diagonal links between cells. Panel (c) shows the centroids and the actual road network. Green segments correspond to national roads, red segments are all other roads, and the width of each segment is proportional to the number of lanes. Panel (d) shows the same centroids and the edges as the baseline graph in Panel (b), where each edge is weighted proportionally to the average number lanes on the cheapest path between each pair of nodes on the road network. The color shade ranges from red to green according to the fraction of the shortest path traveled on a national road.

5.2 Parametrization

We discuss the specific parametric assumptions to implement the general model described in Section 3.

Preferences and Technologies The individual utility over traded and non-traded goods defined in (1) is assumed to be Cobb-Douglas,

$$U = c^\alpha h^{1-\alpha},$$

while the aggregator of traded goods (2) is CES:

$$C_j = \left(\sum_{n=1}^N (C_j^n)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (20)$$

where $\sigma > 0$ is the elasticity of substitution. Labor is the only factor of production and the production technologies (3) are assumed to be linear:

$$Y_j^n = z_j^n L_j^n.$$

We assume $\alpha = 0.4$ to match a standard share of non-traded goods in consumption and $\sigma = 5$ which corresponds to a central value of the demand elasticities reported by [Head and Mayer \(2014\)](#) across estimates from the international trade literature. As we discuss below, the calibrated model gives a reasonable prediction for the distance elasticity of trade, which is closely linked to σ .

Labor Mobility We undertake the country-by-country analysis of misallocation for the cases in which labor is fixed and in which it is perfectly mobile. In this way, we accommodate that internal rates of labor mobility may be different across countries. In the absence of data to discipline this assumption, we opt for reporting the results in these polar cases. For the multi-country application of the last section, we allow in addition for a partial-mobility case where labor is mobile within countries but not across countries.

Transport Technology We adopt the log-linear transport technology (10) with cross-good congestion as described in Section 3.7. We must parametrize the congestion parameter β and the returns to infrastructure parameter γ . Ideally, we would like to set these parameters to elasticities of total trade costs with respect to trade flows and infrastructure. Since such elasticities are not readily available, we narrow the focus to the impact of shipping time on trade costs. Several studies point to a relevant value of time in international shipping.²⁹ Since the majority of inland shipments in the EU are done via road, they likely include goods that are time-sensitive.³⁰

We assume that: i) trade costs are a linear function of shipping time; ii) shipping speed is a log-linear function of the number of vehicles and road lane kilometers; and iii) the number of vehicles is a linear function of the quantity shipped. As shown in Appendix B, under these assumptions we can calibrate β and γ to match the empirical relationship between speed, roads, and vehicles estimated by [Couture et al. \(2018\)](#) in U.S. data. Their estimates imply $\beta = 0.13$ and $\gamma = 0.10$, suggesting

²⁹[Anderson and Van Wincoop \(2004\)](#) calculate 9-percent tax equivalent of the average ocean shipping time cost in the US over the second half of the 20th century. [Hummels and Schaur \(2013\)](#) quantify that one additional day in transit is equivalent to 0.6 to 2.1 percent tariff and [Djankov et al. \(2010\)](#) argue that each additional day of delay is equivalent to a country distancing 70km from its trade partner. [Firth \(2017\)](#) shows that the time delays caused by congestion of railroads impact firm-level outcomes in India, and [Brancaccio et al. \(2019\)](#) argues, using a structurally estimated search model of ships and exporters, that congestion in ports leads to costly delays for exporters.

³⁰75% of the tonne-kilometer shipped via inland transport modes (rail, waterways or road) within the EU-28 are done by road. The share is 50% when considering all transport modes. See https://ec.europa.eu/eurostat/statistics-explained/index.php/Freight_transport_statistics_-_modal_split%20.

decreasing returns to scale.³¹ We use these parameters as a benchmark and also implement the analysis in a case with increasing returns. Specifically, we consider a higher value of γ such that the ratio between γ and β is the mirroring case, $\gamma/\beta = 0.13/0.10$ for $\beta = 0.13$.

Productivities, Endowments, and Geographic Frictions We must impose values for the productivities z_j^n and the endowment of non-traded services H_j . In the case with perfect labor mobility we interpret (L_j^{obs}, GDP_j^{obs}) as outcomes of the planner’s solution for the optimal allocation and optimal flows problems discussed in Section 3.2 taking the observed network I_{jk}^{obs} as given, and use this information to back out the fundamentals (z_j^n, H_j) . In the case with fixed labor, we interpret GDP_j^{obs} as the outcome of the planner solution and use this information to back out the productivities z_j^n , normalizing non-tradeable consumption per capita H_j/L_j^{obs} to 1 and setting the planner’s weights $\omega_j = 1$ everywhere.

Since our data only include aggregate measures of economic activity for each cell, we assume that each location produces only one tradable good. We allow for $N + 1$ different sectors: N differentiated goods, and one homogeneous good. As a benchmark, we assume that each of the differentiated goods is produced in each of the N cells with the largest observed population, and that the homogeneous good is produced by all the remaining cells. We assume 10 different sectors and explore the robustness to alternative values of N . We also implement an alternative calibration where the differentiated products are allocated to the N largest level-2 NUTS regions within each country.

This approach leaves us with J productivity parameters z_j , each corresponding to the productivity of a different location. We choose each location’s productivity (and supply of non-traded goods, when allowing for labor mobility) such that, taking the observed network I_{jk}^{obs} as given, the planner’s solution to the optimal allocation and optimal flows problems from Definition 2 reproduces the observed value added (and population, when allowing for labor mobility).³²

We must also determine the values of the geographic trade frictions δ_{jk}^τ entering in the transport technology (10). We assume that all goods have the same weight ($m_n = 1$) and that frictions depend on distance, $\delta_{jk}^\tau = \delta_0^\tau dist_{jk}$. To calibrate δ_0^τ , we target the level of intra-regional trade in Spain, where regional-level trade data is available. We estimate δ_0^τ jointly with the other fundamentals so that the model matches the 44% share of intra-regional trade in intra-national trade among

³¹Couture et al. (2018) find decreasing returns to scale across their specifications, although the difference between the two parameters is typically small (see Tables 5 and 6 of their paper). They refer to this result as suggesting “modest decreasing returns to scale”. Their preferred estimate, used for our calibration, is column 6 of table 5 of their paper.

³²To compute GDP, we invoke the second welfare theorem from Proposition 4 to recover the prices in the calibrated allocation as the multipliers of the various constraints in the planner’s problem. In the solution of the planner’s problem each location’s value added in tradeable and non-tradeable sectors is $P_j^{n(j)} z_j L_j + P_j^H H_j$, where $n(j)$ denotes the good produced by location j , P_j^n is the price of good n in location j (i.e., multiplier of the flows constraint for good n in j in the planner’s problem), and P_j^H is the price of non-traded services in sector j (i.e., the multiplier of the availability of non-traded goods constraint in the planner’s problem). Value added in the transport sector is not attached to specific nodes and often corresponds to links connecting near empty locations. For accounting purposes, we allocate the national value added in the transport sector proportionally to value added in other sectors, so that regional variation in measured GDP is driven by goods and services.

tradeable sectors across the 15 continental level-2 NUTS regions of Spain from 2001-2005, according to from Spain’s C-Intereg Dataset (Llano et al., 2010). To generate model-based regional trade flows, we aggregate the bilateral node-to-node trade flows in differentiated goods to the bilateral region-to-region level.³³ We then use this value of δ_0^I in the remaining countries when calibrating fundamentals.

Figure A.1 in Appendix B shows the results of the calibration for the convex case of the parameters ($\beta = 0.13, \gamma = 0.10$). Similar relationships hold for the non-convex case ($\beta = 0.13, \gamma = 0.169$). Panels (a) and (b) show the model-implied population and income shares of each location against the data, over all locations in the 24 countries. Except for a few locations, both population and income shares are matched with high precision. The internal trade share for Spain is also precisely matched.³⁴

Panels (c) and (d) show the calibrated fundamentals (productivity and endowment of non-traded good) in the vertical axes against income and population shares in the data, respectively, for the case with labor mobility. Both fundamentals are strongly correlated with observables. A similar positive relationship between productivity and income shares holds in the calibration of the model with fixed labor.

Cost of Building Infrastructure To implement the optimal transport network in counterfactual scenarios, we must parametrize the cost of infrastructure along each edge, δ_{jk}^I . We follow two approaches. In the first approach, we interpret the observed infrastructure I_{jk}^{obs} as the outcome of the planner’s problem, under the assumption that it is equally costly to build in either direction. In this case the observed network, I_{jk}^{obs} , is consistent with the planner’s first-order condition for I_{jk} in (13) under the assumption that $\underline{I}_{jk} = 0$. Imposing symmetry on that first-order condition we then recover the cost of infrastructure as a function of outcomes from the calibrated model (see Appendix A.2). We refer to this measure as the “FOC-based” measure of building costs, $\delta_{jk}^{I,FOC}$.

Our second approach is agnostic about whether the observed network results from any sort of optimization, but takes a stand about how the building costs depend on geographic features. Specifically, we rely on data from Collier et al. (2016), who estimate highway building costs from World Bank infrastructure investment projects across the world, and then relate these costs to a host of geographic and non-geographic frictions.³⁵ We assume that δ_{jk}^I is a function of two geographic features included in their study, distance and ruggedness of the terrain and refer to this building-cost measure as the “geographic” measure, $\delta_{jk}^{I,GEO}$. We interpret an improvement to the connection between any pair of nodes in our counterfactuals as an infrastructure project. In our

³³The model does not make a prediction for bilateral flows of homogeneous goods. Since this good is produced in every region, most of its production is not traded across region boundaries. In the calibration, the intra-regional trade share of this sector is 93%.

³⁴Across the 24 countries the average internal trade share is 38%, with a standard deviation of 12%. We calibrate 0.00156 in the benchmark convex case with mobile labor and 0.00164 in the non-convex case.

³⁵The investment projects in their data are concentrated in low- and middle-income countries, of which three (Lithuania, Georgia, and Macedonia) are in our data. The coefficients from their study introduced in our equation (21) correspond to the average of the coefficients over the distance dummy and the ruggedness index across the 6 specifications in Tables 4 and 5 of their paper.

notation, their estimates imply:

$$\ln \left(\frac{\delta_{jk}^{I,GEO}}{dist_{jk}} \right) = \ln(\delta_0^I) - 0.11 \times (dist_{jk} > 50km) + 0.12 \times \ln(rugged_{jk}), \quad (21)$$

where $dist_{jk}$ is the distance between j and k and $rugged_{jk}$ is the average ruggedness over locations j and k , constructed as detailed in Appendix B. Hence, it is more costly to build on rugged terrain, but less costly per kilometer to build on longer links. We assume that the elasticity of building costs with respect to features of the terrain is the same across all countries.

These steps give two alternative measurements of δ_{jk}^I up to scale in each country. We set $K = 1$ and choose δ_0^I to satisfy the network-building constraint with equality in each country.

5.3 Model-Implied Trade Flows and Congestion

We check the model predictions for bilateral trade flows. We use bilateral trade data across Spanish regions defined at the level-2 NUTS subdivision from Spain’s C-Intereg Dataset. Excluding islands, this gives 15 Spanish regions. Panel (a) of Figure 4 shows observed and model-based trade flows in differentiated products for the calibration where each differentiated good is allocated to a different region. Own-region trade flows are shown as red crosses. The model implied bilateral flows have a correlation of 0.79 with the data.

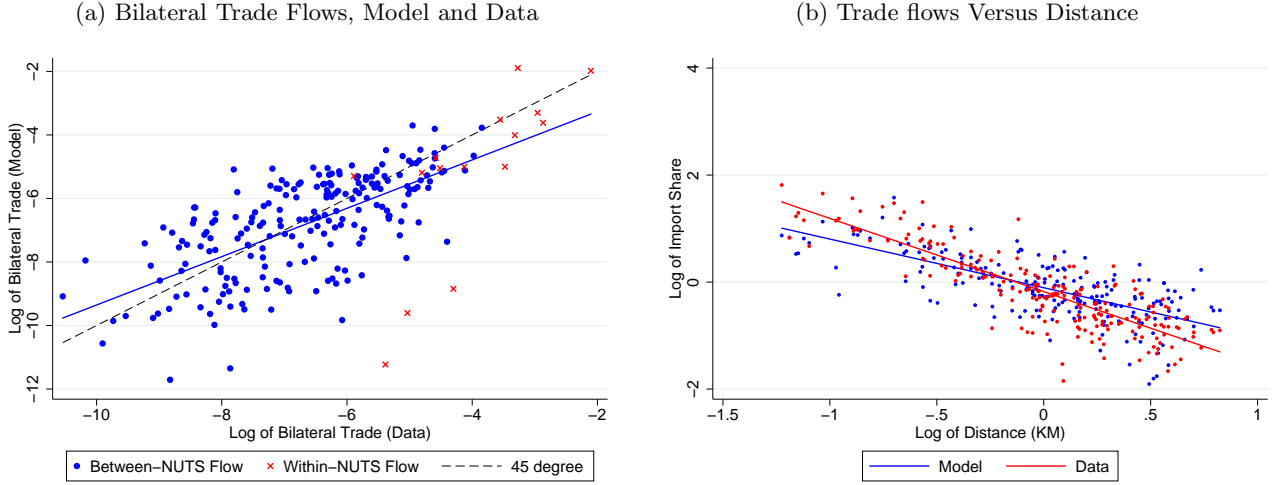
Another way to assess the implied trade flows is to look at the gravity implications. The standard gravity model posits a log-linear relationship between bilateral trade shares and trade costs. The gravity model typically gives a good fit of the international data (Head and Mayer, 2014), and is often applied within countries (Allen and Arkolakis, 2014). A common approach is to parametrize bilateral trade costs as a function of geographic frictions. In our parametrization, bilateral trade costs among locations depend on the distance through the network, but also on the equilibrium levels of congestion, the observed levels of infrastructure, and the calibrated values of β and γ . We can compare the relationship between trade flows and distance implied by the model using the previous aggregation to the level-2 NUTS subdivision in Spain. Panel (b) of Figure 4 shows the bilateral import share among level-2 NUTS and the log of distance in the model and in the data, after controlling for exporter fixed effects.³⁶ A linear regression yields elasticities of -0.91 in the model and -1.37 in the data. Overall, the figures suggest that the model makes reasonable predictions for the distribution of trade flows.³⁷

We examine the congestion taxes needed to implement the allocation. Due to the log-linear

³⁶We run, in both model-generated and observed data, the regression $\ln(\lambda_{jk}^{NUTS}) = \delta \ln(dist_{jk}) + \psi_j + \varepsilon_{jk}$, where λ_{jk} is the import share and $dist_{jk}$ is the bilateral distance between the level-2 NUTS divisions. The figure shows both the import share and distance as residuals from exporter fixed effects. Distance is computed between geographic centroids. We exclude zero flows and flows to the own region.

³⁷The results are similar in calibrations that assign differentiated products to the largest locations in the country, with correlations around 0.8 between model-based non-zero bilateral trade flows and the data. They are also similar in calibrations that assign one good to each region but assume away the homogeneous product. The relationship between trade and distance is very similar for France, suggesting that the key gravity properties are dictated by elasticity parameters rather than by the distribution of fundamentals.

Figure 4: Actual and Predicted Trade



Note: The left panel shows the model-based flows across level-2 NUTS regions of continental Spain against the data in the calibration assigning a differentiated product to the largest city within each level-2 NUTS region. The right panel shows the log of the model-implied and observed import shares against distance as residuals from exporter fixed effects. Linear regression slope (robust SE) is: -0.908 (0.069) in the model and -1.368 (0.058) in the data. Trade flows to the same region appear as red crosses in the first panel.

specification of the transport technology, Proposition 4 implies that the taxes are a fraction $\beta = 13\%$ of the transport costs in every link. Given the total transport costs, we obtain numerically that the mean ad valorem tax across locations is about 0.6% in the cases with and without labor mobility. The total taxes paid represent 0.3% of GDP in both cases.³⁸

5.4 Optimal Expansion and Reallocation

We simulate two types of counterfactuals. First, we compute the aggregate gains from an optimal expansion of the observed road network within each country. We assume that the total resources K are increased by 50% relative to the observed network, constraining the planner to build on top of the existing network, I_{jk}^{obs} . In the notation of restriction (iii) in definitions 1 and 2, this means that $\underline{I}_{jk} = I_{jk}^{obs}$. Second, we compute the losses due to misallocation of current roads within each country. We assume that the total resources K are the same as in the observed network, without constraining the planner to build on top of the existing network, I_{jk}^{obs} . In the notation of restriction (iii) in Definitions 1 and 2, this means that $\underline{I}_{jk} = 0$. We set the upper bound on infrastructure \bar{I}_{jk} to be 50% above the largest level of infrastructure observed in each country.

In short, the first “optimal expansion” counterfactual amounts to optimally expanding the network on top of what is already observed, while the second “optimal reallocation” counterfactual

³⁸In simulations of the calibration for Spain where we randomly increase capacity in random links, we find an elasticity of quantity flows Q_{jk} to infrastructure I_{jk} of 0.503 (0.0202) in the benchmark calibration ($\gamma < \beta$), 0.799 (0.0330) in a calibration that imposes $\gamma = \beta$, and 1.296 (0.0395) in the non-convex calibration ($\gamma > \beta$). Duranton and Turner (2011) report IV estimates of the relationship between total vehicles-kilometers traveled and road capacity across U.S. cities between 0.68 and 1.33 (in their Table 6), with many of these estimates close to 1.

amounts to optimally reallocating the existing roads. The first counterfactual is more policy-relevant, as it prescribes where new roads should be built and yields the aggregate gains of those investments. The second counterfactual is unfeasible in reality, but gives a sense of the losses from misallocation of existing roads.

We implement the optimal expansion under the two measures of building costs, the FOC-based measure $\delta_{jk}^{I,FOC}$ and the geographic measure $\delta_{jk}^{I,GEO}$. The optimal reallocation is only meaningful under the geographic measure, since the observed network is optimal by construction under $\delta_{jk}^{I,FOC}$. We implement each of these counterfactuals for each of the two values of γ , assuming both fixed and mobile labor, separately for each of the 24 countries. We re-calibrate the model for each value of γ , assumption on labor mobility, and country.

Regional Impact Within Countries We inspect first the within-country regional implications for two large countries in our data, Spain and France. Figure 5 depicts the pattern of investment and population change under the geographic measure of building costs $\delta^{I,GEO}$. Panels (a) and (b) show the optimal expansion and panels (c) and (d) show the optimal reallocation under labor mobility. Panels (e) and (f) reproduce the optimal reallocation assuming that labor is not mobile. The thickness of each link increases with the absolute value of the investment, defined as the difference between the counterfactual and the observed infrastructure, $I_{jk}^* - I_{jk}^{obs}$. In the reallocation counterfactual, links with negative investment, $I_{jk}^* - I_{jk}^{obs} < 0$, are shown in red, while all other links are shown in green. In turn, with labor mobility, green nodes denote positive population change, and red nodes denote negative population change. Brighter nodes represent a larger absolute value of population change.

In the optimal reallocation counterfactual, we observe positive investments radiating away from some areas with higher economic activity in the case of France, but a more dispersed investment pattern in Spain. As we compare panels (a) and (b) with panels (c) and (d), we recognize similar investment patterns in the optimal reallocation and expansion counterfactuals within each country: the links identified as having too much infrastructure, shown in red in panels (c) and (d), typically feature no expansion in panels (a) and (b). The comparison between panels (c)-(d) and panels (e)-(f) reveals that allowing labor mobility does not fundamentally affect the optimal infrastructure investments.


In the cases of optimal expansion and optimal reallocation, the population is reallocated to the same set of regions within each country. Due to the labor mobility constraint in the planner's problem, changes in labor are perfectly correlated with changes in consumption of traded commodities per worker, c_j .³⁹ For the cases without labor mobility, there is a similar consistency across the counterfactuals in the changes in consumption of traded commodities per capita c_j across locations.

We inspect, across the 24 countries, how a few typically observable regional outcomes map to infrastructure investment. Panel (a) of Table 1 reports results from regressions of infrastructure


³⁹The labor mobility constraint (vi) from Definition 2 implies $\alpha \Delta \ln c_j = (1 - \alpha) \Delta \ln L_j + \Delta \ln u$, where $\Delta \ln x$ denotes the difference in the log of variable x between the counterfactual and calibrated allocations.

Figure 5: Optimal Network Reallocation and Expansion: Spain and France


(a) Optimal Network Expansion with Labor Mobility, France




(b) Optimal Network Expansion with Labor Mobility, Spain




(c) Optimal Network Reallocation with Labor Mobility, France




(d) Optimal Network Reallocation with Labor Mobility, Spain



(e) Optimal Network Reallocation with Fixed Labor, France



(f) Optimal Network Reallocation with Fixed Labor, Spain



Notes: All counterfactuals use the geographic measure of building cost, $\delta^{I,GEO}$ and the benchmark parametrization of β and γ . Graph nodes aggregate data from 0.5 arc-degree (approximately 50 km) cells. The width and brightness of each link is proportional to the difference between the optimal counterfactual network and the observed network, $I_{jk}^* - I_{jk}^{obs}$, for each link $jk \in \mathcal{E}$ shown in panel (b) of Figures 2 and 3. The color scale is the same as in Figure 2. Red links represent negative investment. With labor mobility, brighter green (red) nodes represent larger population increase (decrease).

growth on each location's initial population and tradeable income per capita. We report here results corresponding to the case with labor mobility and the benchmark parametrization of γ and

β , but the qualitative features we discuss are similar under alternative specifications. Optimal road investments are directed to locations with initially lower levels of infrastructure, reflecting decreasing returns to infrastructure at the link level. The investments are also more intensely directed to locations with initially higher levels of population and income per worker. Since the model implies a complex mapping from the fundamentals to the investments, these observable outcomes guide only a fraction of the optimal investment decisions (R^2 in the order of 24-39% under geographic measure of trade costs and only 4% under the FOC measure).

Table 1: Optimal Infrastructure Investment, Population Growth and Local Characteristics

Panel A: Dependent Variable: Infrastructure Growth			
	(1)	(2)	(3)
	Reallocation	Expansion (GEO)	Expansion (FOC)
Population	0.343***	0.125***	0.002
Tradeable Income per Capita	0.151	0.071	0.007*
Infrastructure	-0.418***	-0.235***	-0.010
Observations	868	868	868
Adjusted R-squared	0.29	0.24	0.04
Panel B: Dependent Variable: Population Growth			
	(1)	(2)	(3)
	Reallocation	Expansion (GEO)	Expansion (FOC)
Population	-0.001	-0.000	-0.000
Tradeable Income per Capita	0.008	0.008	0.001
Consumption per Capita	-0.061***	-0.060***	-0.008***
Infrastructure	-0.001	-0.001	-0.000
Infrastructure Growth	0.002**	0.002*	-0.001
Differentiated Producer	0.010***	0.010***	0.002***
Observations	868	868	868
Adjusted R-squared	0.53	0.54	0.80

Each column corresponds to a different regression pooling all locations across the 24 countries in the benchmark parametrization of γ and β , assuming mobile labor and $N=10$. All regressions include country fixed effects. Standard errors are clustered at the country level. ***=1% significance, **=5%, *=10%. Dependent variables: population growth is defined as $\Delta \ln L_j$, where $\Delta x = x' - x$ denotes the difference between variable x in the counterfactual (x') and in the calibrated allocation (x). Investment growth is defined as the difference over the average, $\Delta I_j / (\frac{1}{2} (I_j + I'_j))$, where total infrastructure at the node level is defined as $I_j = \sum_{k \in \mathcal{N}(j)} I_{jk}$. Independent variables correspond to the log of the level of each variable in the calibrated model. Population and income per capita are the two outcomes matched by the calibration. Consumption per capita corresponds to traded goods c_j in the calibrated model. Differentiated producer is a dummy for whether the location is a producer of differentiated goods in the calibration.

In Panel (b), the dependent variable is population growth. To understand the patterns of optimal reallocation of population, in addition to the variables from the previous regression we also include infrastructure growth, consumption per capita, and a dummy for whether a location is a differentiated producer. The handful of variables in the regression explain between 50% and 80%

of the population changes. Infrastructure growth in a location has a positive impact on population in the counterfactuals that use the geographic measure of trade costs. The magnitude of the effect and its explanatory power on the distribution of population changes is small, reflecting that growth in a location depends on investments in other locations.

Consumption of traded goods per capita is a strong determinant, with a negative elasticity of population growth with respect to initial consumption in the order of 1%-7%. If consumption per capita was excluded, then the coefficient on income per capita would become negative and significant, with a negative elasticity of growth with respect to income per capita of 1% across the three counterfactuals. Hence, the impact of initial income on population growth in the optimal investment plan operates through the level of consumption.

This reallocation pattern reflects that the goal of the optimal investments is to reduce variation in the marginal utility of consumption of traded commodities across locations. Since changes in population and consumption per capita between the counterfactual and initial allocation are perfectly correlated, the optimal investment plan leads to an increase in consumption of traded commodities in locations where consumption per capita is initially low. We conclude that the optimal investment in infrastructure reduces spatial inequalities, although different assumptions on building costs imply different ways of achieving this goal by changing the optimal placement of infrastructure, as implied by our previous discussion.

Aggregate Impact Across Countries We now show the aggregate welfare effects. Table 2 shows the average welfare gain for each counterfactual across the 24 countries in our data. Tables A.2 and A.3 in Appendix B show the results for each country with fixed and mobile labor, respectively. In the benchmark parametrization of γ and β , using the geographic measure of building costs we find average welfare gains across countries of 1.7% – 1.8%. The effects are much smaller under the FOC-based measure because in that case the optimal expansion does not address a suboptimal placement of existing roads. The average gains are increasing in the returns to scale γ , with the average welfare gains increasing to between 2.4% and 2.9% under geographic measure of building costs. These effects vary considerably across countries, ranging from around 0.1% to 8%. There is no clear relationship between misallocation and country size or income. Some Eastern European countries such as Georgia, Lithuania, and Latvia appear with relatively high misallocation in the benchmark case (3.0% to 3.6% relative to a mean of 2.4%), and so do Denmark (7.8%), France (3.4%) and Spain (4.7%). Belgium, Luxembourg and Macedonia appear as the least misallocated countries.

This distribution of welfare gains across countries is in general stable regardless of the parametrization of γ , the assumption on labor mobility, the parametrization of the building costs δ^I , or the type of counterfactual. For example, across the parametrizations of γ and labor mobility, the correlation between the gains from optimally expanding the network under the two measures of building costs, $\delta^{I,GEO}$ and $\delta^{I,FOC}$, is between 0.76 and 0.92. Therefore, the answers to the questions of which countries would gain more from optimally expanding their current road networks and which

Table 2: Average Welfare Gains Across Countries

Returns to Scale:	Benchmark		Non-Convex	
Labor:	Fixed	Mobile	Fixed	Mobile
Optimal Reallocation				
$\delta = \delta^{I,GEO}$	1.7%	1.8%	2.4%	2.5%
Optimal Expansion				
$\delta = \delta^{I,GEO}$	1.7%	1.8%	2.8%	2.9%
$\delta = \delta^{I,FOC}$	0.3%	0.3%	0.9%	1.3%

Each element of the table shows the average welfare gain in the corresponding counterfactual across the 24 countries.

countries suffer larger losses from misallocation of current roads is robust across these cases.

Alternative Assumptions The analysis was implemented assuming $N = 10$ sectors. We also implement the calibration and counterfactuals assuming different numbers of sectors. Table A.4 in Appendix B reports the coefficients from column (2) of Table 1 corresponding to the optimal expansion under calibrations that assume $N = 5$ or $N = 15$. In these alternative cases, the patterns described above remain unchanged, and the magnitude of most of the coefficients does not exhibit large variation.

Similarly, Table A.5 reproduces Table 2 for the benchmark and for $N = 15$. The aggregate gains change little with the number of sectors. The correlation between the aggregate welfare effects across countries under $N = 15$ and under $N = 10$ is above 0.9 for each possible type of counterfactual and assumptions on labor mobility and value of γ . The table also reports average welfare effects under an alternative calibration where each of the largest N regions in each country, defined as level-2 NUTS political subdivisions, is assigned a differentiated product. The correlation in welfare gains across countries between the benchmark case and this alternative allocation is above 0.8 across assumptions of labor mobility, number of goods, and type of counterfactual.

Finally, Table A.6 replicates the benchmark case under the assumption of no congestion across goods. We find very similar average welfare effects in the two cases.

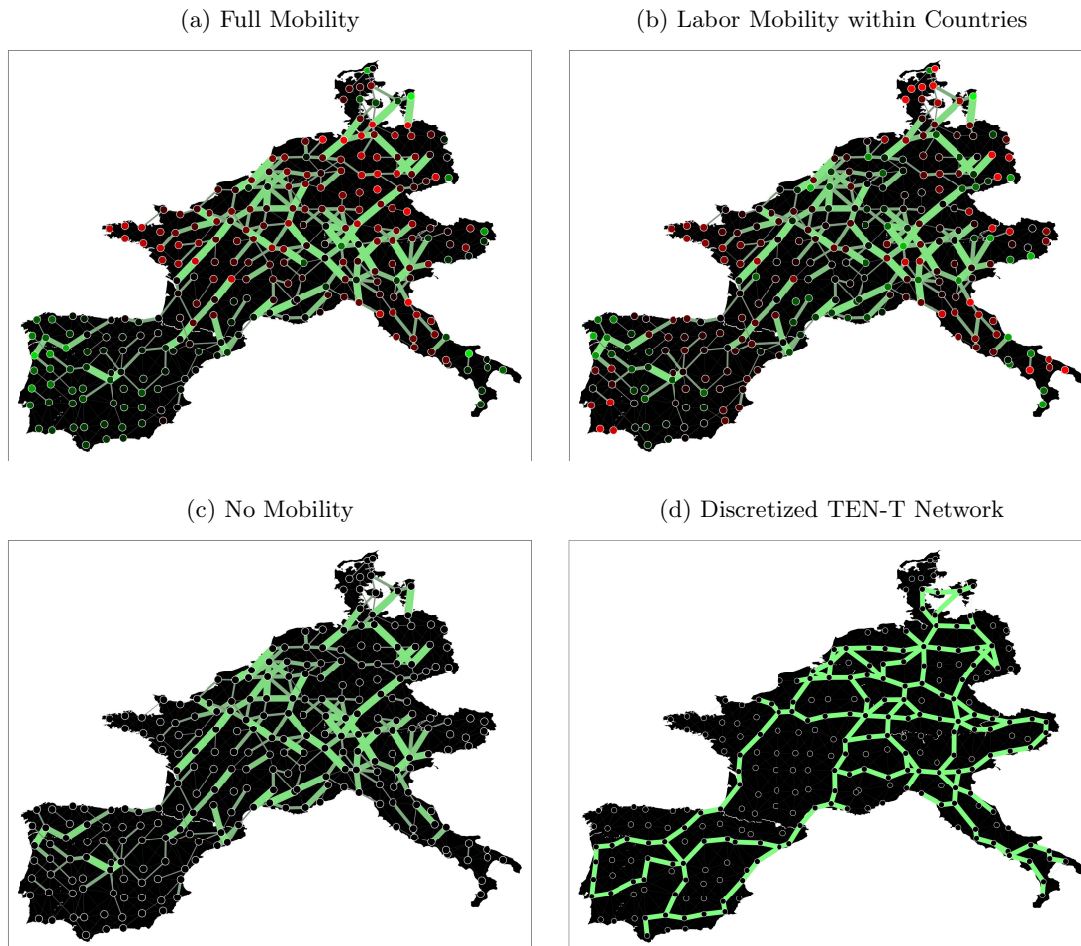
5.5 Application to Multiple Countries within Europe

Our previous applications considered each country in isolation. We now implement the analysis for a region of western Europe.⁴⁰ Appendix Figure A.2 shows the baseline map and the discretized network for this connected set of countries. We assume that each country produces a country-specific differentiated product, in addition to a homogeneous good and use the same parameters as in the benchmark. We re-calibrate the fundamentals assuming that the 5 largest locations in terms of observed population within each country produce the differentiated product of that country,

⁴⁰We include 11 countries: Austria, Belgium, Switzerland, Germany, Denmark, Spain, France, Italy, Luxembourg, Netherlands, and Portugal. We use 1 degree by 1 degree cells, resulting in 261 cells.

while the remaining locations produce the homogeneous product. We now also implement a case with partial mobility where labor is mobile within countries but not across countries, recalibrating the fundamentals each time.

Figure 6: Optimal Network Expansion: Europe



Notes: All counterfactuals use the geographic measure of building costs, $\delta^{I,GEO}$ and the benchmark parametrization of β and γ . The width and brightness of each link is proportional to the difference between the optimal counterfactual network and the observed network, $I_{jk}^* - I_{jk}^{obs}$, for each link $jk \in \mathcal{E}$ shown in panel (b) of Figure A.2. The color scale is the same as in Figure 2. Red links represent negative investment. With labor mobility, brighter green (red) nodes represent larger population increase (decrease). Panel (d) shows a discretized version of the TEN-T Core Network Corridors of the Trans-European Transport Network, based on information available at <http://ec.europa.eu/transport/infrastructure/tentec/tentec-portal/site/en/maps.html>.

Figure 6 shows the optimal network expansion under different assumptions of labor mobility. The counterfactuals highlight the areas where European investments would be more profitable. The investments are concentrated in Benelux countries, France, Germany, and Northern Italy. Within Spain, the optimal expansion looks quite different from what we found in panel (b) of Figure 5, reflecting the European planner's incentives to deal with international trade. The European planner prioritizes two corridors connecting Spain to the north of Portugal and the center of Spain

to France, whereas the Spain-level planner chose a higher density of investment in the south of the country. Within France, the pattern of investments radiating from Paris is similar to the case in (a) of Figure 5, but now we see optimal investments in the connection with Spain, as well as investments in the northeast to connect with neighboring countries.

The optimal network expansion is very similar in the three cases of labor mobility and the welfare gains are close to 2.5% in the three cases. When labor is mobile across Europe, the optimal network investment reallocates workers to southern Spain and Portugal, much like in the country-by-country analysis we found reallocation to areas with relatively lower income per worker. As in the previous cases, these changes in population do not correlate very strongly with the investments.

To conclude, we ask whether these patterns are approximately comparable with the Trans-European Transport Network (TEN-T), a European Commission policy that supports the development of Europe-wide transport networks. The network includes roads, air and inland waterways. The TEN-T defines a core network of “strategic importance” for future investments based on criteria such as eliminating bottlenecks or following suggestions from member states.⁴¹ Panel (d) of Figure 6 shows these corridors for the area of Europe covered by our counterfactual. Broadly speaking, our planning problem identifies some priorities for investment which appear to be similar to what real-world planners have decided, such as the high density of investment in Benelux countries and Germany; the international corridor from Paris to the southwest of France, north of Spain, and Portugal; and the connection between Germany and Denmark. However, we also see some differences, as the solution to our planning problem does not identify the need to invest in roads connecting the southeast of France to the south of Spain and Portugal.

6 Conclusion

In this paper, we develop a framework to study optimal transport networks in spatial equilibrium models. The framework combines a neoclassical environment where each location is a node in a graph, an optimal transport problem subject to congestion in shipping across commodities, and an optimal network design. It nests standard neoclassical trade models and it allows for either fixed or mobile factors across space. We provide conditions such that the full planner’s problem, involving the optimal flow of goods as well as the general-equilibrium and network-design problems, is globally convex and numerically tractable using standard numerical methods typically applied to tackle optimal transport problems.

In the application, we match the model to data on road networks and economic activity across European countries. Using the calibrated model, we compute the gains from road expansion and losses from misallocation. Across countries, we find real consumption losses in the order of 2% associated with misallocation of roads.

⁴¹See https://ec.europa.eu/transport/themes/infrastructure/ten-t-guidelines/maps_en. The planning guidelines are mentioned in the European Commission working document, available in <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52013SC0542&from=EN>.

Our approach using a global planner is particularly well suited for environments where agglomeration spillovers may not be too strong. We have also shown how in principle the model could be used in cases with spillovers. Due to current limitations in computing power and given the level of geographic detail that we handle, we have only applied the model to cases with a limited number of commodities.

We expect the framework to serve as basis for future work. It could be used to study political-economy issues associated with infrastructure, such as spatial competition among planning authorities. We have refrained from identifying sources of misallocation, but it would be interesting to know the role of regional characteristics such as institutional quality. Our application was limited to European countries, but low-income economies are likely to benefit more from infrastructure investment and are perhaps more prone to inefficient investments due to their institutional environments.

The persistence of transport networks also raises interesting issues. The model could be extended to study inefficient network lock-in due to past investments corresponding to dated economic fundamentals. We have also abstracted from decision-making under uncertainty, but it would be interesting to study a planner who decides in anticipation of changing conditions about technology or fundamentals. In the spirit of Barjamovic et al. (2019), who use the prediction of gravity models to infer the location of cities, the model could be used to infer the location of roads in historical data. The framework might also be used to construct instruments for investments in transport infrastructure.

Finally, a number of forces such as commuting or dynamic adjustment were left out of our analysis. We believe these are all interesting avenues to pursue in future research.

References

- Ahuja, R. K., T. L. Magnanti, and J. B. Orlin (1989). Chapter iv network flows. *Handbooks in operations research and management science 1*, 211–369.
- Alder, S. (2019). Chinese roads in india: The effect of transport infrastructure on economic development. *Manuscript, Univ. North Carolina, Chapel Hill*.
- Allen, T. and C. Arkolakis (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics 1085*, 1139.
- Allen, T. and C. Arkolakis (2019). The welfare effects of transportation infrastructure improvements. *Manuscript, Dartmouth and Yale*.
- Allen, T., C. Arkolakis, and Y. Takahashi (2019). Universal gravity. *Journal of Political Economy*.
- Anderson, J. E. and E. Van Wincoop (2003). Gravity with gravitas: a solution to the border puzzle. *American Economic Review 93*(1), 170–192.
- Anderson, J. E. and E. Van Wincoop (2004). Trade costs. *Journal of Economic literature 42*(3), 691–751.
- Arrow, K. J. and A. C. Enthoven (1961). Quasi-concave programming. *Econometrica: Journal of the Econometric Society*, 779–800.
- Asturias, J., M. García-Santana, and R. Ramos Magdaleno (2016). Competition and the welfare gains from transportation infrastructure: Evidence from the golden quadrilateral of india.

- Atkeson, A. and A. T. Burstein (2010). Innovation, firm dynamics, and international trade. *Journal of political economy* 118(3), 433–484.
- Atkin, D. and D. Donaldson (2015). Who’s getting globalized? the size and implications of intra-national trade costs. Technical report, National Bureau of Economic Research.
- Banavar, J. R., F. Colaiori, A. Flammini, A. Maritan, and A. Rinaldo (2000). Topology of the fittest transportation network. *Physical Review Letters* 84(20), 4745.
- Barjamovic, G., T. Chaney, K. A. Coşar, and A. Hortaçsu (2019). Trade, merchants, and the lost cities of the bronze age. *The Quarterly Journal of Economics* 134(3), 1455–1503.
- Bartelme, D. (2015). Trade costs and economic geography: Evidence from the U.S. Technical report, University of Michigan.
- Baum-Snow, N. (2007). Did highways cause suburbanization? *The Quarterly Journal of Economics*, 775–805.
- Beckmann, M. (1952). A continuous model of transportation. *Econometrica: Journal of the Econometric Society*, 643–660.
- Bernot, M., V. Caselles, and J.-M. Morel (2009). *Optimal transportation networks: models and theory*, Volume 1955. Springer Science & Business Media.
- Bertsekas, D. P. (1998). *Network optimization: continuous and discrete models*. Citeseer.
- Boyd, S. and L. Vandenberghe (2004). *Convex optimization*. Cambridge university press.
- Brancaccio, G., M. Kalouptsi, and T. Papageorgiou (2019). Geography, transportation, and endogenous trade costs. *Econometrica*.
- Brandt, L., T. Tombe, and X. Zhu (2013). Factor market distortions across time, space and sectors in China. *Review of Economic Dynamics* 16(1), 39–58.
- Burstein, A. and J. Cravino (2015). Measured aggregate gains from international trade. *American Economic Journal: Macroeconomics* 7(2), 181–218.
- Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2017). The impact of regional and sectoral productivity changes on the U.S. economy. *The Review of Economic Studies* 85(4), 2024–2096.
- Carlier, G. (2010). Optimal transportation and economic applications. Technical report, IMA.
- Chandra, A. and E. Thompson (2000). Does public infrastructure affect economic activity?: Evidence from the rural interstate highway system. *Regional Science and Urban Economics* 30(4), 457–490.
- Chaney, T. (2014a). The network structure of international trade. *The American Economic Review* 104(11), 3600–3634.
- Chaney, T. (2014b). Networks in international trade. In *The Oxford Handbook of the Economics of Networks*.
- Collier, P., M. Kirchberger, and M. Söderbom (2016). The cost of road infrastructure in low-and middle-income countries. *The World Bank Economic Review* 30(3), 522–548.
- Combes, P.-P. and M. Lafourcade (2005). Transport costs: measures, determinants, and regional policy implications for france. *Journal of Economic Geography* 5(3), 319–349.
- Coşar, A. K. and B. Demir (2016). Domestic road infrastructure and international trade: Evidence from turkey. *Journal of Development Economics* 118, 232–244.
- Costinot, A. and A. Rodríguez-Clare (2013). Trade theory with numbers: Quantifying the consequences of globalization, vol. 4 of handbook of international economics.
- Couture, V., G. Duranton, and M. A. Turner (2018). Speed. *Review of Economics and Statistics* 100(4), 725–739.

- Desmet, K. and E. Rossi-Hansberg (2013). Urban accounting and welfare. *American Economic Review* 103(6), 2296–2327.
- Djankov, S., C. Freund, and C. S. Pham (2010). Trading on time. *The Review of Economics and Statistics* 92(1), 166–173.
- Doll, C., H. van Essen, et al. (2008). Road infrastructure cost and revenue in europe. *Produced within the study Internalisation Measures and Policies for all external cost of Transport (IMPACT) P Deliverable 2*.
- Donaldson, D. (2015). The gains from market integration. *Annual Review of Economics* 7(1), 619–647.
- Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review* 108(4-5), 899–934.
- Donaldson, D. and R. Hornbeck (2016). Railroads and american economic growth: A market access approach. *The Quarterly Journal of Economics*, qjw002.
- Duranton, G., P. M. Morrow, and M. A. Turner (2014). Roads and trade: Evidence from the us. *The Review of Economic Studies* 81(2), 681–724.
- Duranton, G. and M. A. Turner (2011). The fundamental law of road congestion: Evidence from us cities. *The American Economic Review* 101(6), 2616–2652.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Ekeland, I. (2010). Notes on optimal transportation. *Economic Theory* 42(2), 437–459.
- Faber, B. (2014). Trade integration, market size, and industrialization: evidence from china’s national trunk highway system. *The Review of Economic Studies*, rdu010.
- Fajgelbaum, P. D., E. Morales, J. C. Suárez Serrato, and O. Zidar (2018). State taxes and spatial misallocation. *The Review of Economic Studies* 86(1), 333–376.
- Felbermayr, G. J. and A. Tarasov (2015). Trade and the spatial distribution of transport infrastructure.
- Fernald, J. G. (1999). Roads to prosperity? assessing the link between public capital and productivity. *American Economic Review*, 619–638.
- Figuroa, C., B. Fotsch, S. M. Hubbard, and J. Haddock (2013). Assessment procedures for paved and gravel roads. Technical report, Purdue University School of Civil Engineering.
- Firth, J. (2017). I’ve been waiting on the railroad: The effects of congestion on firm production.
- Galichon, A. (2016). *Optimal Transport Methods in Economics*. Princeton University Press.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. *Handbook of International Economics, Vol. 4*.
- Helpman, E. (1998). The size of regions: transport and housing as factors in agglomeration. In D. Pines, E. Sadka, and I. Zilcha (Eds.), *Topics in Public Economics*, pp. 33–54. Cambridge University Press Cambridge.
- Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing TFP in China and India. *The Quarterly Journal of Economics* 124(4), 1403–1448.
- Hsieh, C.-T. and E. Moretti (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics* 11(2), 1–39.
- Hummels, D. L. and G. Schaur (2013). Time as a trade barrier. *American Economic Review* 103(7), 2935–59.
- Kantorovich, L. V. (1942). On the translocation of masses. In *Dokl. Akad. Nauk SSSR*, Volume 37, pp. 199–201.
- Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy* 99(3), 483–499.

- Lai, E. L.-C., H. Fan, and H. S. Qi (2015). Global gains from reduction of trade costs.
- Limao, N. and A. J. Venables (2001). Infrastructure, geographical disadvantage, transport costs, and trade. *The World Bank Economic Review* 15(3), 451–479.
- Llano, C., A. Esteban, J. Pérez, and A. Pulido (2010). Opening the interregional trade black box: The c-interreg database for the spanish economy (1995–2005). *International Regional Science Review*.
- Maibach, M., C. Schreyer, D. Sutter, H. Van Essen, B. Boon, R. Smokers, A. Schrotten, C. Doll, B. Pawlowska, and M. Bak (2013). Handbook on estimation of external costs in the transport sector.
- Martincus, C. V., J. Carballo, and A. Cusolito (2017). Roads, exports and employment: Evidence from a developing country. *Journal of Development Economics* 125, 21–39.
- Mohring, H. and M. Harwitz (1962). Highway benefits: An analytical framework. Technical report, Northwestern University.
- Monge, G. (1781). *Mémoire sur la théorie des déblais et des remblais*. De l’Imprimerie Royale.
- Nagy, D. (2016). City location and economic development. *Manuscript, CREI*.
- Nesterov, Y. and A. Nemirovskii (1994). *Interior-point polynomial algorithms in convex programming*. SIAM.
- Newbery, D. M. (1988). Road damage externalities and road user charges. *Econometrica: Journal of the Econometric Society*, 295–316.
- Ramondo, N., A. Rodríguez-Clare, and M. Saborío-Rodríguez (2016). Trade, domestic frictions, and scale effects. *American Economic Review* 106(10), 3159–3184.
- Redding, S. J. (2016). Goods trade, factor mobility and welfare. *Journal of International Economics* 101, 148–167.
- Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics* 9, 21–58.
- Redding, S. J. and M. A. Turner (2015). Transportation costs and the spatial organization of economic activity, vol. 5 of handbook of regional and urban economics.
- Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics* 11(4), 707–720.
- Riley, S., S. DeGloria, and R. Elliot (1999). A terrain ruggedness index that quantifies topographic heterogeneity. *Intermountain Journal of Sciences* 5(1–4), 23–27.
- Roback, J. (1982). Wages, rents, and the quality of life. *The Journal of Political Economy*, 1257–1278.
- Santambrogio, F. (2015). Optimal transport for applied mathematicians. *Birkhäuser, NY*.
- Sotelo, S. (2016). Domestic trade frictions and agriculture. *Manuscript, University of Michigan*.
- Swisher IV, S. (2015). Reassessing railroads and growth: Accounting for transport network endogeneity. Technical report, mimeo.
- Takayama, A. (1985). *Mathematical economics*. Cambridge University Press.
- Tay, R. and A. Churchill (2007). Effect of different median barriers on traffic speed. *Canadian Journal of Transportation* 1(1).
- Trew, A. W. (2016). Endogenous infrastructure development and spatial takeoff.
- Villani, C. (2003). *Topics in optimal transportation*. Number 58. American Mathematical Soc.
- Winston, C. (1985). Conceptual developments in the economics of transportation: an interpretive survey. *Journal of Economic Literature* 23(1), 57–94.

Online Appendix

A Appendix to Section 3 (Model)

A.1 Planner's Problem

In this section we present the first-order conditions to the planner's problem. We refer to these conditions in some of the characterizations in the text and in the proofs below. We present the problem adopting a formulation of the transport technology that nests the approach with own-good congestion in which the transport cost is denominated in units of the good being shipped, as well as the approach with congestion across goods in which the transport cost is denominated in units of the bundle of traded goods (discussed in Section 3.7). In the formulations below, the parameter $\chi \in \{0, 1\}$ takes a value of 0 in the case with own-good congestion and a value of 1 with cross-good congestion. The case $\chi = 0$ corresponds to the equations presented in the body of the paper.

Immobile Labor

The Lagrangian of the problem in Definition 1 is

$$\begin{aligned} \mathcal{L} = & \sum_j \omega_j L_j U(c_j, h_j) - \sum_j P_j^D \left[c_j L_j + \chi \sum_{k \in \mathcal{N}(j)} \tau_{jk}(Q_{jk}, I_{jk}) Q_{jk} - D_j(D_j^1, \dots, D_j^N) \right] - \sum_j P_j^H (h_j L_j - H_j) \\ & - \sum_j \sum_n P_j^n \left[D_j^n + \sum_{k \in \mathcal{N}(j)} (Q_{jk}^n + (1 - \chi) \tau_{jk}^n(Q_{jk}^n, I_{jk}) Q_{jk}^n) - F_j^n(L_j^n, \mathbf{v}_j^n, \mathbf{X}_j^n) - \sum_{i \in \mathcal{N}(j)} Q_{ij}^n \right] \\ & - \sum_j W_j \left[\sum_n L_j^n - L_j \right] - \sum_j \sum_l R_j^l \left[\sum_n V_j^{ln} - V_j^l \right] \\ & - P_K \left(\sum_j \sum_{k \in \mathcal{N}(j)} \delta_{jk}^I I_{jk} - K \right) + \sum_{j,k} \zeta_{jk}^I (I_{jk} - \bar{I}_{jk}) + \sum_{j,k} \zeta_{jk}^{\bar{I}} (\bar{I}_{jk} - I_{jk}) \\ & + \sum_{j,k,n} \zeta_{jkn}^Q Q_{jk}^n + \sum_{j,n} \zeta_{jn}^L L_j^n + \sum_{j,n,l} \zeta_{jnl}^V V_j^{ln} + \sum_{j,n,l} \zeta_{jnl}^X X_j^{ln} + \sum_{j,n} \zeta_{jn}^D D_j^n + \sum_j \zeta_j^c c_j + \sum_j \zeta_j^h h_j \end{aligned}$$

where $Q_{jk} = \sum_{n=1}^N m^n Q_{jk}^n$ and where $P_j^D, P_j^H, P_j^N, W_j, R_j^l, P_K, \zeta_{jk}^I, \zeta_{jkn}^Q, \zeta_{jn}^L, \zeta_{jnl}^V, \zeta_{jn}^C, \zeta_j^c, \zeta_j^h$ are the multipliers of all constraints implied by (i)-(v) in Definition 1. The first-order conditions with respect to consumption and production are:

$$\begin{aligned} [c_j] \quad & \omega_j L_j U_C(c_j, h_j) + \zeta_j^c = P_j^D L_j \\ [h_j] \quad & \omega_j L_j U_H(c_j, h_j) + \zeta_j^h = P_j^H L_j \\ [D_j^n] \quad & P_j^D \frac{\partial D_j}{\partial D_j^n} + \zeta_{jn}^D = P_j^n \\ [L_j^n] \quad & P_j^n \frac{\partial Y_j^n}{\partial L_j^n} + \zeta_{jn}^L = W_j \\ [V_j^n] \quad & P_j^n \frac{\partial Y_j^n}{\partial V_j^{ln}} + \zeta_{jnl}^V = R_j^l \\ [X_j^n] \quad & P_j^n \frac{\partial Y_j^n}{\partial X_j^{ln}} + \zeta_{jnl}^X = P_j^l \end{aligned}$$

The first-order condition with respect to flows is:

$$\begin{aligned} [Q_{jk}^n] \quad & -\chi P_j^D \left(\tau_{jk}^n(Q_{jk}, I_{jk}) + \frac{\partial \tau_{jk}^n(Q_{jk}, I_{jk})}{\partial Q_{jk}} Q_{jk} \right) \\ & - (1 - \chi) P_j^n \left(\tau_{jk}^n(Q_{jk}^n, I_{jk}) + \frac{\partial \tau_{jk}^n(Q_{jk}^n, I_{jk})}{\partial Q_{jk}^n} Q_{jk}^n \right) + P_k^n - P_j^n + \zeta_{jkn}^Q = 0 \end{aligned} \quad (\text{A.1})$$

which, along with the complementary slackness condition for Q_{jk}^n , implies (8) in the main text.

Finally, the first order condition with respect to the network investment is

$$[I_{jk}] \quad \chi \sum_{n=1}^N P_j^D Q_{jk}^n \left(-\frac{\partial \tau_{jk}^n(Q_{jk}^n, I_{jk})}{\partial I_{jk}} \right) + (1 - \chi) \sum_{n=1}^N P_j^n Q_{jk}^n \left(-\frac{\partial \tau_{jk}^n(Q_{jk}^n, I_{jk})}{\partial I_{jk}} \right) + (\zeta_{jk}^L - \zeta_{jk}^{\bar{I}}) = P_K \delta_{jk}^I \quad (\text{A.2})$$

which, along with the complementary slackness condition for I_{jk}^n , implies (9) in the text.

Mobile Labor

The Lagrangian of the problem in Definition 2 is

$$\begin{aligned} \mathcal{L} = & u - \sum_j \tilde{\omega}_j L_j (u - U(c_j, h_j)) - W^L \left(\sum_j L_j - L \right) \\ & - \sum_j P_j^D \left[c_j L_j + \chi \sum_{n=1}^N \sum_{k \in \mathcal{N}(j)} \tau_{jk}^n \left(\sum_{n=1}^N m^n Q_{jk}^n, I_{jk} \right) Q_{jk}^n - D_j(D_j^1, \dots, D_j^N) \right] - \sum_j P_j^H (h_j L_j - H_j) \\ & - \sum_j \sum_n P_j^n \left[D_j^n + \sum_{k \in \mathcal{N}(j)} (Q_{jk}^n + (1 - \chi) \tau_{jk}^n(Q_{jk}^n, I_{jk}) Q_{jk}^n) - F_j^n(L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n) - \sum_{i \in \mathcal{N}(j)} Q_{ij}^n \right] \\ & - \sum_j W_j \left[\sum_n L_j^n - L_j \right] - \sum_j \sum_l R_j^l \left[\sum_n V_j^n - V_j \right] \\ & - P_K \left(\sum_j \sum_{k \in \mathcal{N}(j)} \delta_{jk}^I I_{jk} - K \right) + \sum_{j,k} \zeta_{jk}^L (I_{jk} - \underline{I}_{jk}) + \sum_{j,k} \zeta_{jk}^{\bar{I}} (\bar{I}_{jk} - I_{jk}) \\ & + \sum_{j,k,n} \zeta_{jkn}^Q Q_{jk}^n + \sum_{j,n} \zeta_{jn}^L L_j^n + \sum_{j,n,l} \zeta_{jnl}^V V_j^{ln} + \sum_{j,n,l} \zeta_{jnl}^X X_j^{ln} + \sum_{j,n} \zeta_{jn}^D D_j^n + \sum_j \zeta_j^c c_j + \sum_j \zeta_j^h h_j \end{aligned}$$

where, in addition to the previous notation for the multipliers, in the first line we have defined $\tilde{\omega}_j$ and W^L as the multipliers of constraints (vi) and (vii) in Definition 2.

The first-order conditions with respect to consumption of traded services $[C_j^n]$, factor allocation within locations $[L_j^n]$, $[V_j^n]$ and $[X_j^n]$, optimal transport $[Q_{jk}^n]$, and optimal investment $[I_{jk}]$ are the same as in the problem without labor mobility. The first-order conditions with respect to u and L_j are:

$$\begin{aligned} [u] \quad & 1 = \sum_j L_j \tilde{\omega}_j \\ [L_j] \quad & P_j^D c_j + P_j^H h_j - \tilde{\omega}_j [U(c_j, h_j) - u] = W_j - W^L \end{aligned}$$

where from monotonicity of $U(c_j, h_j)$ it follows that

$$U(c_j, h_j) = \begin{cases} u & \text{if } L_j > 0, \\ 0 & \text{if } L_j = 0. \end{cases}$$

In addition, the first-order conditions with respect to consumption of traded and non-traded services, $[c_j]$ and $[h_j]$, are the same as in the problem without labor mobility replacing the planner's weights ω_j with the multipliers of the mobility constraint $\tilde{\omega}_j$. Combining $[L_j]$ with $[c_j]$ and $[h_j]$ gives the multiplier on the labor-mobility constraint. For populated locations:

$$\tilde{\omega}_j = \frac{W_j - W^L}{U_C(c_j, h_j) c_j + U_H(c_j, h_j) h_j}.$$

A.2 Symmetry in Infrastructure Investments

For the applications in Section 5 we impose symmetry in infrastructure levels as an additional restriction in the planner's problem, i.e., $I_{jk} = I_{kj}$. This section provides the first-order condition for I_{jk} in that case. The first-order condition with respect to I_{jk} is

$$[I_{jk}] \quad -\chi \sum_{n=1}^N \left(P_j^D Q_{jk}^n \frac{\partial \tau_{jk}^n(Q_{jk}, I_{jk})}{\partial I_{jk}} + P_k^D Q_{kj}^n \frac{\partial \tau_{kj}^n(Q_{kj}, I_{jk})}{\partial I_{jk}} \right) \\ - (1-\chi) \sum_{n=1}^N \left(P_j^n Q_{jk}^n \frac{\partial \tau_{jk}^n(Q_{jk}^n, I_{jk})}{\partial I_{jk}} + P_k^n Q_{kj}^n \frac{\partial \tau_{kj}^n(Q_{kj}^n, I_{jk})}{\partial I_{jk}} \right) + (\zeta_{jk}^L - \zeta_{jk}^T) = P_K (\delta_{jk}^I + \delta_{kj}^I). \quad (\text{A.3})$$

Assuming symmetry leaves all the remaining first-order conditions presented in Section A.1 unchanged. Under the log-linear specification (10) of the transport technology, the optimal infrastructure investment, conditional on $I_{jk} \in (\zeta_{jk}^L, \zeta_{jk}^T)$, is

$$I_{jk}^* = \left[\frac{\gamma}{P_K (\delta_{jk}^I + \delta_{kj}^I)} \left(\chi (\delta_{jk}^\tau P_j^D Q_{jk}^{1+\beta} + \delta_{kj}^\tau P_k^D Q_{kj}^{1+\beta}) + \sum_{n=1}^N (1-\chi) (\delta_{jk}^\tau P_j^n (Q_{jk}^n)^{1+\beta} + \delta_{kj}^\tau P_k^n (Q_{kj}^n)^{1+\beta}) \right) \right]^{\frac{1}{1+\gamma}}. \quad (\text{A.4})$$

As discussed in Section 5.2, to build $\delta_{jk}^{I,FOC}$ we use (A.4) under the symmetry assumption $\delta_{jk}^I = \delta_{kj}^I$. Setting $I_{jk}^* = I_{jk}^{obs}$, $\delta_{jk}^{I,FOC}$ can be backed out as function of calibrated parameters, the observed network I_{jk}^{OBS} , and the equilibrium prices generated by the calibrated model. Note that, to generate these prices, we use the model calibrated given the network I_{jk}^{OBS} , as discussed in Section 5.2.

A.3 Proofs of the Propositions

Proposition 1. (*Convexity of the Planner's Problem*) (i) Given the network $\{I_{jk}\}$, the joint optimal transport and allocation problem in the fixed (respectively mobile) labor case is a convex (respectively quasiconvex) optimization problem if $Q\tau_{jk}(Q, I_{jk})$ is convex in Q for all j and $k \in \mathcal{N}(j)$; and (ii) if in addition $Q\tau_{jk}(Q, I)$ is convex in both Q and I for all j and $k \in \mathcal{N}(j)$, then the full planner's problem including the network design problem from Definition (1) (resp. Definition (2)) is a convex (respectively quasiconvex) optimization problem. In either the joint transport and allocation problem, or the full planner's problem, strong duality holds when labor is fixed.

Proof. Consider the planner's problem from Definition 1. We can write it as

$$\max_{\left\{ C_j, \left\{ D_j^n, \left\{ Q_{jk}^n, I_{jk} \right\}_{k \in \mathcal{N}(j)} \right\}_{\forall j} \right\}} f = \sum_j \omega_j L_j U \left(\frac{C_j}{L_j}, \frac{H_j}{L_j} \right)$$

subject to: (i) availability of traded commodities,

$$g_j^1 = C_j + \chi \sum_{k \in \mathcal{N}(j)} \tau_{jk}(Q_{jk}, I_{jk}) Q_{jk} - D_j \left(D_j^1, \dots, D_j^N \right) \leq 0 \text{ for all } j;$$

(ii) the balanced-flows constraint,

$$g_{jn}^2 \equiv D_j^n + \sum_{k \in \mathcal{N}(j)} Q_{jk}^n [1 + (1-\chi) \tau_{jk}(Q_{jk}^n, I_{jk})] - F_j^n (L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n) - \sum_{i \in \mathcal{N}(j)} Q_{ij}^n \leq 0 \text{ for all } j, n;$$

(iii) the network-building constraint,

$$\sum_j \sum_{k \in \mathcal{N}(j)} \delta_{jk}^I I_{jk} \leq K;$$

and conditions (iv)-(v) in the text. Since constraints (iii)-(v) are linear, we need f to be concave and g_j^1 and g_{jn}^2 to

be convex. Since U is jointly concave in both its arguments, f is concave. $D_j(\{D_j^n\})$ is concave, hence g_j^1 is convex. If $Q\tau_{jk}(Q, I)$ is convex then g_{jn}^2 is the sum of linear and convex functions, hence it is convex. To show that this problem admits strong duality, a constraint qualification is required. Note first that constraints g_j^1 and g_{jn}^2 must hold with equality at an optimum and therefore can be substituted into the objective function. The remaining constraints (iii)-(v) are all linear and thus satisfy the Arrow-Hurwicz-Uzawa qualification constraint (Takayama (1985), Theorem 1.D.4). Hence, the global optimum must satisfy the KKT conditions and the duality gap is 0.⁴²

Consider now the planner's problem with labor mobility from Definition 2. Because U is homothetic, we can express it as $U = G(U_0(c, h))$, where G is an increasing continuous function and U_0 is homogeneous of degree 1. Therefore, imposing the change of variables $\tilde{u} = G^{-1}(u)$, the planner's problem can be restated as maximizing \tilde{u} subject to the convex constraints (i)-(v) and $L_j\tilde{u} \leq U_0(C_j, H_j)$. To make the latter constraint convex, let us denote $U_j = L_j\tilde{u}$ and replace \tilde{u} in the objective function by $\min_{j|L_j>0} \left\{ \frac{U_j}{L_j} \right\}$,⁴³ so that the problem becomes

$$\max_{C_j, \{D_j^n, L_j^n, \mathbf{V}_j^n, \{Q_{jk}^n, I_{jk}^n\}_{k \in \mathcal{N}(j)}\}, U_j, L_j} \min_{j|L_j>0} \left\{ \frac{U_j}{L_j} \right\}$$

subject to the convex restrictions (i)-(v) above as well as

$$U_j \leq U_0(C_j, H_j) \text{ for all } j.$$

The objective function is quasiconcave because U_j/L_j is quasiconcave and the minimum of quasiconcave functions is quasiconcave. In addition, all the restrictions are convex. Arrow and Enthoven (1961) then implies that the Karush-Kuhn-Tucker conditions are sufficient if the gradient of the objective function is different from zero at the candidate for an optimum, and here the gradient never vanishes. □

Proposition 2. (Optimal Network in Log-Linear Case) *When the transport technology is given by (10), the full planner's problem is a convex (resp. quasiconvex) optimization problem if $\beta \geq \gamma$. The optimal infrastructure is given by (12).*

Proof. First, note that if $\beta \geq \gamma$ then $Q\tau(Q, I) \propto Q^{1+\beta}I^{-\gamma}$ is convex in $Q \in \mathbb{R}_+$ and $I \in \mathbb{R}_+$. To see that, note that the determinant of the Hessian of $Q^{1+\beta}I^{-\gamma}$ is $(1+\beta)\gamma(\beta-\gamma)Q^{2\beta}I^{-2(\gamma+1)}$, which is positive for $Q \in \mathbb{R}_+$ and $I \in \mathbb{R}_+$ if $\beta \geq \gamma \geq 0$. Next, from the first-order condition for optimal infrastructure (9), if the solution to the planning problem implies $I_{jk} = \underline{L}_{jk}$ so that there is no investment, then:

$$\begin{aligned} P_K &\geq -\frac{1}{\delta_{jk}^I} \sum_n P_j^n Q_{jk}^n \frac{\partial \tau_{jk}^n}{\partial I_{jk}} \Big|_{I_{jk}=\underline{L}_{jk}} \\ &\geq \gamma \frac{\delta_{jk}^\tau \sum_n P_j^n (Q_{jk}^n)^{1+\beta}}{\delta_{jk}^I \underline{L}_{jk}^{\gamma+1}} \\ &\geq \frac{\gamma(1+\beta)^{-\frac{1+\beta}{\beta}} \sum_{n: P_k^n > P_j^n} P_j^n \left(\frac{P_k^n}{P_j^n} - 1 \right)^{\frac{1+\beta}{\beta}}}{\delta_{jk}^I (\delta_{jk}^\tau)^{\frac{1}{\beta}} \underline{L}_{jk}^{\frac{\beta-\gamma}{\beta}}}, \end{aligned}$$

where the second line follows from (10) and the third line follows from (11). The last inequality is equivalent to $\underline{L}_{jk} \geq I_{jk}^*$ for I_{jk}^* defined in (14). Therefore, if $\underline{L}_{jk} < I_{jk}^*$ then $I_{jk} > \underline{L}_{jk}$ and $I_{jk} = I_{jk}^*$. Moreover, if there is any n such that $P_k^n \neq P_j^n$ then $I_{jk}^* > 0$. □

⁴²Despite having substituted constraints g_j^1 and g_{jn}^2 into the objective function, the multipliers for these constraints, P_j^D and P_j^n , can be recovered from the above KKT conditions such that $\omega_j U_C(c_j, h_j) = P_j^D$ and $P_j^D \partial C_j^T / \partial C_j^n = P_j^n$.

⁴³Since the objective function is strictly increasing in \tilde{u} and because \tilde{u} only shows up in the constraints $L_j\tilde{u} \leq U_0(C_j, H_j)$ for all j , it is necessarily the case that $\tilde{u} = \min_{j|L_j>0} U_j/L_j$.

Proposition 3. (Tree Property) Assume that $\lim_{c \rightarrow 0^+} U_C(c, h) = \infty$. In the absence of a pre-existing network (i.e., $L_{jk} = 0$), if the transport technology is given by (10) and satisfies $\gamma > \beta$, and if there is a unique commodity produced in a single location, then the optimal transport network is a tree.

Proof. See Section G of the Supplementary Material. □

Definition 3. The decentralized equilibrium without labor mobility consists of quantities $c_j, h_j, D_j, D_j^n, L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n, \{Q_{jk}^n\}_{k \in \mathcal{N}(j)}$, goods prices $\{p_j^n\}_n, p_j^D, p_j^H$ and factor prices $w_j, \{r_j^m\}_m$ in each location j such that:

(i)(a) consumers optimize:

$$\begin{aligned} \{c_j, h_j\} &= \arg \max_{\hat{c}_j, \hat{h}_j} U(\hat{c}_j, \hat{h}_j) \\ p_j^D \hat{c}_j + p_j^H \hat{h}_j &= e_j \equiv w_j + t_j, \end{aligned}$$

where e_j are expenditures per worker in j and where p_j^D is the price index associated with D_j (D_j^1, \dots, D_j^N) at prices $\{p_j^n\}_n$ and t_j is a transfer per worker located in j . The set of transfers satisfy

$$\sum_j t_j L_j = \Pi$$

where Π adds up the aggregate returns to the portfolio of fixed factors and the government tax revenue,

$$\Pi = \sum_j p_j^H H_j + \sum_j \sum_m r_j^m V_j^m + \sum_j \sum_{k \in \mathcal{N}(j)} \sum_n t_{jk}^n p_k^n Q_{jk}^n;$$

(i)(b) firms optimize:

$$\{L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n\} = \underset{\hat{L}_j^n, \hat{\mathbf{V}}_j^n, \hat{\mathbf{X}}_j^n}{\operatorname{argmax}} p_j^n F_j^n(\hat{L}_j^n, \hat{\mathbf{V}}_j^n, \hat{\mathbf{X}}_j^n) - w_j \hat{L}_j^n - \sum_m r_j^m \hat{V}_j^{mn};$$

(i)(c) the transport companies optimize,

$$\pi_{od}^n = \max_{r=(j_0, \dots, j_\rho) \in \mathcal{R}_{od}} p_d^n - p_o^n - \sum_{k=0}^{\rho-1} \left(\chi P_{j_k}^D m^n \tau_{j_k j_{k+1}} + (1 - \chi) P_{j_k}^n \tau_{j_k j_{k+1}}^n \right) - \sum_{k=0}^{\rho-1} p_{j_{k+1}} t_{j_k j_{k+1}}^n,$$

for all $(o, d) \in \mathcal{J}^2$, where $\mathcal{R}_{od} = \{(j_0, \dots, j_\rho) \in \mathcal{J}^{\rho+1}, \rho \in \mathbb{N} \mid j_0 = o, j_\rho = d, j_{k+1} \in \mathcal{N}(j_k) \text{ for all } 0 \leq k < \rho\}$ is the set of routes from o to d , and there is free entry to delivering products from every source to every destination: $\pi_{od}^n \leq 0$ for all $(o, d) \in \mathcal{J}^2$, = if good n is shipped from o to d .

(i)(d) producers of final commodities optimize:

$$\{D_j^n\} = \underset{\hat{D}_j^n}{\operatorname{argmax}} D_j \left(\{\hat{D}_j^n\} \right) - \sum_j p_j^n \hat{D}_j^n;$$

as well as the market-clearing and non-negativity constraints (i), (ii), (iv), and (v) from Definition 1.

If, in addition, labor is mobile, then the decentralized equilibrium also consists of utility u and employment $\{L_j\}$ such that

$$u = U_j(c_j, h_j)$$

whenever $L_j > 0$, and the labor market clearing condition (vii) from Definition 2 holds.

Proposition 4. (First and Second Welfare Theorems) If the tax on shipments of product n from j to k , denominated in the same unit as transport costs, is $t_{jk}^n = \chi m^n \varepsilon_{Q, jk}^\tau \tau_{jk} (Q_{jk}, I_{jk}) + (1 - \chi) \varepsilon_{Q, jkn}^\tau \tau_{jk}^n (Q_{jk}^n, I_{jk})$ where $\varepsilon_{Q, jk}^n = \partial \log \tau_{jk}^n / \partial \log Q_{jk}^n$ ($\chi = 0$) and $\varepsilon_{Q, jk} = \partial \log \tau_{jk} / \partial \log Q_{jk}$ ($\chi = 1$), then: (i) if labor is immobile, the competitive allocation coincides with the planner's problem under specific planner's weights ω_j and, conversely, the

planner's allocation can be implemented by a market allocation with specific transfers t_j ; and (ii) if labor is mobile, the competitive allocation coincides with the planner's problem if and only if all workers own an equal share of fixed factors and tax revenue, i.e., $t_j = \frac{\Pi}{L}$. In either case, the price of good n in location j , p_j^n , equals the multiplier on the balanced-flows constraint in the planner's allocation, P_j^n .

Proof. Equivalence of the First-order Conditions. Condition (i)(c) from the definition of the market allocation implies that the free entry condition of shippers holds for every pair of neighbors; i.e., for every $j \in \mathcal{J}$ and $k \in \mathcal{N}(j)$,

$$p_k^n \leq p_j^n + (1 - \chi) p_j^n (\tau_{jk}^n + t_{jk}^n) + \chi P_j^D (m^n \tau_{jk} + t_{jk}^n), \quad \text{if } Q_{jk}^n > 0. \quad (\text{A.5})$$

This condition is consistent with the first-order condition (8) from the planner's problem if and only if the tax scheme is defined as in the proposition. We must further show that, under this tax scheme, a route is the solution to (i)(c) if and only if it is used in the solution to the planner's problem, which we establish at the end of this proof.

Without labor mobility, the rest of the allocation corresponds to a standard neoclassical economy with convex technologies and preferences where the welfare theorems hold. Specifically, the first-order conditions from the consumer and firm optimization problems (i)(a) and (i)(b) yield:

$$\begin{aligned} [\hat{c}_j] \quad & \left(\frac{1}{\lambda_j} \right) U_C(c_j, h_j) = p_j^D \\ [\hat{h}_j] \quad & \left(\frac{1}{\lambda_j} \right) U_H(c_j, h_j) = p_j^H \\ [\hat{D}_j^n] \quad & p_j^D \frac{\partial D_j}{\partial D_j^n} = p_j^n \\ [\hat{L}_j^n] \quad & \frac{\partial Y_j^n}{\partial L_j^n} P_j^n \leq w_j, \quad \text{if } L_{jn} > 0 \\ [\hat{V}_j^{mn}] \quad & \frac{\partial Y_j^n}{\partial V_j^{mn}} P_j^n \leq r_j^m, \quad \text{if } V_j^{mn} > 0. \end{aligned}$$

Since the market clearing constraints are the same in the market's and the planner's allocation, the planner's allocation coincides with the market if the planner's weights are such that the planner's FOC for c_j coincide with the market. This is the case if the weight ω_j from the planner's problem equals the inverse of the multiplier on the budget constraint from the consumer's optimization problem (i)(a) in the market allocation. To find that weight, using that U is homothetic we can write $U = G(U_0(c, h))$, where U_0 is homogeneous of degree 1. Then, the planner's allocation coincide with the market's under weights

$$\omega_j = \frac{e_j}{G'(U_0(c_j, h_j)) U_0(c_j, h_j)},$$

where e_j is the expenditure per worker and c_j, h_j are the consumption per worker of the traded and non-traded good in the market allocation. If U is homogeneous of degree one, then $\omega_j = P_j^U$, where P_j^U is the price index associated with $U(c_j, h_j)$ at the market equilibrium prices p_j^D, p_j^H . In the opposite direction, given arbitrary weights ω_j , the market allocation implements the planner's under the transfers $t_j = P_j^D c_j + P_j^H h_j - W_j$ constructed using the quantities $\{c_j, h_j\}$ from the planner's allocation and the multipliers $\{P_j^D, P_j^H\}$ and W_j corresponding to the constraints (i) and (iv) of the planner's problem, respectively.

For the case with labor mobility, note that, for populated locations, the planner's first-order condition with respect to L_j implies:

$$P_j^D c_j + P_j^H h_j = W_j - W^L.$$

Therefore, the market allocation and the planner's solution coincide if and only if in the market allocation expenditure per worker in location j takes the form $e_j = w_j + \text{Constant}$ for all j . The only transfer scheme delivering the same transfer per capita is $t_j = \frac{\Pi}{L}$.

Equivalence of Least Cost Routes. We want to establish that any route used in the planner's problem is a

solution to (i)(c) in Definition 3 under the proposed tax scheme. Fix good n . We introduce the following notation: for a given route $r = (j_0, \dots, j_\rho) \in \mathcal{R}_{od}$, we denote by f_r^n the matrix of flows

$$f_r^n(j, k) = \begin{cases} 1 & \text{if } \exists l, 0 \leq l \leq \rho - 1 \text{ and } j = j_l, k = j_{l+1} \\ 0 & \text{otherwise} \end{cases}.$$

Consider an optimal route from o to d , $r^* = (j_0^*, \dots, j_\rho^*) \in \mathcal{R}_{od}$, i.e., such that $Q_{j_k^* j_{k+1}^*}^n > 0$ at the optimum of the planner's problem ($\zeta_{j_k^* j_{k+1}^*}^Q = 0$). We now consider redirecting a marginal amount of goods $\varepsilon > 0$ from r^* to some other route $r = (j_0, \dots, j_\rho) \in \mathcal{R}_{od}$. In other words, denoting $\mathbf{Q}^n = (Q_{jk}^n)_{j \in \mathcal{J}, k \in \mathcal{N}(j)}$, we consider the perturbation $\mathbf{Q}^n + \varepsilon f_r^n - \varepsilon f_{r^*}^n$. The first-order effect of the deviation around the optimum must reduce the Lagrangian:

$$\mathcal{L}(\mathbf{Q}^n + \varepsilon f_r^n - \varepsilon f_{r^*}^n) - \mathcal{L}(\mathbf{Q}^n) \leq 0.$$

To translate this condition into a minimum cost route problem we decompose the first-order impact on the Lagrangian,

$$\mathcal{L}(\mathbf{Q}^n + \varepsilon f_r^n - \varepsilon f_{r^*}^n) - \mathcal{L}(\mathbf{Q}^n) = \nabla_Q \mathcal{L}(\mathbf{Q}^n) \cdot (f_r^n - f_{r^*}^n) \varepsilon + o(\varepsilon),$$

and evaluate each term separately. The first deviation term,

$$\begin{aligned} \nabla_Q \mathcal{L}(\mathbf{Q}^n) \cdot f_r^n &= \sum_{l=0}^{\rho-1} \left[-P_{j_l}^n + P_{j_{l+1}}^n + \zeta_{j_l j_{l+1} n}^Q - \chi P_{j_l}^D \left(m^n \tau_{j_l j_{l+1}}^n + Q_{j_l j_{l+1}} m^n \frac{\partial \tau_{j_l j_{l+1}}^n}{\partial Q_{j_l j_{l+1}}^n} \right) \right. \\ &\quad \left. - (1 - \chi) P_{j_l}^n \left(\tau_{j_l j_{l+1}} + Q_{j_l j_{l+1}}^n \frac{\partial \tau_{j_l j_{l+1}}^n}{\partial Q_{j_l j_{l+1}}^n} \right) \right], \end{aligned}$$

simplifies to

$$\begin{aligned} \nabla_Q \mathcal{L}(\mathbf{Q}^n) \cdot f_r^n &= \left[P_d^n - P_o^n + \sum_{l=0}^{\rho-1} \zeta_{j_l j_{l+1} n}^Q - \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D m^n \tau_{j_l j_{l+1}} + (1 - \chi) P_{j_l}^n \tau_{j_l j_{l+1}}^n \right) \right. \\ &\quad \left. - \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D Q_{j_l j_{l+1}} m^n \frac{\partial \tau_{j_l j_{l+1}}^n}{\partial Q_{j_l j_{l+1}}^n} + (1 - \chi) P_{j_l}^n Q_{j_l j_{l+1}}^n \frac{\partial \tau_{j_l j_{l+1}}^n}{\partial Q_{j_l j_{l+1}}^n} \right) \right]. \end{aligned}$$

Using the definition of the optimal tax, we have that

$$t_{jk}^n = \chi Q_{jk} m^n \frac{\partial \tau_{jk}^n}{\partial Q_{jk}^n} + (1 - \chi) Q_{jk}^n \frac{\partial \tau_{jk}^n}{\partial Q_{jk}^n}.$$

Substituting into the previous deviation term, we obtain

$$\begin{aligned} \nabla_Q \mathcal{L}(\mathbf{Q}^n) \cdot f_r^n &= \left[P_d^n - P_o^n - \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D m^n \tau_{j_l j_{l+1}} + (1 - \chi) P_{j_l}^n \tau_{j_l j_{l+1}}^n \right) \right. \\ &\quad \left. - \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D + (1 - \chi) P_{j_l}^n \right) t_{j_l j_{l+1}}^n + \sum_{l=0}^{\rho-1} \zeta_{j_l j_{l+1} n}^Q \right]. \end{aligned}$$

By assumption, the total deviation $\mathbf{Q}^n + \varepsilon f_r^n - \varepsilon f_{r^*}^n$ has a negative impact on the Lagrangian for the feasible deviation

$\varepsilon > 0$, so that $\nabla_Q \mathcal{L}(\mathbf{Q}^n) \cdot (f_r^n - f_r^{n*}) \leq 0$. Using that $\zeta_{j_i^* j_{i+1}^*}^Q = 0$, we get:

$$\begin{aligned} & P_o^n + \sum_{l=0}^{\rho^*-1} \left(\chi P_{j_l^*}^D m^n \tau_{j_l^* j_{l+1}^*}^n + (1-\chi) P_{j_l^*}^n \tau_{j_l^* j_{l+1}^*}^n \right) + \sum_{l=0}^{\rho^*-1} \left(\chi P_{j_l^*}^D + (1-\chi) P_{j_l^*}^n \right) t_{j_l^* j_{l+1}^*}^n \\ & \leq P_o^n + \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D m^n \tau_{j_l j_{l+1}}^n + (1-\chi) P_{j_l}^n \tau_{j_l j_{l+1}}^n \right) + \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D + (1-\chi) P_{j_l}^n \right) t_{j_l j_{l+1}}^n - \sum_{l=0}^{\rho-1} \zeta_{j_l j_{l+1}}^Q \\ & \leq P_o^n + \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D m^n \tau_{j_l j_{l+1}}^n + (1-\chi) P_{j_l}^n \tau_{j_l j_{l+1}}^n \right) + \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D + (1-\chi) P_{j_l}^n \right) t_{j_l j_{l+1}}^n. \end{aligned}$$

Hence, the optimal route r^* is solution to the least-cost route problem

$$\min_{r=(j_0, \dots, j_\rho) \in \mathcal{R}_{od}} P_o^n + \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D m^n \tau_{j_l j_{l+1}}^n + (1-\chi) P_{j_l}^n \tau_{j_l j_{l+1}}^n \right) + \sum_{l=0}^{\rho-1} \left(\chi P_{j_l}^D + (1-\chi) P_{j_l}^n \right) t_{j_l j_{l+1}}^n,$$

where we recognize condition (i)(c) of Definition 3.

Finally, that the minimum cost route problem in the case of own-good congestion ($\chi = 0$) is equivalent to

$$\min_{r=(j_0, \dots, j_\rho) \in \mathcal{R}_{od}} P_o^n \prod_{l=0}^{\rho-1} \left(1 + \tau_{j_l j_{l+1}}^n \right) + \sum_{l=0}^{\rho-1} P_{j_{l+1}}^n t_{j_l j_{l+1}}^n \prod_{k=l}^{\rho-1} \left(1 + \tau_{j_k j_{k+1}}^n \right),$$

since $P_{j_{l+1}}^n = \left(1 + \tau_{j_l j_{l+1}}^n \right) P_{j_l}^n + P_{j_{l+1}}^n t_{j_l j_{l+1}}^n$ along any used path. □

Proposition 5. *If the global convexity condition of Proposition 1 is satisfied and the toll is consistent with the optimal Pigouvian tax ($\theta_{jk}^n = \chi P_j^D m^n \varepsilon_{Q, jkn}^{\tau} \tau_{jk}^n + (1-\chi) P_j^n \varepsilon_{Q, jkn}^{\tau} \tau_{jk}^n$) then the decentralized infrastructure choice implements the optimal network investment.*

Proof. To ease notation we focus on the case with own-good congestion ($\chi = 0$), but the case with cross-goods congestion ($\chi = 1$) can be derived following similar steps. Consider the problem of a regulated monopoly on link jk allowed to charge a per-unit toll θ_{jk}^n on good n . The monopolist can purchase asphalt at price p_K . We assume that the government forbids entry on unused links or sets a price too low for entry, allowing us to focus on links used at the social optimum ($Q_{jk}^n > 0$ for some $n, I_{jk} > 0$). Free-entry of shipping companies on link jk yields

$$p_k^n \leq p_j^n \left(1 + \tau_{jk}^n(Q_{jk}^n, I_{jk}^n) \right) + \theta_{jk}^n, = \text{if } Q_{jk}^n > 0.$$

Under the assumption that τ_{jk}^n is strictly increasing in Q_{jk}^n , the demand for transport at any level of infrastructure I_{jk} given prices is

$$Q_{jk}^n(I_{jk}; p) = \text{inv}_Q \tau_{jk} \left(\frac{p_k^n - p_j^n - \theta_{jk}^n}{p_j^n}, I_{jk} \right) \text{ for } p_k^n \geq p_j^n + \theta_{jk}^n, \quad (\text{A.6})$$

where $\text{inv}_Q \tau(Q, I)$ denotes the inverse of function τ with respect to Q . The monopoly solves the profit maximization problem

$$\max_{I_{jk}} \sum_n \theta_{jk}^n Q_{jk}^n(I_{jk}) - p_K \delta_{jk}^I I_{jk}$$

subject to (A.6). It can be shown that convexity of $\tau(Q, I)$ in Q and I is sufficient for the problem to be concave. The first-order condition over infrastructure is

$$\sum_n \frac{\partial}{\partial I_{jk}} \left[\text{inv}_Q \tau \left(\frac{p_k^n - p_j^n - \theta_{jk}^n}{p_j^n}, I_{jk} \right) \right] \theta_{jk}^n = p_K \delta_{jk}^I.$$

From the implicit function theorem, we have that $\frac{\partial}{\partial I_{jk}} \left[\text{inv}_Q \tau_{jk} \left(\frac{p_k^n - p_j^n - \theta_{jk}^n}{p_j^n}, I_{jk} \right) \right] = -\frac{\partial \tau_{jk}}{\partial I_{jk}} / \frac{\partial \tau_{jk}}{\partial Q_{jk}^n}$. In turn, imple-

menting the efficient flows Q_{jk}^n requires the toll θ_{jk}^n corresponding to the Pigouvian congestion tax from Proposition 4, that is

$$\theta_{jk}^n = p_j^n t_{jk}^n = p_j^n \varepsilon_{Q,jkn}^\tau \tau_{jk} (Q_{jk}^n, I_{jk}) = p_j^n Q_{jk}^n \frac{\partial \tau_{jk}^n}{\partial Q_{jk}^n}.$$

Combining the last two expressions, the monopolist's first-order condition thus becomes the same as (13), except for the price of asphalt p_K . Imposing market clearing in K , the set of conditions in the decentralized allocation coincide with the planner, implying that p_K equals the Lagrange multiplier P_K and other prices equal their corresponding multipliers, $p_j^n = P_j^n$ and $p_j^D = P_j^D$. Moreover, under the global convexity condition from Proposition 1, first-order conditions of the builders are sufficient, which demonstrates the result. \square

B Appendix to Section 4 (Calibration and Counterfactuals)

Construction of $\mathcal{P}(j, k)$ The definition of the weights $\omega_{jk}(s)$ assigned to the construction of I_{jk}^{obs} involves the cheapest path $\mathcal{P}(j, k)$ for all $j \in \mathcal{J}$ and $k \in \mathcal{N}(j)$ in every country. To find $\mathcal{P}(j, k)$, we first convert the shapefile with all the road segments from EuroGeographics into a weighted graph, where each edge corresponds to a segment s on the road network. We define $\mathcal{P}(j, k)$ as the shortest path between j and k under the segment-specific weights $length_s * lanes_s^{-\chi_{lane}} * \chi_{use}^{1-nat_s} * \chi_{paved}^{1-paved_s} * \chi_{median}^{1-median_s}$, where $length_s$ is the length of the segment, $lanes_s$ is the number of lanes, nat_s equals 1 if the segment belongs to a national road, $paved_s$ equals 1 if the segment is paved, and $median_s$ equals 1 if the segment has a median. When information on number of lanes is missing we assign a number of lanes equal to the minimum of 1 observed in the data. When the information is missing we define the road use as non-national. We parametrize χ_{lane} , χ_{use} , χ_{paved} , and χ_{median} based on the extent by which adding a lane, using a national road, using paved road, or using a road with a median reduces road user costs. Table 4 of Combes and Lafourcade (2005) reports that, in France, the reference cost per km. in a national road with at least 4 lanes is 25% higher than in other national roads. In our road network data for France, the average number of lanes in national roads with at least 4 lanes is 4.43, and the average number of lanes in national roads with less than 4 lanes is 1.9. From this, we infer that adding 2.5 lanes on top of 2 lanes, a 125% increase in the number of lanes, reduces costs by 25%, implying an elasticity of user costs with respect to number of lanes of $\chi_{lane} = \frac{25\%}{125\%} = 0.2$ in absolute value. In addition, Table 4 in Combes and Lafourcade (2005) reports that the total reference cost is about 7% higher on “secondary roads” relative to “other national roads”, from which we infer $\chi_{use} = 1.07$. According to Figueroa et al. (2013), road user costs are 35% higher on gravel relative to paved roads, implying $\chi_{paved} = 1.35$, and according to Tay and Churchill (2007), adding a median increases speed by 5%, implying $\chi_{median} = 1.05$.

Calibration of β and γ To parametrize β and γ we assume that: i) trade costs are a linear function of shipping time, $\tau_{jk} = a_{jk}^\tau \frac{dist_{jk}}{S_{jk}}$; ii) shipping speed is a log-linear function of the number of vehicles and road lane kilometers,

$$S_{jk} = a_{jk}^S I_{jk}^\gamma V_{jk}^{-\beta}; \quad (\text{A.7})$$

and iii) the total number of vehicles is a linear function of the quantity of goods that is shipped, $V_{jk} = a_{jk}^D Q_{jk}$. These assumptions are consistent with the functional form 10 for $\tau_{jk}(Q, I)$. To recover the parameters γ and β , one would ideally like to estimate the relationship between speed, roads, and vehicles in (A.7) across links. This relationship is estimated by Couture et al. (2018) across cities.⁴⁴ Equation (2) in their paper assumes a log-linear relationship

⁴⁴Couture et al. (2018) estimate the parameters from data on movements of vehicles, regardless of the trip purpose (i.e., it may include transport of passengers or goods). We assume that the speed of vehicles transporting goods responds to traffic and to highway lanes similarly to vehicles transporting passengers, and that the total number of vehicles V_{jk} is a linear function of the number vehicles transporting goods. Data from the 2015 E-Road traffic census in Europe shows a correlation of 0.81 between the average daily traffic of all vehicles and of vehicles used to transport goods across measuring posts in European highways. On average across measuring stations, 16% of all vehicles are used for transport.

between speed, roads, and vehicle travel time, defined as vehicle-kilometers (i.e., vehicles times distance) over speed. To measure roads they use the log of lane-kilometers of interstate highways in the U.S., which corresponds to our measure of I_{jk} . Assuming that their estimates would hold at the level of a connection between populated areas in our data, in our notation equation (2) of their paper can be written:⁴⁵

$$\ln S_{jk} = \alpha^{CDT} \ln I_{jk} - \theta^{CDT} \ln \left(\frac{V_{jk} * dist_{jk}}{S_{jk}} \right) + \varepsilon_{jk}, \quad (\text{A.8})$$

where we use α^{CDT} and θ^{CDT} to refer to α and β in their paper. These parameters translate to ours as follows: $\alpha^{CDT} \equiv \frac{\gamma}{1+\beta}$ and $\theta^{CDT} \equiv \frac{\beta}{1+\beta}$. When $\alpha^{CDT} < \theta^{CDT}$ there are decreasing returns to scale in the provision of vehicle kilometers traveled. Couture et al. (2018) find decreasing returns to scale ($\alpha^{CDT} < \theta^{CDT} \rightarrow \gamma < \beta$) across all their specification (Tables 5 and 6 of their paper) using a variety of OLS and IV approaches. Their preferred estimate (column 6 of table 5) yields $\alpha^{CDT} = 0.09$ and $\theta^{CDT} = 0.13$, implying $\gamma = 0.10$ and $\beta = 0.13$.

Construction of Ruggedness Measure We use elevation data from the ETOPO1 Global Relief Model. The ETOPO1 dataset corresponds to a 1 arc-minute degree grid. We construct ruggedness for each cell as the average ruggedness across the 900 arc-minute cells from the ETOPO1 dataset contained in each 0.5 arc-degree cell in our discretized maps. We use the standard ruggedness index by Riley et al. (1999). Letting $\mathcal{J}^{etopo}(j)$ be the set of cells in ETOPO1 contained in each cell $j \in \mathcal{J}$ of our discretization and $\mathcal{N}^{etopo}(i)$ be the 8 neighboring cells to each cell in ETOPO1, this index is defined as: $rugged_j = \left(\sum_{i \in \mathcal{J}^{etopo}(j)} \sum_{k \in \mathcal{N}^{etopo}(i)} (elev_i - elev_k)^2 \right)^{1/2}$; i.e., it is the standard deviation of the difference in elevation across neighboring cells. Then, we define $rugged_{jk}$ in (21) as $rugged_{jk} = \frac{1}{2} (rugged_j + rugged_k)$.

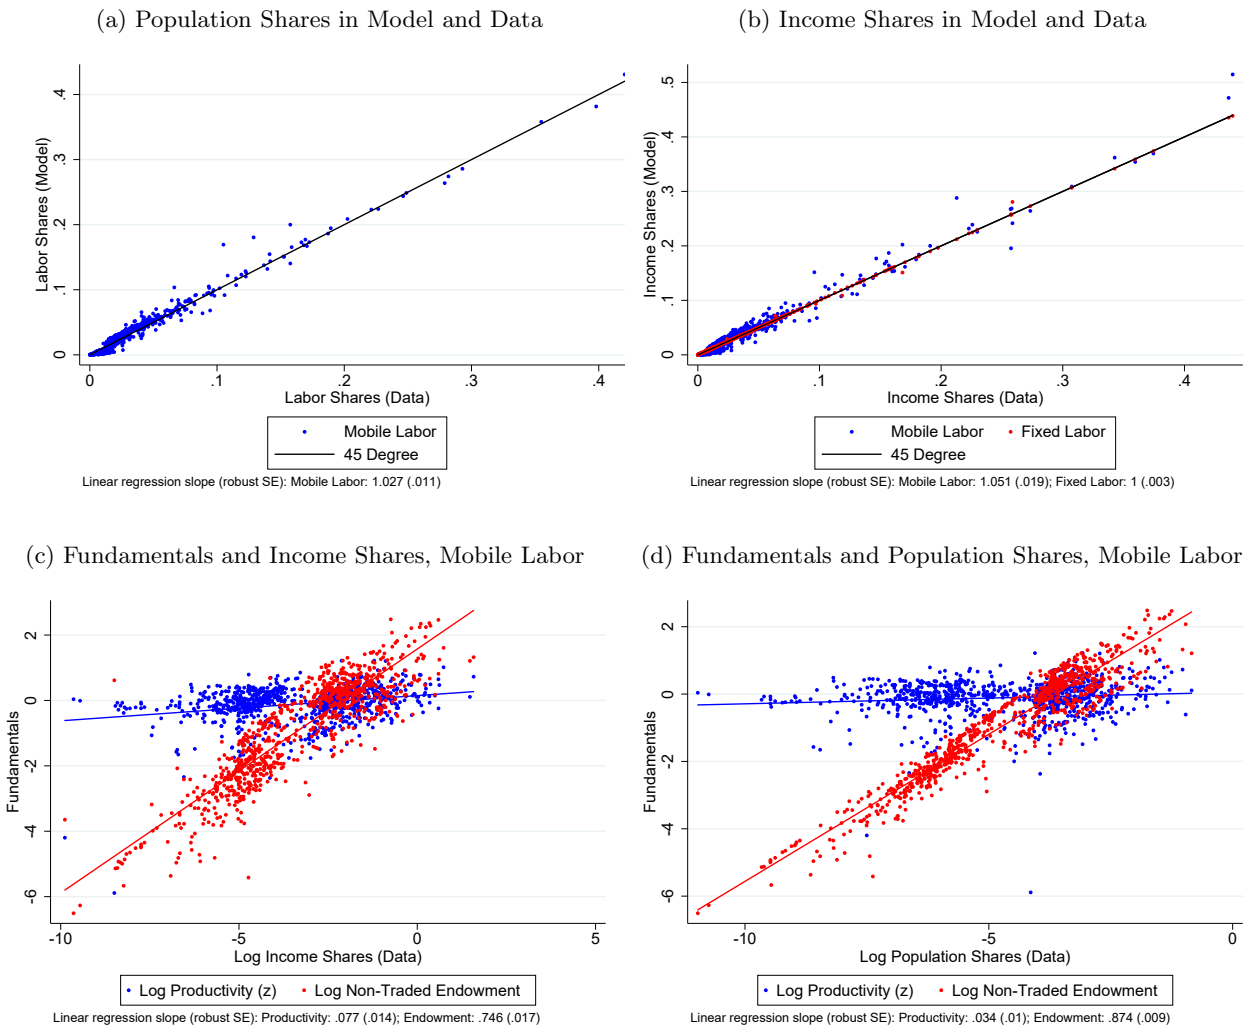
Country	Code	Actual Road Network			Discretization		
		Length (Km.)	Number of Segments	Average Lanes per Km.	Number of Cells	Length (Km.)	Average Infrastructure Index
		(1)	(2)	(3)	(4)	(5)	(6)
Austria	AT	17229	6156	2.36	46	9777	1.53
Belgium	BE	19683	10472	2.49	21	3535	2.12
Switzerland	CHLI	14451	11043	2.26	24	5307	1.71
Cyprus	CY	2818	947	2.29	20	2077	0.69
Czech Republic	CZ	28651	10183	2.20	48	10593	1.01
Germany	DE	115173	66423	2.42	54	27227	2.61
Denmark	DK	15358	6297	2.20	30	5418	0.99
Spain	ES	101990	18048	2.39	61	30686	1.80
Finland	FI	70394	9221	2.04	70	34057	0.28
France	FR	128787	38668	2.05	74	38754	2.12
Georgia	GE	28682	9009	1.95	39	8420	0.34
Hungary	HU	32728	9235	2.10	50	11572	0.95
Ireland	IE	24952	4144	2.10	47	10671	0.63
Italy	IT	77586	44155	2.32	36	14742	2.18
Lithuania	LT	10682	1586	2.39	43	9408	0.59
Luxembourg	LU	1778	863	2.31	8	574	0.73
Latvia	LV	11495	2102	2.03	47	9857	0.33
Moldova	MD	8527	1457	2.21	20	4359	0.38
Macedonia	MK	5578	908	2.15	13	2205	0.27
Northern Ireland	ND	7087	1888	2.18	12	1780	0.67
Netherlands	NL	14331	8383	2.68	20	3884	2.55
Portugal	PT	15034	4933	2.10	43	11116	1.49
Slovenia	SI	7801	2441	2.19	12	1525	1.25
Slovakia	SK	11406	2601	2.18	30	5880	1.04

Table A.1: Summary Statistics of Actual and Discretized Road Network by Country

Note: Columns (1) to (3) report statistics from EuroGlobalMap, and Columns (4) to (6) report statistics from the discretization of road networks described in Section 5.1.

⁴⁵In their notation, vehicle travel time (number of vehicles times time of travel) is $VTT_i = \frac{VKT_i}{S_i}$, where VKT is vehicle kilometers (number of vehicles times distance) and S_i is speed (distance over time of travel). In our notation, therefore, $VKT_i = V_{jk} \times dist_{jk}$.

Figure A.1: Calibration of Population and Income Shares, all Locations and Countries



Notes: The figures pool the 868 cells across the 24 countries in the convex case of the parameters for the calibration with 10 differentiated goods. Similar relationships hold for the non-convex case. In the panels (c) to (e), log-productivity and log-endowment of the non-traded good per capita are demeaned within each country.

ICC	Non-Convex			Convex		
	Expansion (GEO)	Expansion (FOC)	Misallocation	Expansion (GEO)	Expansion (FOC)	Misallocation
Austria	3.5%	1.0%	3.1%	2.3%	0.3%	2.3%
Belgium	1.0%	0.4%	0.8%	0.6%	0.1%	0.5%
Cyprus	1.5%	0.6%	1.3%	1.0%	0.2%	1.0%
Czech Republic	1.6%	0.7%	1.3%	1.0%	0.2%	1.0%
Denmark	8.3%	1.0%	7.8%	4.3%	0.4%	4.3%
Finland	5.3%	1.2%	4.7%	2.6%	0.4%	2.3%
France	4.0%	1.0%	3.4%	2.7%	0.4%	2.5%
Georgia	3.4%	1.3%	3.1%	2.2%	0.3%	2.2%
Germany	2.7%	1.2%	2.3%	1.8%	0.3%	1.7%
Hungary	3.1%	0.9%	2.7%	1.7%	0.2%	1.7%
Ireland	2.8%	0.9%	2.4%	1.8%	0.3%	1.7%
Italy	2.5%	1.2%	1.9%	1.6%	0.4%	1.3%
Latvia	4.0%	0.9%	3.6%	2.5%	0.3%	2.7%
Lithuania	3.4%	1.0%	3.0%	2.4%	0.3%	2.3%
Luxembourg	0.2%	0.2%	0.1%	0.2%	0.1%	0.1%
Macedonia	1.0%	0.6%	0.9%	0.5%	0.1%	0.5%
Moldova	1.8%	0.6%	1.5%	1.0%	0.2%	0.9%
Netherlands	1.6%	0.5%	1.4%	1.2%	0.2%	1.1%
Northern Ireland	1.2%	0.6%	1.0%	0.8%	0.1%	0.8%
Portugal	2.6%	0.9%	2.1%	1.4%	0.3%	1.2%
Slovakia	2.7%	1.3%	2.3%	1.9%	0.3%	1.9%
Slovenia	1.8%	0.6%	1.6%	1.4%	0.2%	1.4%
Spain	5.6%	2.3%	4.7%	3.7%	0.5%	3.5%
Switzerland	1.8%	0.7%	1.5%	1.0%	0.2%	0.9%
Average	2.8%	0.9%	2.4%	1.7%	0.3%	1.7%

Table A.2: Welfare Gains From Optimal Reallocation or Expansion of Current Networks, Fixed Labor

ICC	Non-Convex			Convex		
	Expansion (GEO)	Expansion (FOC)	Misallocation	Expansion (GEO)	Expansion (FOC)	Misallocation
Austria	3.4%	1.9%	2.8%	2.4%	0.4%	2.5%
Belgium	1.0%	0.5%	0.8%	0.6%	0.1%	0.5%
Czech Republic	1.6%	1.3%	1.3%	1.1%	0.2%	1.1%
Denmark	9.6%	1.8%	8.9%	4.5%	0.4%	4.5%
France	4.3%	2.2%	3.4%	3.1%	0.5%	2.9%
Germany	2.4%	2.1%	2.3%	2.0%	0.3%	1.8%
Hungary	3.1%	1.1%	2.7%	1.8%	0.3%	1.8%
Ireland	2.8%	1.5%	2.3%	1.8%	0.3%	1.7%
Finland	6.4%	1.6%	5.7%	3.1%	0.4%	2.8%
Italy	2.6%	2.4%	2.0%	1.6%	0.4%	1.3%
Latvia	4.0%	1.1%	3.6%	2.5%	0.3%	2.6%
Lithuania	3.8%	1.1%	3.3%	2.6%	0.3%	2.6%
Moldova	1.8%	0.6%	1.6%	1.1%	0.2%	1.0%
Luxembourg	0.2%	0.1%	0.1%	0.1%	0.1%	0.1%
Macedonia	0.9%	0.6%	0.8%	0.5%	0.1%	0.4%
Northern Ireland	1.1%	0.6%	0.9%	0.8%	0.1%	0.8%
Netherlands	1.6%	0.5%	1.3%	1.4%	0.2%	1.3%
Slovakia	3.1%	1.9%	2.7%	2.2%	0.3%	2.2%
Portugal	2.4%	1.3%	2.1%	1.3%	0.3%	1.1%
Slovenia	1.7%	0.7%	1.5%	1.3%	0.1%	1.4%
Switzerland	1.7%	0.9%	1.4%	1.1%	0.2%	1.0%
Spain	5.2%	3.2%	4.6%	4.0%	0.6%	3.8%
Georgia	3.6%	1.6%	3.2%	2.3%	0.3%	2.4%
Cyprus	1.5%	0.5%	1.3%	1.0%	0.2%	0.9%
Average	2.9%	1.3%	2.5%	1.8%	0.3%	1.8%

Table A.3: Welfare Gains From Optimal Reallocation or Expansion of Current Networks, Mobile Labor

Table A.4: Optimal Infrastructure Investment, Population Growth and Local Characteristics for Different Number of Sectors

	(1)	(2)	(3)	(4)	(5)	(6)
	Investment	Investment	Investment	Pop. Growth	Pop. Growth	Pop. Growth
Population	0.114***	0.125***	0.125***	-0.000	-0.000	-0.000
Tradeable Income per Capita	0.146**	0.071	0.083	0.013*	0.008	0.011
Infrastructure	-0.214***	-0.235***	-0.236***	-0.001	-0.001	-0.001
Consumption per Capita				-0.064***	-0.060***	-0.065***
Infrastructure Growth				0.000	0.002*	0.002*
Differentiated Producer				0.009***	0.010***	0.010***
Observations	868	868	868	868	868	868
Adjusted R-squared	0.22	0.24	0.25	0.56	0.54	0.57

Each column corresponds to a different regression pooling all locations in the optimal expansion counterfactual across the 24 countries in the convex case with mobile labor and $\delta = \delta^{I,GEO}$. The dependent variable is investment in columns (1)-(3) and population growth in columns (4)-(6). All regressions include country fixed effects. Standard errors are clustered at the country level. ***=1% significance, **=5%, *=10%. Dependent variables: population growth is defined as $\Delta \ln L_j$, where Δx denotes the difference between variable x in the counterfactual and in the calibrated allocation. Investment growth is defined as the difference over the average, $\Delta I_j / (\frac{1}{2}(I_j + \Delta I_j))$, where total infrastructure at the node level defined as $I_j = \sum_{k \in \mathcal{N}(j)} I_{jk}$. Independent variables correspond to the log of the level of each variable in the calibrated model. Population and income per capita are the two outcomes matched by the calibration. Consumption per capita corresponds to traded goods, c_j in the calibrated model. Differentiated producer is a dummy for whether the location is a producer of differentiated goods in the calibration.

Table A.5: Average Welfare Gains for Different Number and Allocation of Differentiated Goods

Allocation of goods	Benchmark				Within NUTS			
	N=10		N=15		N=10		N=15	
Labor:	Fixed	Mobile	Fixed	Mobile	Fixed	Mobile	Fixed	Mobile
Optimal Reallocation								
$\delta = \delta^{I,GEO}$	1.7%	1.8%	1.7%	1.8%	1.8%	1.9%	2.2%	2.2%
Optimal Expansion								
$\delta = \delta^{I,GEO}$	1.7%	1.8%	1.8%	1.8%	1.9%	2.0%	2.3%	2.3%
$\delta = \delta^{I,FOC}$	0.3%	0.3%	0.3%	0.3%	0.3%	0.3%	0.4%	0.4%

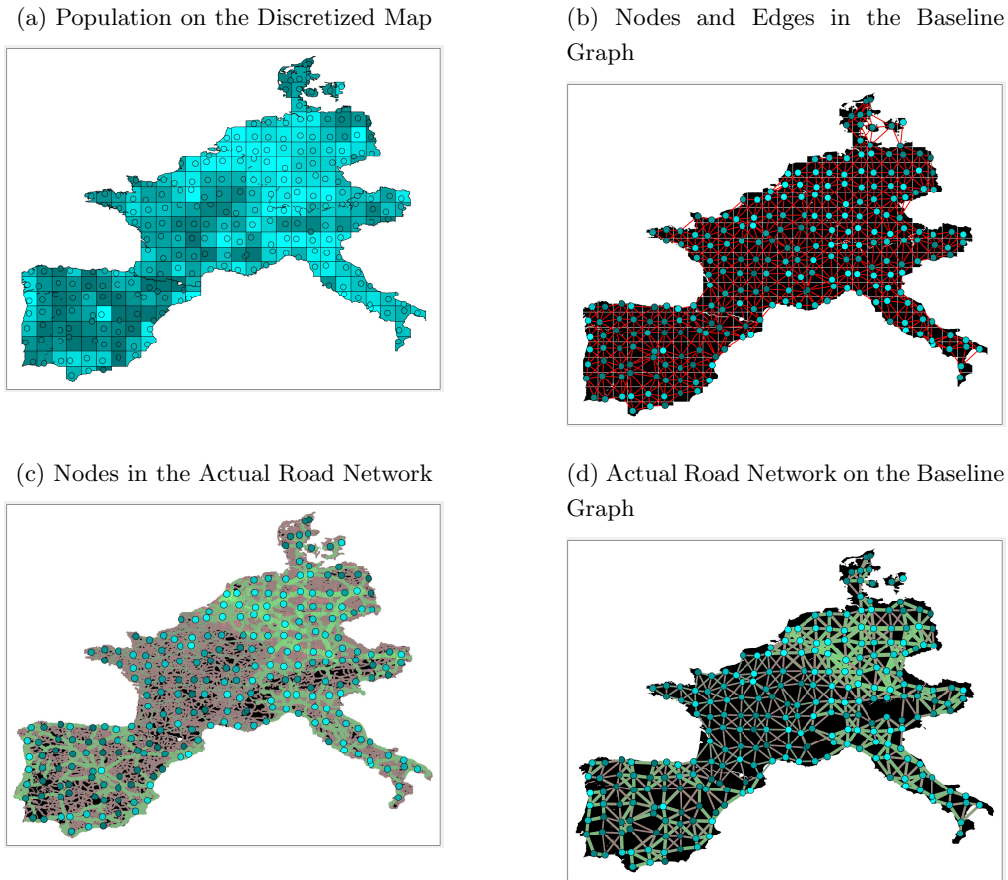
Each element of the table shows the average welfare gain in the corresponding counterfactual across the 24 countries for the convex case.

Table A.6: Average Welfare Gains with and without Congestion across Goods

Congestion	Across Goods		Own Good (Iceberg)	
	Fixed	Mobile	Fixed	Mobile
Labor:				
Optimal Reallocation				
$\delta = \delta^{I,GEO}$	1.7%	1.8%	1.5%	1.4%
Optimal Expansion				
$\delta = \delta^{I,GEO}$	1.7%	1.8%	1.6%	1.5%
$\delta = \delta^{I,FOC}$	0.3%	0.3%	0.2%	0.2%

Each element of the table shows the average welfare gain in the corresponding counterfactual across the 24 countries for the convex case with N=10 goods.

Figure A.2: Discretization of the European Road Network

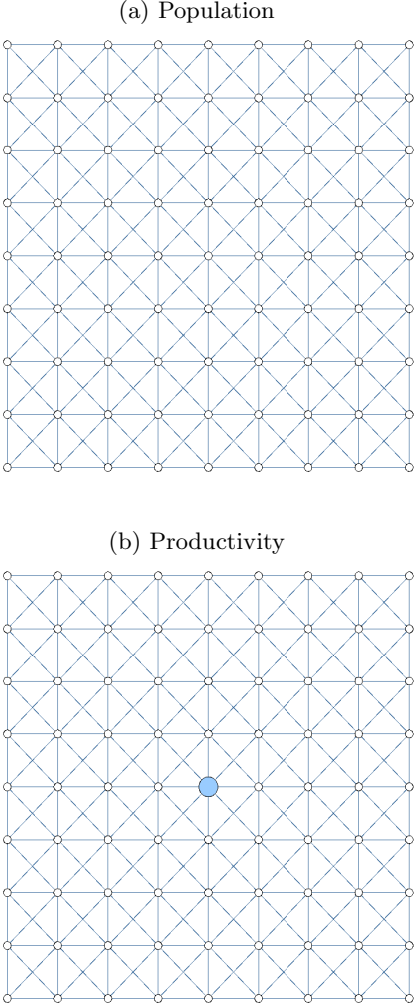


Notes: Panel (a) shows total population from GPW aggregated into 1 arc-degree (approximately 100 km) cells. Panel (b) shows the nodes \mathcal{J} corresponding to the population centroids of each cell in Panel (a), reallocated to their closest point on the actual road network, and the edges \mathcal{E} corresponding to all the vertical and diagonal links between cells. Panel (c) shows the centroids and the actual road network. Green segments correspond to national roads, red segments are all other roads, and the width of each segment is proportional to the number of lanes. Panel (d) shows the same centroids and the edges as the baseline graph in Panel (b), where each edge is weighted proportionally to the average number lanes on the cheapest path between each pair of nodes on the road network. The color shade ranges from red to green according to the fraction of the shortest path traveled on a national road.

Supplementary Material
(For authors' websites)

C Appendix to Section 4 (Illustrative Examples)


Figure A.3: A Simple Underlying Geography




Notes: On panel (a), each circle represents a location. The links represent the underlying network, i.e., links upon which the transport network may be built. Population and housing are uniform across space, normalized to 1. On panel (b), the size of the circles represent the productivity of each location.

Figure A.4: The Optimal Network for $K = 1$ and $K = 100$

(a) $K=1$




(b) $K=100$




Notes: On each panel, the thickness and color of the segments reflects the level of infrastructure built on a given link. Thicker and darker colors represent more infrastructure. On the bottom panels, the heat map represents the level of prices and consumption, normalized to 1 at the center. Lighter color represents higher values for prices and consumption. Prices and consumption levels are linearly interpolated across space to obtain smooth contour plots.

Figure A.5: Optimal Network with Randomly Located Cities


(a) Convex Case: $\gamma = \beta = 1$



(b) Non-Convex Case: $\gamma = 2 > \beta = 1$

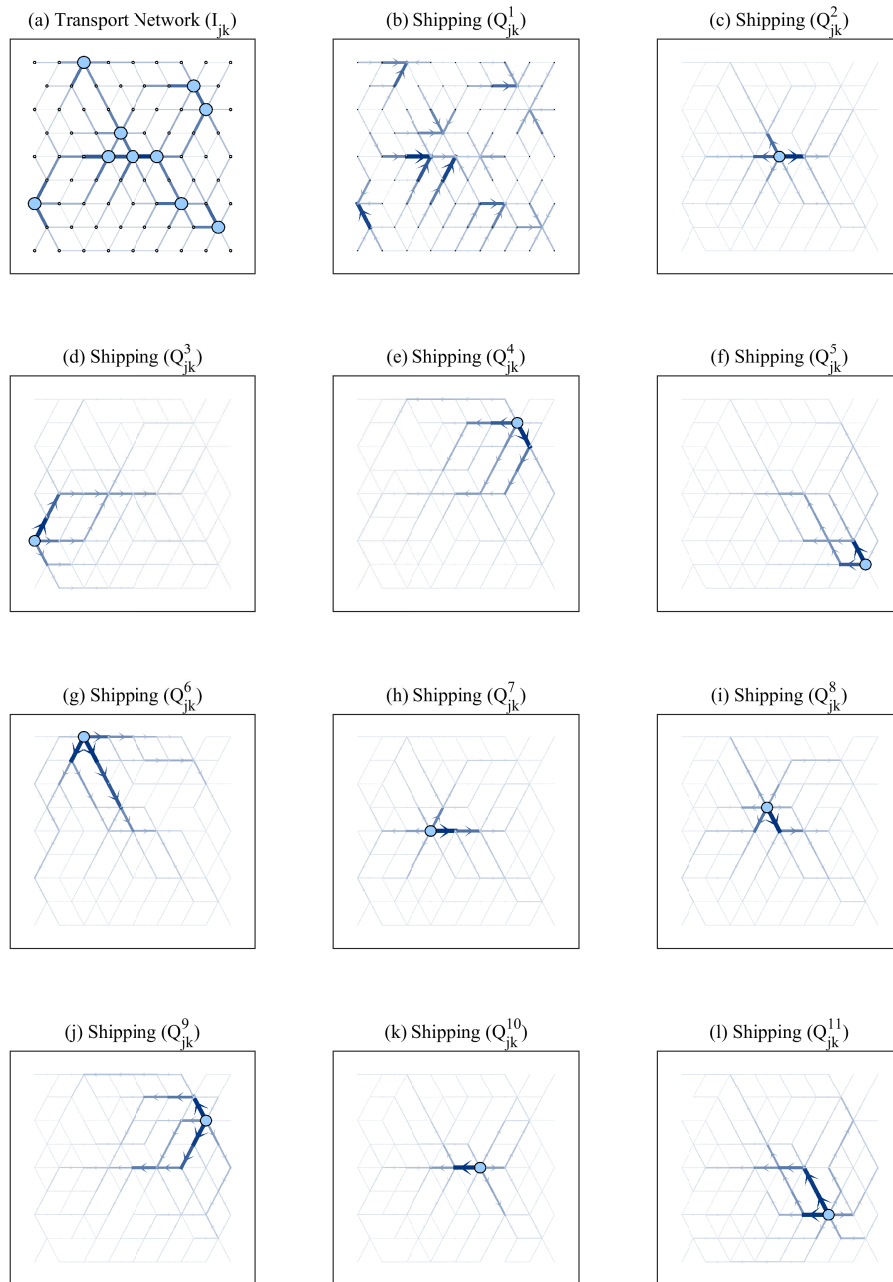


(c) Optimal Network Before and After Annealing Refinement in Non-Convex Case



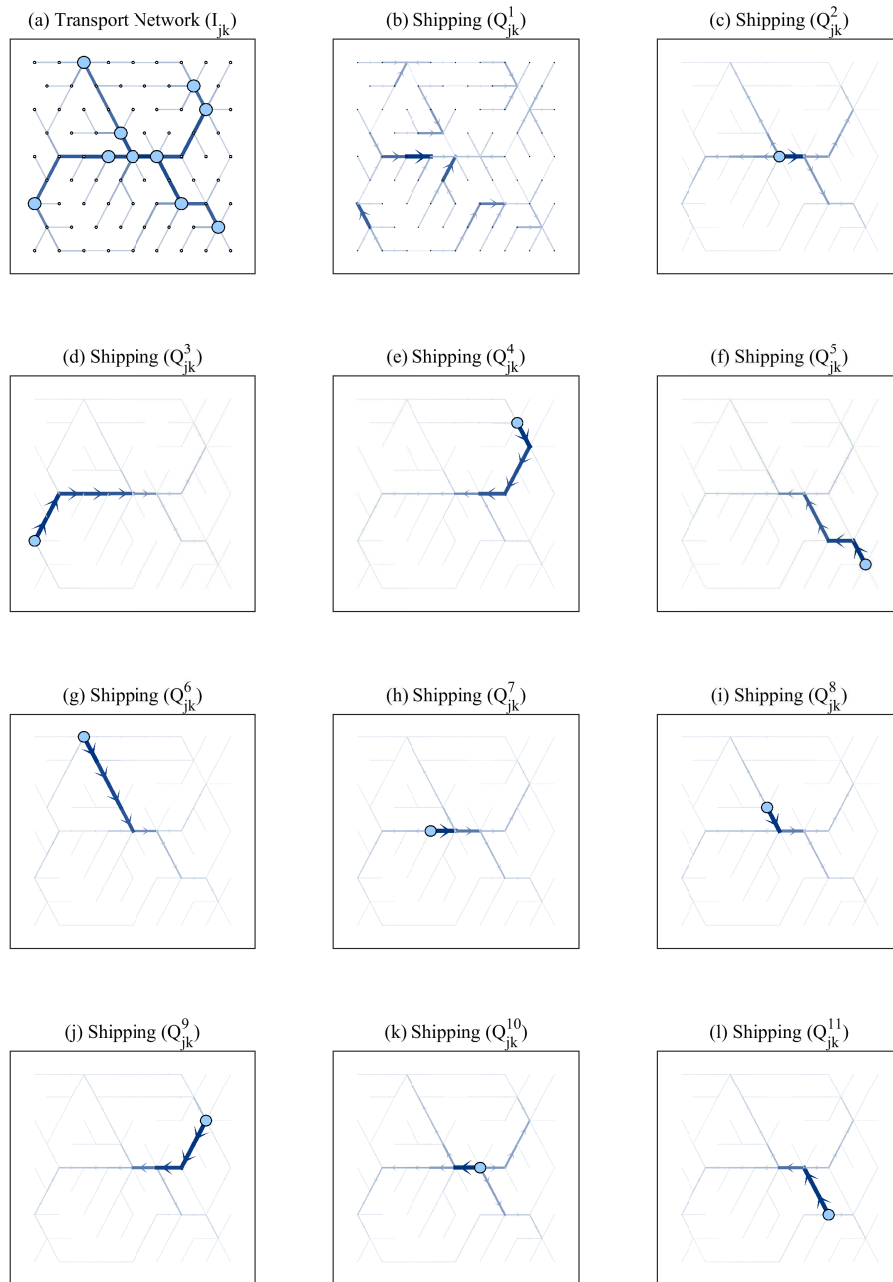
Notes: On each panel, the thickness and color of the segments reflects the level of infrastructure built or the shipment sent on a given link. Thicker and darker colors represent higher infrastructure or quantity.

Figure A.6: Optimal Network with 10+1 Goods, Convex Case ($\beta = \gamma = 1$), Labor Mobility



Notes: On panel (a), the thickness and color of the segments reflects the level of infrastructure built on a given link, and the size of each circle is the population share. On the other panels, the segments represent the quantity shipped through each link and the circles represent the location of producers.


Figure A.7: Optimal Network with 10+1 Goods, Nonconvex Case ($\beta = 1, \gamma = 2$), Labor Mobility




Notes: On panel (a), the thickness and color of the segments reflects the level of infrastructure built on a given link, and the size of each circle is the population share. On the other panels, the segments represent the quantity shipped through each link and the circles represent the location of producers. This figure represents a local optimum.

Figure A.8: The Optimal Transport Network under Alternative Building Costs


(a) Baseline Geography




(b) Adding a Mountain




(c) Adding a River and a Bottleneck Access by Land




(d) Allowing for Endogenous Bridges



(e) Allowing for Water Transport



(f) Non-Convex Case ($\gamma = 2$; $\beta = 1$) with Annealing



Notes: The thickness and color of the segments reflect the level of infrastructure built on a given link. Thicker and darker colors represent more infrastructure and quantities. The circles represent the 20 cities randomly allocated across spaces. The larger red circle represents the city with the highest productivity. The different panels vary in the parametrization of the cost of building infrastructure. In panel (a), it is only a function of Euclidean distance. In panel (b), we add a mountain and assume that the cost also depend on difference in elevation. In panel (c), we add a river with a natural land crossing and assume that the cost of building along or across the river is infinite. In panel (d) there is no natural land crossing but allow for construction of bridges. In panel (e) we additionally allow for investment in water transport. Panel (d) makes the assumptions as Panel (e) but assumes increasing returns to network building.

D Numerical Implementation

In this section, we provide a more detailed explanation of the numerical algorithms we use to solve the model. A Matlab toolbox implementing our model with detailed documentation and a few examples is available on the authors' websites.⁴⁶

D.1 Resolution method

Convex case and duality approach

As explained in section 3.6, our preferred approach to solve the model relies on solving the dual Lagrangian problem of the planner. We provide, here, a simple example of how to solve the joint optimal allocation and transport problem taking the infrastructure network $\{I_{jk}\}$ as given. This example can easily be generalized to the full problem, including the network design problem, in the convex case, but is also part of our resolution method for the nonconvex case. We focus on the case studied in the quantitative part of the paper, in which: i) we use the log-linear specification of transport costs, $\tau_{jk}^n = \delta_{jk}^\tau (Q_{jk}^n)^\beta I_{jk}^{-\gamma}$ (own-good congestion, $\chi = 0$) or $\tau_{jk}^n = m^n \delta_{jk}^\tau (Q_{jk}^n)^\beta I_{jk}^{-\gamma}$ with $Q_{jk} = \sum_{n=1}^N m^n Q_{jk}^n$ (cross-good congestion, $\chi = 1$); ii) labor is the sole production factor, $F_j^n (L_j^n) = z_j^n (L_j^n)^a$; and iii) C^T is a CES aggregator with elasticity of substitution σ . We consider the case with immobile labor.⁴⁷

We write the Lagrangian of the problem

$$\begin{aligned} \mathcal{L} = & \sum_j \omega_j L_j U(c_j, h_j) - \sum_j P_j^D \left[c_j L_j + \chi \sum_{k \in \mathcal{N}(j)} \delta_{jk}^\tau \left(\sum_{n=1}^N m^n Q_{jk}^n \right)^{1+\beta} I_{jk}^{-\gamma} - \left(\sum_n (D_j^n)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right] \\ & - \sum_j \sum_n P_j^n \left[D_j^n + \sum_{k \in \mathcal{N}(j)} \left(Q_{jk}^n + (1-\chi) \delta_{jk}^\tau (Q_{jk}^n)^{1+\beta} I_{jk}^{-\gamma} \right) - z_j^n (L_j^n)^a - \sum_{i \in \mathcal{N}(j)} Q_{ij}^n \right] \\ & - \sum_j W_j \left[\sum_n L_j^n - L_j \right] + \sum_{j,k,n} \zeta_{jkn}^Q Q_{jk}^n + \sum_{j,n} \zeta_{jn}^L L_j^n + \sum_{j,n} \zeta_{jn}^C D_j^n + \sum_j \zeta_j^c c_j. \end{aligned}$$

Recall that the dual problem consists of solving

$$\inf_{\lambda \geq \mathbf{0}} \sup_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}).$$

We start by expressing our control variables $\mathbf{x} = (c_j, D_j^n, Q_{jk}^n, L_j^n)$ as functions of the Lagrange multipliers $\boldsymbol{\lambda} = (P_j^D, P_j^n, W_j, \zeta_{jkn}^Q, \zeta_{jn}^L, \zeta_{jn}^C, \zeta_j^c)$. Using the optimality conditions, one obtains the following expressions:

$$\begin{aligned} c_j &= U_c^{-1} \left(\omega_j^{-1} \left(\sum_{n'} (P_{j'}^{n'})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, h_j \right) \\ D_j^n &= \left[\frac{P_j^n}{\left(\sum_{n'} (P_{j'}^{n'})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}} \right]^{-\sigma} \left(c_j L_j + \chi \sum_{k \in \mathcal{N}(j)} \delta_{jk}^\tau \left(\sum_{n=1}^N m^n Q_{jk}^n \right)^{1+\beta} I_{jk}^{-\gamma} \right) \\ L_j^n &= \frac{(P_j^n z_j^n)^{\frac{1}{1-a}}}{\sum_{n'} (P_{j'}^{n'} z_{j'}^{n'})^{\frac{1}{1-a}}} L_j. \end{aligned}$$

⁴⁶Latest version available at <https://sites.google.com/site/edouardschaal/OptimalTransportNetworkToolbox.zip>.

⁴⁷In the mobile labor case, we can only show that the planner's problem is a quasi-convex optimization problem. Hence, a duality gap may exist. We therefore adopt a (slower) primal approach in that case.

In the case of own-good congestion ($\chi = 0$), the flows can be easily inverted,

$$Q_{jk}^n = \left[\frac{1}{1 + \beta} \frac{I_{jk}^\gamma}{\delta_{jk}^\tau} \max \left(\frac{P_k^n}{P_j^n} - 1, 0 \right) \right]^{\frac{1}{\beta}},$$

but in the case of cross-good congestion, one first needs to invert the aggregate flows $Q_{jk} = \sum_{n=1}^N m^n Q_{jk}^n$, which is sufficient to evaluate the Lagrangian at any point,⁴⁸

$$Q_{jk} = \max_n \left(\frac{P_k^n - P_j^n}{(1 + \beta) P_j^D m^n \delta_{jk}^\tau I_{jk}^{-\gamma}} \right)^{\frac{1}{\beta}}.$$

As these expressions illustrate, we have been able to eliminate a large number of the multipliers directly, so that the only remaining Lagrange multipliers are $\boldsymbol{\lambda} = (P_j^n)_{j,n}$. We may now compute the inner part of the saddle-point problem:⁴⁹

$$\begin{aligned} \mathcal{L}(\mathbf{x}(\boldsymbol{\lambda}), \boldsymbol{\lambda}) = & \sum_j \omega_j L_j U(c_j(\boldsymbol{\lambda}), h_j) - \sum P_j^D \left(c_j(\boldsymbol{\lambda}) L_j + \chi \sum_{k \in \mathcal{N}(j)} \delta_{jk}^\tau Q_{jk}^{1+\beta}(\boldsymbol{\lambda}) I_{jk}^{-\gamma} \right) \\ & - \sum_j \sum_n P_j^n \left[D_j^n(\boldsymbol{\lambda}) + \sum_{k \in \mathcal{N}(j)} \left(Q_{jk}^n(\boldsymbol{\lambda}) + (1 - \chi) \delta_{jk}^\tau (Q_{jk}^n(\boldsymbol{\lambda}))^{1+\beta} I_{jk}^{-\gamma} \right) - z_j^n (L_j^n(\boldsymbol{\lambda}))^a - \sum_{i \in \mathcal{N}(j)} Q_{ij}^n(\boldsymbol{\lambda}) \right]. \end{aligned}$$

The dual problem then consists of the simple unconstrained, convex minimization problem in $J \times N$ unknowns:⁵⁰

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

This problem can be readily fed into numerical optimization software. Faster convergence can be achieved by providing the software with an analytical gradient and hessian. Note that, as a direct implication of the envelope theorem, the gradient of the dual problem is simply the vector of constraints:

$$\nabla \mathcal{L}(\mathbf{x}(\boldsymbol{\lambda}), \boldsymbol{\lambda}) = - \begin{pmatrix} \vdots \\ C_j^n(\boldsymbol{\lambda}) + \sum_{k \in \mathcal{N}(j)} \left(Q_{jk}^n(\boldsymbol{\lambda}) + \delta_{jk}^\tau (Q_{jk}^n(\boldsymbol{\lambda}))^{1+\beta} I_{jk}^{-\gamma} \right) - z_j^n (L_j^n(\boldsymbol{\lambda}))^a - \sum_{i \in \mathcal{N}(j)} Q_{ij}^n(\boldsymbol{\lambda}) \\ \vdots \end{pmatrix}.$$

Nonconvex cases

When the conditions for convexity fail to obtain, the full planner's problem is not a convex optimization problem. It is, however, easy to find local optima by using the following iterative procedure. We then search for a global maximum using a simulated annealing method that we describe below.

Finding Local Optima Despite the failure of global convexity for the full planner's problem, the joint optimal allocation and transport problems, taking the network as given, is always convex as long as $\beta \geq 0$. We thus use our duality approach to solve for $(c_j, D_j^n, Q_{jk}^n, L_j^n)$ for a given level of infrastructure I_{jk} , and then iterate on the (necessary) first order conditions that characterize the optimal network. The procedure can be summarized in pseudo-code as follows.

⁴⁸The specific values Q_{jk}^n can be recovered at the end of the optimization by inverting the linear system (for given multipliers $\boldsymbol{\lambda}$) which corresponds to the balanced-flows constraints, the constraints $Q_{jk}(\boldsymbol{\lambda}) = \sum_{n=1}^N m^n Q_{jk}^n$ and the complementary slackness conditions, $\zeta_{jkn}^Q Q_{jk}^n = 0$.

⁴⁹Note that, due to complementary slackness, we can drop the constraints that correspond to all the Lagrange multipliers that we were able to solve by hand. As a result, only the balanced flows constraints remain.

⁵⁰Dual problems are always convex, by construction, even when the primal problem is not.

1. Let $l := 1$. Guess some initial level of infrastructure $\{I_{jk}^{(1)}\}$ that satisfies the network building constraint.
2. Given the network $\{I_{jk}^{(l)}\}$, solve for $(c_j, D_j^n, Q_{jk}^n, L_j^n)$ using a duality approach.
3. Given the flows Q_{jk}^n and the prices P_j^n , get a new guess $I_{jk}^{(l+1)} = \left[\frac{\gamma}{P_K} \frac{\delta_{jk}^\tau}{\delta_{jk}^I} \left(P_j^D Q_{jk}^{1+\beta} \right) \right]^{\frac{1}{1+\gamma}}$ for $\chi = 1$ (or $I_{jk}^{(l+1)} = \left[\frac{\gamma}{P_K} \frac{\delta_{jk}^\tau}{\delta_{jk}^I} \left(\sum_n P_j^n (Q_{jk}^n)^{1+\beta} \right) \right]^{\frac{1}{1+\gamma}}$ for $\chi = 0$) and set P_K such that $\sum \delta_{jk}^I I_{jk}^{(l+1)} = K$.
4. If $\sum_{j,k} \left| I_{jk}^{(l+1)} - I_{jk}^{(l)} \right| \leq \varepsilon$, then we have converged to a potential candidate for a local optimum. If not, set $l := l + 1$ and go back to (2).

Simulated Annealing In the absence of global convexity results, the above iterative procedure is likely to end up in a local extremum. Unfortunately, there exists to our knowledge few global optimization methods that would guarantee convergence to a global maximum in a reasonable amount of time.⁵¹ We opt for the simple but widely used heuristic method of simulated annealing, which is a very popular probabilistic method to search for the global optimum of high dimensional problems such as, for instance, the traveling salesman problem. Simulated annealing can be described as follows:

1. Let $l := 1$. Set the initial network $\{I_{jk}^{(1)}\}$ to a local optimum from the previous section and compute its welfare $v^{(1)}$. Set the initial “temperature” T of the system to some number.
2. Draw a new candidate network $\{\hat{I}_{jk}\}$ by perturbing $\{I_{jk}^{(l)}\}$ (see below). (Optional: deepen the network.) Compute the corresponding optimal allocation and transport $\{c_j, D_j^n, Q_{jk}^n, L_j^n\}$. Compute associated welfare \hat{v} .
3. Accept the new network, i.e., set $I_{jk}^{(l+1)} = \hat{v}$ and $v^{(l+1)} = \hat{v}$ with probability $\min \left[\exp \left(\left(\hat{v} - v^{(l)} \right) / T \right), 1 \right]$, if not keep the same network, $\{I_{jk}^{(l+1)}\} = \{I_{jk}^{(l)}\}$ and $v^{(l+1)} = v^{(l)}$.
4. Stop when $T < T_{min}$. Otherwise set $l := l + 1$ and $T := \rho_T T$ and return to (2),

where $\rho_T < 1$ controls the speed of convergence. Note that we allow to “deepen” the network in step (2), meaning that we additionally apply the iterative procedure from the previous section for a pre-specified number of iterations so that the candidate network $\{\hat{I}_{jk}\}$ is more likely to be a local optimum.

Drawing Candidate Networks The performance of the simulated annealing depends on how new candidate networks are drawn. Because of the complex network structure, purely random perturbations are likely to be rejected and the algorithm may easily fail to improve the initial network. We propose the two different perturbation methods that have consistently produced the best results throughout our simulations.

Algorithm #1 “Rebranching” This algorithm exploits the structure of the problem to make educated guesses for the candidate networks. The algorithm builds on the idea that, under increasing returns, a welfare improvement can be achieved by directly connecting locations to more central locations. Since a lower price level indicates that a location has higher relative availability of goods produced anywhere in the economy, we use the price level as a proxy for centrality. We thus construct candidate networks where random locations are better connected to their lowest-price neighbors. The algorithm can be described as follows:

1. Given an initial network $I_{jk}^{(l)}$, draw a random set of locations $I \subset \mathcal{J}$ for “random rebranching” or set $I = \mathcal{J}$ for “deterministic rebranching”.

⁵¹Techniques such as the branch-and-bound method are guaranteed to converge to the global optimum, but remain heavy to implement and computationally intensive.

2. For each $j \in I$, identify $m(j)$ as the neighbor with the lowest price index for the bundle of tradable goods, $m(j) = \operatorname{argmin}_{k \in \mathcal{N}(j)} P_k^D$, and $n(j)$ as the “parent” with the highest level of infrastructure, $n(j) = \operatorname{argmax}_{k \in \mathcal{N}(j) | P_k^n \leq P_j^n} I_{kj}$.
3. For each $j \in I$, define the candidate network $I_{jk}^{(l+1)}$ by switching the infrastructure levels of $m(j)$ and $n(j)$:

$$I_{kj}^{(l+1)} = \begin{cases} I_{m(j)j}^{(l)} & \text{if } k = n(j), j \in I \\ I_{n(j)j}^{(l)} & \text{if } k = m(j), j \in I \\ I_{kj}^{(l)} & \text{if } j \notin I \text{ or } (j \in I \text{ and } k \notin \{m(j), n(j)\}), \end{cases}$$

which, by construction, satisfies the network building constraint.

Algorithm #2 “Hybrid Alder-FS” This algorithm attempts to implement the spirit of the algorithm proposed in Alder (2019), while exploiting the continuity of infrastructure investments and differentiability of the problem. The algorithm can be described as follows:

1. Compute the welfare gradient with respect to infrastructure investments and some given price of asphalt P_K . Delete the 5% links that correspond to the lowest elements of the gradient.
2. Compute the new economic allocation given the new network.
3. Compute the new welfare gradient and add a predefined quantity of asphalt to the link that has the highest element in the welfare gradient among inactive links.
4. Rescale the network so that the total quantity of asphalt is used and make sure that the network is connected (otherwise add the most beneficial link).

D.2 Numerical Evaluation of the Nonconvex Algorithm

We evaluate our simulated annealing approach in the nonconvex case in a range of random economies for which one can compute or approximate the global optimum.

Brute Force Approach An important caveat when running this performance evaluation is that infrastructure investments are continuous, so that the space of feasible networks is infinite. There is, unfortunately, no readily available brute force algorithm that guarantees finding the global optimum in this case. We explore every discrete combination of bilateral links (on/off), which we then use as initial conditions to iterate on the first-order conditions that characterize optimal infrastructure investment in our model. This brute force algorithm can only be used for small economies with a limited number of locations, as the number of combinations $2^{n(n-1)/2}$ explodes rapidly.

Alder (2019) We also compare the performance of our algorithm to Alder (2019), who proposes another heuristic algorithm in a spatial economic model. Since the Alder algorithm is discrete in nature (locations are connected or not, and there is no intensive aspect to infrastructure investments), there is no obvious way to perform the comparison. We propose two approaches.

In the first approach, which we refer to as “Alder”, we implement a simple version of the Alder (2019) algorithm in our model. Specifically, we restrict infrastructure investments to be discrete (0 or I), where I is constant across links and chosen such that the asphalt budget constraint is satisfied with equality. Under this constraint, we implement every step of the Alder algorithm, that is: 1) start from the full network, 2) compute welfare by removing each link one by one, 3) remove the 5% least profitable links, 4) add the most profitable link, 5) make sure that the network is connected, and iterate over (2)-(5) until no further step improves welfare. As a final step, we use the outcome of the Alder algorithm as initial condition in the first-order condition iteration from our model, making sure that welfare

keeps improving along the way. As a result we obtain a continuous version of the optimal network that is comparable to ours.

A caveat in the first approach is that the first part of the algorithm is constrained by the discreteness of investments. We thus devise a new algorithm that combines the appealing aspects of Alder (2019) with our continuous-investment approach which relies on differentiation and on the simulated annealing to escape local optima. We refer to this second approach as “Hybrid Alder-FS”. To be more precise, we modify our simulated annealing algorithm described in the previous section so that the stage in which we perturb the network is done in the spirit of the Alder algorithm. Specifically, the network is perturbed in the following way: 1) compute the welfare gradient with respect to infrastructure investment, 2) remove the 5% links with the lowest welfare gradient, 3) recompute the welfare and welfare gradient in the network after removal, 4) add the link that has the highest welfare gradient among non-existing links, and 5) make sure that the network is connected. The rest of the algorithm is kept identical to ours. In particular, it iterates on the first-order conditions associated with infrastructure investment before accepting the perturbation. In that sense, the intensive dimension of infrastructure investment is preserved throughout.

Note finally that the original Alder (2019) algorithm is designed to maximize aggregate income net of the dollar value of infrastructure investments. In this comparison, we change the objective function so that it maximizes welfare instead, subject to a fixed asphalt budget.

Simulated annealing There are many ways one can perturb the network in each loop of the simulated annealing method. To explore the role of the perturbation stage, we experiment with various ways to perturb the network: 1) purely random perturbations, 2) deterministic rebranching, and 3) random rebranching. “Rebranching” refers to the algorithm that we described in the previous section. In the deterministic version, all nodes are selected for rebranching every loop (i.e., reconnected to their best neighbors), while they are picked at random in the random version. This comparison allows us to evaluate the importance of randomness in the perturbation process as a way to escape local optima.

Small random economies

To perform the algorithm comparison, we draw random geographies for a given number of replications N_{reps} . Specifically, for each geography, we draw n locations from a uniform distribution over $[0, 1] \times [0, 1]$. The underlying network is fully connected. We parametrize the transport cost parameters δ_{jk}^T and network building costs δ_{jk}^I to be the euclidean distance between each pair of locations. To ease comparison with Alder (2019), who uses a gravity model without labor mobility, we study an Armington version of our model in which each location produces its variety with fixed labor. The rest of the model is otherwise identical to the one described in section 3. In particular, there is congestion in the transport technology, that is $\beta > 0$ using technology (10). We focus exclusively on nonconvex cases with $\gamma > \beta$.⁵²

Table A.7 below reports the performance of the various numerical approaches for several numbers of locations and replications. Due to the computation time of the brute force algorithm, we must restrict ourselves to 6 locations at most.

Before starting the individual comparison of each algorithm, it should be noted that the welfare losses are overall quite small with numbers in the magnitude of 10^{-4} in consumption equivalent. This suggests that, all things considered, each of these algorithms perform quite well for small economies. Looking at their individual performance, we find that the simple Alder algorithm is the least effective method with a welfare loss of about 0.02-0.03% in consumption equivalent compared to the brute force algorithm in most cases. A potential reason is that the discreteness constraint heavily distorts the search for the optimum in an economy with continuous investments. Turning to the other algorithms, our findings suggest that the random version of our “rebranching” algorithm and the Hybrid Alder-FS algorithm get the best results with average welfare losses of about 0.004% or less. The Hybrid

⁵²In the convex case $\beta \geq \gamma$, a simple iteration over our model’s first-order conditions or descent along the gradient suffice to find the global optimum quickly, so the algorithm comparison is irrelevant.

Alder-FS algorithm does slightly better in our simulations with 5 locations, but the random rebranching algorithm has the edge in the 6 locations case. In terms of computation time, the Hybrid algorithm requires slightly more time than the FS algorithm due to an additional evaluation of the equilibrium during the updating process. The Alder algorithm is the fastest algorithm given the small amount of locations, but the difference quickly reverts when we increase their number, as we show next.

Large random economies

We extend our comparison to economies with more locations. Unfortunately, we can no longer use the brute force algorithm, but we can still compare the performance between the other algorithms. In this last exercise, we focus on the version of our model used for the European road exercise in Section 5. Namely, we consider $w \times h$ rectangular networks with Moore neighborhood (each location is connected to its immediate horizontal, vertical and diagonal neighbors). Since we study economies with a large number of locations, we limit the number of traded goods to 2. As we did in the European exercise, we assume that one of the goods is a homogeneous “agricultural” good that can be produced in any location. The other good is a differentiated variety that can only be produced in a number of locations N_{cities} drawn uniformly among all locations. The rest of the model is otherwise identical to the previous section.

The results of this exercise, in Table A.8 below, are reported in terms of the welfare gains with respect to the local optimum that results from iterating over our model’s first-order conditions without using annealing. Our findings suggest larger differences between the algorithms as spatial complexity increases. While the Hybrid Alder-FS algorithm seems to perform quite well overall (and definitely better than random perturbations), our “rebranching” algorithm, in its deterministic and random versions, yields the best results in all cases. The random version performs the best with welfare improvements ranging between 0.01% and 0.1%. The simple (discrete) Alder algorithm does not do as well, most likely for the reason mentioned before.

Table A.7: Comparison of the welfare losses for small economies

#locations	N_{reps}	Simulated annealing				
		Alder	Random	Rebranching (deterministic)	Rebranching (random)	Hybrid Alder-FS
4	100	-0.0360%	-0.0036%	-0.0036%	-0.0020%	-0.0020%
5	100	-0.1281%	-0.0046%	-0.0047%	-0.0043%	-0.0032%
6	10	-0.0216%	-0.0016%	-0.0019%	0.0004%	-0.0019%

Notes: This table reports welfare in percentage consumption equivalent for the optimal network resulting from each algorithm compared to the brute force algorithm. The parameters chosen for the simulation are $\beta = 1$, $\gamma = 2$, $K = 1$, $\sigma = 5$, uniform amenities, productivity and population, no cross-good congestion, Cobb-Douglas utility with share $\alpha = 0.5$ on traded goods.

Table A.8: Comparison of the welfare gains for large economies

#locations	N_{cities}	N_{reps}	Alder	Simulated annealing			Hybrid Alder-FS
				Random	Rebranching (deterministic)	Rebranching (random)	
25	10	100	-0.0272%	0.0000%	0.0122%	0.0153%	0.0035%
64	20	100	-0.0133%	0.0000%	0.0378%	0.0431%	0.0017%
64	40	100	-0.0068%	0.0003%	0.0057%	0.0061%	0.0012%
100	20	100	-0.0033%	0.0011%	0.0858%	0.0954%	0.0024%
100	40	100	-0.0062%	0.0015%	0.0259%	0.0303%	0.0018%
100	60	100	-0.0048%	0.0005%	0.0042%	0.0061%	0.0008%

Notes: This table reports the welfare gain in percentage consumption equivalent compared to the outcome of our approach in non-convex cases without simulated annealing, simply obtained by iterating over the first-order conditions of our model starting from a full network. The parameters chosen for the simulation are $\beta = 1$, $\gamma = 2$, $K = 1$, $\sigma = 5$, uniform amenities, productivity and population, no cross-good congestion, and Cobb-Douglas utility with share $\alpha = 0.5$ on traded goods.

E Intuition for Global Convexity Condition

We show that the infrastructure investment and shipping decisions can be viewed as the outcome of a game between a decentralized planning agency and a shipping company. Due to the complementarity between infrastructure and trade, an investment in infrastructure in a link leads to an increase in commodity flows. These additional flows in turn encourage further investments. The convexity condition on the transport technology determines whether this “race” is stable or divergent. When the transport technology is convex, each additional increase in investment or flow becomes smaller until convergence to a non-trivial interior point.

For clarity, we focus on the single good case. We focus on a link (o, d) from origin o to destination d . We first consider the problem of a shipping company who takes prices and infrastructure as given and fully internalizes the congestion externality. From section 3 that this is equivalent to having the shipping company not internalize the congestion externality while facing an appropriately defined Pigouvian tax. The problem faced by the shipping company is

$$\max_{Q \geq 0} P_d Q - P_o Q (1 + \tau(Q, I)),$$

where P_i is the price in location $i = o, d$. The first-order condition is

$$P_d - P_o - P_o \frac{\partial(Q\tau)}{\partial Q} \leq 0 \text{ with equality if } Q > 0. \quad (\text{A.9})$$

This first-order condition is identical to (8) from the full planning problem. Denote $Q^*(I)$ the optimal amount of commodity flow Q that solves equation (A.9).⁵³ Consider now the problem of a myopic decentralized agency. It is myopic in the sense that it takes the trade flow Q as given and does not internalize the fact that Q responds to infrastructure in equilibrium, and decentralized in the sense that it solely invests in link (o, d) , taking as given the price P_K of asphalt (that is, the Lagrange multiplier of the asphalt budget constraint in the full planning problem) and the price of commodities. The objective of the planning agency is to minimize the sum of the market value of transport costs and the cost of asphalt. We write the problem as follows

$$\min_{I \geq 0} P_o Q \tau(Q, I) + P_K \delta^I I,$$

where δ^I is the marginal cost of investing infrastructure in the link. We obtain the following first-order condition,

$$P_o \frac{\partial(Q\tau)}{\partial I} + P_K \geq 0 \text{ with equality if } I > 0. \quad (\text{A.10})$$

This condition is identical to equation (9) in the main text. We denote $I^*(Q)$ the optimal infrastructure investment that solves equation (A.10).⁵⁴

Figure A.9 shows the two best response functions $Q^*(I)$ and $I^*(Q)$ for the convex and nonconvex cases using the functional form $\tau(Q, I) = \delta^\tau Q^\beta I^{-\gamma}$.⁵⁵ As the figure illustrates, there are two equilibria: the null equilibrium and an interior equilibrium with $I > 0$ and $Q > 0$. The reason why the game features two equilibria even in the convex case is that the planning agency is myopic and takes the commodity flow Q as given. In the full model and in the convex case, the social planner internalizes the response in Q and never chooses the null equilibrium under the Inada conditions.


The curve $Q^*(I)$ crosses $I^*(Q)$ from above in the convex case ($\beta \geq \gamma$) and from below otherwise. In the convex case, only the interior equilibrium is stable, as a small increase (resp. decrease) in infrastructure prompts an increase

⁵³Under the requirement that $Q\tau$ is convex in Q and satisfies some Inada conditions $\lim_{Q \rightarrow 0} \frac{\partial(Q\tau)}{\partial Q} = 0$ and $\lim_{Q \rightarrow \infty} \frac{\partial(Q\tau)}{\partial Q} = \infty$, the solution $Q^*(I)$ exists and is unique.

⁵⁴Under the assumption that $Q\tau$ is convex in I and satisfies the Inada conditions $\lim_{I \rightarrow 0} \frac{\partial(Q\tau)}{\partial I} = -\infty$ and $\lim_{I \rightarrow \infty} \frac{\partial(Q\tau)}{\partial I} = 0$, $I^*(Q)$ exists and is unique.

⁵⁵In this case we obtain: $Q^*(I) = \left(\frac{1}{1+\beta} \frac{1}{\delta^\tau} \frac{P_d - P_o}{P_o} \right)^{\frac{1}{\beta}} I^{\frac{\gamma}{\beta}}$ and $I^*(Q) = \left(\frac{\gamma}{P_K} \frac{\delta^\tau}{\delta^I} P_o \right)^{\frac{1}{1+\gamma}} Q^{\frac{\beta+1}{\gamma+1}}$.

Figure A.9: Nash equilibria of the infrastructure investment game



Note: Stable equilibria indicated by a black dot, unstable equilibria by a white dot.

(resp. decrease) in Q , but the response is limited due to the strong congestion forces. In turn, the moderate increase in flows justifies reducing (resp. increasing) the level of infrastructure back to the interior equilibrium. In the nonconvex case, the interior equilibrium is no longer stable but the null equilibrium is.

This discussion shows that the convexity property on the transport technology controls the complementarities between the planning agency and the private economy. In the nonconvex case, complementarities are strong. The convex case corresponds to weak complementarities: a deviation by one player is accompanied by a deviation by the other player in the same direction (they remain complements) but the response is too weak to push the economy away from the interior equilibrium. The interior equilibrium is stable when $Q^*(I)$ crosses $(I^*)^{-1}(I)$ from above, that is

$$\frac{\partial Q^*}{\partial I} \leq \left(\frac{\partial I^*}{\partial Q} \right)^{-1}. \quad (\text{A.11})$$

Differentiating the first-order conditions (A.9) and (A.10), we have $\frac{\partial Q^*}{\partial I} = -\frac{\partial^2(Q\tau)}{\partial Q \partial I} / \frac{\partial^2(Q\tau)}{(\partial Q)^2}$ and $\frac{\partial I^*}{\partial Q} = -\frac{\partial^2(Q\tau)}{\partial Q \partial I} / \frac{\partial^2(Q\tau)}{(\partial I)^2}$. The stability condition (A.11) is thus equivalent to

$$\frac{\frac{\partial^2(Q\tau)}{\partial Q \partial I}}{\frac{\partial^2(Q\tau)}{(\partial Q)^2}} \cdot \frac{\frac{\partial^2(Q\tau)}{\partial Q \partial I}}{\frac{\partial^2(Q\tau)}{(\partial I)^2}} \leq 1. \quad (\text{A.12})$$

In turn, (A.12) is equivalent to $\det(\text{Hess}(Q\tau)) \geq 0$, which along with the maintained assumption that $Q\tau$ is convex in each individual argument, $\frac{\partial^2(Q\tau)}{(\partial Q)^2} \geq 0$ and $\frac{\partial^2(Q\tau)}{(\partial I)^2} \geq 0$, is equivalent to requiring that the total transport costs $Q\tau(Q, I)$ is jointly convex over Q and I .

F Dealing with Inefficiencies

We show how to extend our approach to deal with cases with externalities in which Pigouvian taxes are not available or incorrectly set. We provide an example with partially internalized congestion in transport as well as externalities in production and amenities.

F.1 Environment

We modify the environment by introducing agglomeration externalities or spillovers in production

$$Y_j^n = F(L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n; \bar{L}_j)$$

where \bar{L}_j captures the production spillovers. \bar{L}_j is taken as given by firms and equals total employment in the location,

$$\bar{L}_j = \sum_n L_j^n.$$

We continue assuming that $F_j^n(\cdot)$ has constant returns to scale in its first three arguments, and is increasing and concave in each. Similarly, we allow for externalities in amenities by assuming that utility is given by

$$U(c_j, h_j; \bar{L}_j).$$

Finally, we assume that the congestion externalities in transport are only partially internalized by shipping companies. We focus on the convex case studied in our numerical exercise with mobile labor and cross-good congestion, but the analysis can be easily applied to the other cases. Specifically, we assume that the transport technology is given by

$$\tau_{jk}^n(Q_{jk}, I_{jk}; \bar{Q}_{jk}) = m^n \delta_{jk}^\tau (\bar{Q}_{jk})^{\beta - \tilde{\beta}} (Q_{jk})^{\tilde{\beta}} I_{jk}^{-\gamma}, \quad \beta \geq \gamma, \quad \beta \geq \tilde{\beta} \geq 0,$$

where $Q_{jk} = \sum_{n=1}^N m^n Q_{jk}^n$. In this specification, the total congestion externalities amount to a total exponent of β , but shipping companies only internalize them up to exponent $\tilde{\beta}$, taking the aggregate flow \bar{Q}_{jk} as given.⁵⁶ In equilibrium, the perceived flow on a given link \bar{Q}_{jk} is required to satisfy $\bar{Q}_{jk} = Q_{jk}$.

F.2 Fictitious planning approach

Solving for the economic allocation $\{c_j^n, L_j^n, \dots\}$ and trade flows $\{Q_{jk}^n\}$ given an infrastructure network $\{I_{jk}\}$ is an important step in the optimization. The method described in the main text involves solving for the allocation and flows jointly using a planning approach. Unfortunately, the presence of externalities does not allow us to invoke the First Welfare Theorem directly. We nonetheless propose an approach in which a ‘‘fictitious planner’’ does not internalize the effect of her decisions on \bar{L}_j and \bar{Q}_{jk} .

The fictitious planning problem for our economy can be characterized as the solution of a fixed-point problem. We first define the problem faced by the fictitious planner given $\{\bar{L}_j, \bar{Q}_{jk}\}$:

$$W\left(\{\bar{L}_j, \bar{Q}_{jk}\}_{j \in \mathcal{J}, k \in \mathcal{N}(j)}; \{I_{jk}\}\right) = \max_{\substack{u, c_j, h_j, L_j, D_j^n, \\ L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n, Q_{jk}, Q_{jk}^n}} u \quad (\text{A.13})$$

subject to (i) availability of traded commodities,

$$c_j L_j + \sum_{k \in \mathcal{N}(j)} \tau_{jk}(Q_{jk}, I_{jk}; \bar{Q}_{jk}) Q_{jk} \leq D_j \left(D_j^1, \dots, D_j^N\right) \text{ for all } j,$$

where $Q_{jk} = \sum_{n=1}^N m^n Q_{jk}^n$, and non-traded commodities,

$$h_j L_j \leq H_j \text{ for all } j;$$

⁵⁶This example is akin to having congestion taxes set to a fraction $\tilde{\beta}/\beta$ of the efficient Pigouvian taxes relative to unit transport costs.

(ii) the balanced-flows constraint,

$$D_j^n + \sum_{n'} X_j^{nn'} + \sum_{k \in \mathcal{N}(j)} Q_{jk}^n \leq F_j^n(L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n; \bar{L}_j) + \sum_{i \in \mathcal{N}(j)} Q_{ij}^n \text{ for all } j, n;$$

(iii) local factor market clearing,

$$\sum L_j^n \leq L_j \text{ for all } j;$$

and

$$\sum_n V_j^{mn} \leq V_j^n \text{ for all } j, m;$$

(iv) free labor mobility,

$$L_j u \leq L_j U(c_j, h_j; \bar{L}_j) \text{ for all } j;$$

(v) aggregate labor market clearing,

$$\sum_j L_j = L;$$

and (vi) non-negativity constraints on consumption, flows and factors.

For a given \bar{L}_j and \bar{Q}_{jk} , the above fictitious planning problem is a well-behaved convex problem under the same conditions as in the main text and can be efficiently solved using a convex solver. Imposing that $\bar{L}_j = L_j$ and $\bar{Q}_{jk} = Q_{jk}$ in equilibrium, the first-order conditions of this planning problem are identical to that of the competitive equilibrium. We are now ready to define a solution to the fictitious planning problem.

Definition 4. An allocation $\Omega = \{u, c_j, h_j, L_j, D_j^n, L_j^n, \mathbf{V}_j^n, \mathbf{X}_j^n, Q_{jk}, Q_{jk}^n\}_{j \in \mathcal{J}, k \in \mathcal{N}(j), 1 \leq n \leq N}$ is a solution to the fictitious planning problem of this economy for a given infrastructure network $\{I_{jk}\}$ if it is a solution of the problem (A.13),

$$\Omega \in \operatorname{argmax} W(\{\bar{L}_j, \bar{Q}_{jk}\}; \{I_{jk}\})$$

and satisfies

$$\begin{aligned} \bar{L}_j &= L_j \text{ for all } j, \\ \bar{Q}_{jk} &= Q_{jk} \text{ for all } j, k \in \mathcal{N}(j). \end{aligned}$$

A fictitious planning allocation is thus the solution of a fixed point problem over \bar{L}_j and \bar{Q}_{jk} , which can be solved using an iterative procedure.⁵⁷ This solution yields the competitive equilibrium of an economy with spillovers. In cases where uniqueness of the competitive equilibrium is guaranteed, there is an equivalence between the competitive equilibrium and the fictitious planning economy. In cases when uniqueness is not guaranteed, the fictitious planning approach remains a tractable way to compute candidate equilibria.

F.3 Infrastructure gradient

After solving for the economic allocation and trade flows given a certain infrastructure network $\{I_{jk}\}$, we are now ready to compute the welfare gradient with respect to infrastructure investments from the point of view of an overall planner who internalizes that the allocation corresponds to the solution to the fictitious planning problem and is therefore not generically efficient. To achieve this, we use the fact the competitive equilibrium is locally characterized as the solution to the fictitious planner's first-order conditions. We then differentiate these first-order conditions to evaluate the local impact of infrastructure changes.

We consider the case studied in our numerical exercise from section 5 with (i) labor as the only factor of production, (ii) D_j is a CES aggregator with elasticity of substitution σ , (iii) a unique good produced in each location (there is

⁵⁷In practice, we start with a guess on $\{\bar{L}_j, \bar{Q}_{jk}\}$, solve problem (A.13), update our guess and repeat until convergence. We do not have any proof guaranteeing the convergence of this procedure, but our experience has shown that it converges most of the time for reasonable externality parameters.

at most a unique $n(j)$ such that $z_j^{n(j)} > 0$ for each j) and the production function is given by $Y_j^n = z_j^n L_j^n \bar{L}_j^\varepsilon$ where ε governs the production externality and (iv) partially internalized transport externality ($0 \leq \tilde{\beta} < \beta$). In that case, the fictitious planning problem described in (A.13) simplifies to

$$\max_{\Omega=(u, c_j, h_j, L_j, D_j^n, Q_{jk}^n)} u$$

subject to

$$L_j u \leq L_j U(c_j, h_j; \bar{L}_j), \forall j \quad [\times \omega_j] \quad (\text{A.14})$$

$$c_j L_j + \sum_{k \in \mathcal{N}(j)} \delta_{jk}^\tau \bar{Q}_{jk}^{\beta-\beta} \left(\sum_n m^n Q_{jk}^n \right)^{1+\beta} I_{jk}^{-\gamma} \leq \left(\sum_n (D_j^n)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \forall j \quad [\times P_j^D] \quad (\text{A.15})$$

$$h_j L_j \leq H_j, \forall j \quad [\times P_j^H] \quad (\text{A.16})$$

$$D_j^n + \sum_{k \in \mathcal{N}(j)} Q_{jk}^n \leq z_j^n L_j \bar{L}_j^\varepsilon + \sum_{k \in \mathcal{N}(j)} Q_{kj}^n, \forall (j, n) \quad [\times P_j^n] \quad (\text{A.17})$$

$$\sum_j L_j = 1 \quad [\times W] \quad (\text{A.18})$$

$$Q_{jk}^n \geq 0, L_j \geq 0 \quad [\times \nu_{jk}^n, \xi_j]$$

where the variables in brackets denote the associated Lagrange multipliers. Constraints (A.14)-(A.18) will bind in equilibrium and define a system of $2 + 3J + J \times N$ equations that we summarize as

$$G_{cons}(\Omega; \bar{Q}_{jk}, \bar{L}_j, I_{jk}) = 0.$$

The first-order conditions of the above problem are

$$[u] \quad 1 - \sum_j \omega_j L_j = 0$$

$$[c_j] \quad \omega_j L_j U_c(c_j, h_j; \bar{L}_j) - P_j^D L_j = 0$$

$$[h_j] \quad \omega_j L_j U_h(c_j, h_j; \bar{L}_j) - P_j^H L_j = 0$$

$$[L_j] \quad -\omega_j (u - U(c_j, h_j)) - P_j^D c_j - P_j^H h_j + P_j^{n(j)} z_j^n \bar{L}_j^\varepsilon - W + \xi_j = 0$$

$$[D_j^n] \quad P_j^D D_j^{\frac{1}{\sigma}} (D_j^n)^{-\frac{1}{\sigma}} - P_j^n = 0$$

$$[Q_{jk}^n] \quad -P_j^D m^n (1 + \beta) \delta_{jk}^\tau \bar{Q}_{jk}^{\beta-\beta} Q_{jk}^\beta I_{jk}^{-\gamma} - P_j^n + P_k^n + \nu_{jk}^n = 0$$

along with the complementary slackness conditions

$$\nu_{jk}^n Q_{jk}^n = 0$$

$$\xi_j L_j = 0.$$

Together with the complementary slackness conditions, the first-order conditions define a system of $1 + 4J + J \times N + 4 \times nlinks \times N$ equations, which we denote by

$$G_{FOC}(\Omega; \bar{Q}_{jk}, \bar{L}_j, I_{jk}) = 0.$$

Imposing $\bar{Q}_{jk} = Q_{jk}$ and $\bar{L}_j = L_j$, the constraints and first-order conditions jointly define the system

$$G(\Omega; Q_{jk}, I_{jk}) = \begin{pmatrix} G_{cons}(\Omega; Q_{jk}, I_{jk}) \\ G_{FOC}(\Omega; Q_{jk}, I_{jk}) \end{pmatrix} = 0$$

of $2 + 7J + 2J \times N + 4 \times nlinks \times N$ equations in $2 + 4J + 2J \times N + 4 \times nlinks \times N$ unknowns (more precisely, the

allocation $u, c_j, h_j, L_j, D_j^n, Q_{jk}^n$ and the multipliers $\omega_j, P_j^D, P_j^H, P_j^n, W, \nu_{jk}^n, \xi_j$.

By differentiating our equilibrium conditions around a solution to the fictitious planning problem described above, we can now compute how each element of the allocation Ω is affected by a change in infrastructure. Using the implicit function theorem, we obtain

$$J_I(\Omega) = -[J_\Omega(G)]^{-1} J_I(G),$$

where $J_X(F)$ denotes the Jacobian of some function F with respect to some vector of variables X . Since welfare u is one of the variables in the allocation Ω , the gradient of welfare with respect to a change in infrastructure I_{jk} assuming a shadow value P_K for asphalt/concrete is simply given by

$$\frac{du}{dI_{jk}} - \delta_{jk}^I P_K.$$

With this welfare gradient, it is possible to optimize over the infrastructure network using a simple gradient-descent algorithm. There are, however, two new caveats in comparison to the efficient case covered in the main text: i) the computations can be quite time-consuming since the fictitious planning problem needs to be solved at every step of the optimization, ii) Proposition 1 on the convexity of the overall planning problem no longer applies and the resulting network has no guarantee to be the global optimum.

F.4 Numerical comparison

To quantitatively evaluate the importance of inefficiencies, we return to the benchmark case of Spain in our main quantitative exercise.⁵⁸ We put aside agglomeration externalities in production and consumption to focus on partially internalized congestion in transport as the only source of inefficiency. The model parameters are those estimated in the calibration and we simply vary the perceived congestion parameter $\tilde{\beta}$ while keeping the overall level of congestion β constant. We consider three cases: $\tilde{\beta} = 0.25\beta, 0.5\beta$ and 0.75β , which correspond to having the Pigouvian taxes set to 25%, 50% and 75% of what would be their optimal level relative to transport costs.

Table A.9 reports some summary statistics. Unsurprisingly, overall welfare resulting from the fictitious planner is lower compared to the efficient planner the less shipping companies internalize congestion (lower $\tilde{\beta}$). The difference is nonetheless quite small with welfare differences in the order of 10^{-4} in consumption equivalent terms. Similarly, we find that the resulting equilibrium trade flows Q_{jk} and infrastructure gradients are almost identical, with correlations close to 1 between the efficient and inefficient allocations.

$\tilde{\beta}$	Welfare	$\text{corr}(Q_{jk}^{\text{efficient}}, Q_{jk}^{\text{inefficient}})$	$\text{corr}\left(\frac{dW}{dI_{jk}}^{\text{efficient}}, \frac{dW}{dI_{jk}}^{\text{inefficient}}\right)$
Efficient	0	1	1
0.75β	-0.0057%	0.9999	0.9951
0.5β	-0.0235%	0.9998	0.9902
0.25β	-0.0545%	0.9996	0.9862

Notes: dW/dI_{jk} denotes the welfare gradient. Welfare is stated in percentage deviation from the efficient benchmark in consumption equivalent.

Table A.9: Statistical comparison between efficient and inefficient allocations

Since shipping companies fail to internalize their impact on congestion, we find that the inefficient allocation tends to feature larger trade flows the less the externality is internalized in comparison to the efficient one (Figure A.10). The left panel of Figure A.12 shows that the difference in flows are especially important around Madrid, Barcelona and Valencia. The discrepancies between the welfare gradients in infrastructure do not show a systematic pattern, as Figure A.11 illustrates. The right panel of Figure A.12 shows nonetheless a tendency for the planner to

⁵⁸More precisely, we consider the case $\beta = 0.13, \gamma = 0.10$ with labor mobility and cross-good congestion.

invest in infrastructure, in the inefficient case, in alternative routes around the most congested centers in order to alleviate congestion on the main axes. This is most easily seen in the cases of Madrid and Valencia.




Figure A.10: Comparison of trade flows Q_{jk} between efficient and inefficient cases




Figure A.11: Comparison of infrastructure gradients between efficient and inefficient cases



Notes: The differences are plotted in green when higher in the inefficient allocation and in red when higher in the efficient allocation. The figure displays the case $\tilde{\beta} = 0.5\beta$.

Figure A.12: Comparison of infrastructure gradients between efficient and inefficient cases

G Supplementary Material to Proposition 3

G.1 Definitions

Let $G = (\mathcal{I}, \mathcal{E})$ be an undirected graph. We say that a *path* of length $n \in \mathbb{N}^*$ from a node $a \in \mathcal{I}$ to $b \in \mathcal{I}$ is a finite sequence of nodes (i_1, \dots, i_n) such that $i_k \in \mathcal{I}$ for $1 \leq k \leq n$, $i_1 = a$ and $i_n = b$ and $\{i_k, i_{k+1}\} \in \mathcal{E}$. A *simple path* is a path that contains no repeated node, i.e., $i_k \neq i_l$ for all $1 \leq k, l \leq n$ and $k \neq l$. A *cycle* of length n is a path $p = (i_1, \dots, i_n)$ such that $i_1 = i_n$. A *simple cycle* of length n is a cycle that contains no repeated node other than the starting and ending nodes, i.e., $i_k \neq i_l$ for $1 \leq k, l \leq n-1$ and $k \neq l$. A *tree* is a connected graph that has no simple cycle. Equivalently, in a tree, there is a unique simple path connecting any two nodes.

G.2 Propositions and Lemmas

Proposition 3. (*Tree Property*) Assume that $\lim_{c \rightarrow 0^+} U_C(c, h) = \infty$. In the absence of a pre-existing network (i.e., $L_{jk} = 0$), if the transport technology is given by (10) and satisfies $\gamma > \beta$, and if there is a unique commodity produced in a single location, then the optimal transport network is a tree.

Proof. To establish the result, we focus on the case with fixed labor.⁵⁹ Note further that in the case with a unique traded commodity the own-good and cross-good congestion cases are identical. We assume WLOG that production Y_j is exogenous (endowment economy) since there is only one commodity and factors are immobile. To fix ideas, let us assume that $\mathcal{I} = \{0, 1, \dots, J-1\}$ and $Y_0 > 0$ but $Y_j = 0$ for $j \geq 1$. We write down the Lagrangian of the problem

$$\begin{aligned} \mathcal{L} = & \sum_j \omega_j L_j U(c_j, h_j) - \sum_j P_j \left[L_j c_j + \sum_{k \in N(j)} \left(1 + \delta_{jk}^r \frac{Q_{jk}^\beta}{I_{jk}^\gamma} \right) Q_{jk} - Y_j - \sum_{k \in N(j)} Q_{kj} \right] \\ & - P_K \left(\sum_j \sum_{k \in N(j)} \delta_{jk}^I I_{jk} - K \right) + \sum_{j,k} \zeta_{jk}^Q Q_{jk} + \sum_{j,k} \zeta_{jk}^I I_{jk}, \quad \zeta_{jk}^Q \geq 0, \zeta_{jk}^I \geq 0. \end{aligned}$$

Despite being a nonconvex optimization problem, there must exist a vector of Lagrange multipliers such that the KKT conditions hold.⁶⁰ As a preliminary step, we eliminate the infrastructure investment I_{jk} using (13), so that $I_{jk} = \left(\frac{\gamma}{P_K} \frac{\delta_{jk}^r}{\delta_{jk}^I} P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}}$ whenever $Q_{jk} > 0$, otherwise $I_{jk} = 0$. Solving for the multiplier P_K such that (7) is

⁵⁹In the labor mobility case, an identical argument can be made by taking the optimal allocation of L_j as given and replacing the Pareto weights ω_j with the Lagrange multipliers of the constraints $L_j u \leq L_j U(c_j, h_j)$.

⁶⁰The resource constraint can be substituted in the objective function to yield $P_j = \omega_j U_C(c_j, h_j)$. The Arrow-Hurwicz-Uzawa theorem (Takayama (1985), Theorem 1.D.4) implies that, the remaining constraints being affine, there must exist a vector of Lagrange multipliers such that the KKT conditions hold.

satisfied, we reformulate the problem with allocations and flows only as follows

$$\begin{aligned} \mathcal{L} = & \sum_j \omega_j L_j U(c_j, h_j) - \sum_j P_j \left[L_j c_j + \sum_{k \in \mathcal{N}(j)} Q_{jk} - Y_j - \sum_{k \in \mathcal{N}(j)} Q_{kj} \right] \\ & - K^{-\gamma} \left[\sum_{j,k} \left(\delta_{jk}^I / \delta_{jk}^\tau \right)^{\frac{\gamma}{\gamma+1}} \left(P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}} \right]^{\gamma+1} + \sum_{j,k} \zeta_{jk}^Q Q_{jk}, \quad \zeta_{jk}^Q \geq 0. \end{aligned} \quad (\text{A.19})$$

□

The source of nonconvexities is the term $\left[\sum_{j,k} \left(\delta_{jk}^I / \delta_{jk}^\tau \right)^{\frac{\gamma}{\gamma+1}} \left(P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}} \right]^{\gamma+1}$, which is convex when $\beta \geq \gamma$, but neither convex nor concave when $\gamma > \beta$. Let us now assume that $(\mathbf{c}^*, \mathbf{Q}^*)$ with $\mathbf{c}^* = (c_0^*, \dots, c_{j-1}^*)'$ and $\mathbf{Q}^* = (Q_{jk}^*)_{j,k \in \mathcal{N}(j)}$ is a local optimum, i.e., it satisfies the FOCs and SOC of the Lagrangian (A.19). We are going to show that the graph associated with \mathbf{Q}^* is a tree. Define the (undirected) graph associated to \mathbf{Q}^* as the tuple $(\mathcal{I}, \mathcal{E}^*)$ such that $\mathcal{E}^* \subset \mathcal{E}$ is a subset of the edges of the underlying geography such that

$$\mathcal{E}^* = \{ \{j, k\} \in \mathcal{E} \mid Q_{jk}^* > 0 \}.$$

Note that since I_{jk} is non-zero whenever $Q_{jk} > 0$ or $Q_{kj} > 0$, the support of graph $(\mathcal{I}, \mathcal{E}^*)$ coincides with that of the transport network $\{I_{jk}\}$. After this preparatory work, we now refer the reader to the Proposition 6 that follows.

Proposition 6. *\mathcal{E}^* is a tree.*

Proof. Because node 0 is the unique productive center and there is an Inada condition in consumption, there must exist a path connecting each node to 0. Hence, \mathcal{E}^* must be connected. It remains to show that \mathcal{E}^* cannot have simple cycles. We proceed by contradiction. Assume there exists a simple cycle $p = (i_1, \dots, i_n)$. Figure A.13 illustrates the different types of cycles that can arise. Case (i) is a cycle with circular flows that run in only one direction. Lemma 1 tells us that such a cycle cannot arise if $(\mathbf{c}^*, \mathbf{Q}^*)$ is locally optimal, as they inefficiently waste goods in transportation. Cases (ii) and (iii) correspond to cycles along which flows run into different directions. Lemma 3 establishes that whenever there is a cycle of type (iii), then there must exist a cycle of type (ii). We conclude with Lemma 4 by showing that cycles of type (ii) cannot arise if $(\mathbf{c}^*, \mathbf{Q}^*)$ is locally optimal. The reason is that one is better off redirecting flows into one of the two branches because of economies of scale in the transport technology when $\gamma > \beta$. Hence, simple cycles may not exist and \mathcal{E}^* is a tree.




Figure A.13: Different types of simple cycles

□

Lemma 1. *If $(\mathbf{c}^*, \mathbf{Q}^*)$ is a local optimum with $(\mathcal{I}, \mathcal{E}^*)$ its associated graph, then there exists no simple cycle $p = (i_1, \dots, i_n)$ such that $Q_{i_k, i_{k+1}}^* > 0$ for all $1 \leq k \leq n-1$.*

Proof. Case (i) in Figure A.13 presents the type of cycle with circular flows that cannot exist in a local optimum. By contradiction, assume that there exists such a cycle $p = (i_1, \dots, i_n)$. Then, for $\varepsilon > 0$ small, consider the allocation of flows

$$Q_{jk}^\varepsilon = \begin{cases} Q_{jk} - \varepsilon & \text{if } \exists l, 1 \leq l \leq n-1, \text{ such that } j = i_l \text{ and } k = i_{l+1} \\ Q_{jk}^* & \text{elsewhere} \end{cases}.$$

If $\varepsilon \leq \min_{1 \leq k \leq n-1} Q_{i_k, i_{k+1}}$, then $(\{c_j\}, \{Q_{jk}^\varepsilon\})$ is a feasible allocation that is strictly preferable to $(\{c_j\}, Q^*)$ since it yields the same utility at a lower transport cost. Hence, the gradient of the Lagrangian with respect to ε is strictly greater than 0 (recall that $P_j = u_c(c_j, h_j) > 0$), contradicting the assumption that (c^*, Q^*) is a local optimum. \square

Lemma 2. *For every node $a \in \mathcal{I}$ distinct from the productive center $0 \in \mathcal{I}$ and such that $L_a > 0$, there exists a simple path $p = (i_1, \dots, i_n)$ that connects 0 to a such that $Q_{i_k, i_{k+1}} > 0$ for $1 \leq k \leq n-1$.*

Proof. The proof is constructive. We build a simple path $p = (i_1, \dots, i_n)$ with $i_1 = a$, $i_k \neq i_l$ for all $1 \leq k, l \leq n$ and $k \neq l$ such that $Q_{i_k, i_{k-1}} > 0$. We proceed by recursion on the length of path p , which we denote by $|p| = n$. We start the recursion by setting $i_1 = a$. Because of the Inada conditions in the utility function and $L_a > 0$, we know that $c_a L_a > 0$. The balanced flow constraint in a ,

$$c_a L_a = \sum_{k \in \mathcal{N}(a)} Q_{ka} - \sum_{k \in \mathcal{N}(a)} Q_{ak} \left[1 + \delta_{ak}^\tau \frac{Q_{ak}^\beta}{I_{ak}^\gamma} \right] > 0,$$

tells us that location a must be a net recipient of goods from its neighbors. Hence, there exists $k \in \mathcal{N}(a)$ such that $Q_{ka} > 0$. Let $i_2 = k$. If $i_2 = 0$, then we have found a simple path connecting a to 0 with positive flows from 0 to a . If not, we now have a path $p_2 = (i_1, i_2)$ of length 2 such that $i_1 = a$, $i_1 \neq i_2 \neq 0$ and $Q_{i_2, i_1} > 0$. Assume now that $n \geq 2$ and, by recursion hypothesis, that we have a path $p_n = (i_1, \dots, i_n)$ with $i_1 = a$, $i_k \neq i_l \neq 0$ for all $1 \leq k, l \leq n$ and $k \neq l$ and such that $Q_{i_k, i_{k-1}} > 0$. Consider location i_n . The balanced flow constraint at i_n tells us that

$$c_{i_n} L_{i_n} = \sum_{k \in \mathcal{N}(i_n)} Q_{k, i_n} - \sum_{k \in \mathcal{N}(i_n)} Q_{i_n, k} \left[1 + \delta_{i_n, k}^\tau \frac{Q_{i_n, k}^\beta}{I_{i_n, k}^\gamma} \right] \geq 0.$$

Since we know by recursion hypothesis that $Q_{i_n, i_{n-1}} > 0$, then there exists a $k \in \mathcal{N}(i_n)$ such that $Q_{k, i_n} > 0$. We know that $k \neq i_l$ for all $1 \leq l \leq n$ because otherwise there would exist a cycle with circular flows, which is ruled out by Lemma 1. If $k = 0$, then we have found a path $p_{n+1} = (i_1, \dots, i_n, 0)$ that connects a to 0 with only positive flows from 0 to a . If not, then set $i_{n+1} = k$. We then have a path $p_{n+1} = (i_1, \dots, i_{n+1})$ with $i_1 = a$, $i_k \neq i_l \neq 0$ for all $1 \leq k, l \leq n+1$ and $k \neq l$ and such that $Q_{i_k, i_{k-1}} > 0$.

We conclude as follows. Since \mathcal{I} is finite, the above recursion must finish in a finite number of iterations. Since the recursion only stops after finding a path that ends in 0, then there must exist a simple path p of size $n < |\mathcal{I}|$ with $p = (i_1, \dots, i_n)$ such that $i_1 = a$, $i_n = 0$ and $Q_{i_k, i_{k+1}} > 0$ for $1 \leq k \leq n-1$. By construction, the path $\tilde{p} = (i_n, i_{n-1}, \dots, i_1)$ proves the statement. \square

Lemma 3. *Assume there exists a simple cycle $p = (i_1, \dots, i_n)$. Then, there exists $(a, b) \in \mathcal{I}^2$, $a \neq b$, such that there exists two distinct simple paths from a to b , $p_1 = (i_k^1)_{1 \leq k \leq n_1}$ and $p_2 = (i_k^2)_{1 \leq k \leq n_2}$ with $i_1^1 = i_1^2 = a$ and $i_{n_1}^1 = i_{n_2}^2 = b$, such that the flows are strictly positive from a to b along both paths, i.e., $Q_{i_k^l, i_{k+1}^l}^* > 0$ for $l \in \{1, 2\}$ and $1 \leq k \leq n_l - 1$.*

Proof. The objective of this lemma is to establish that if there exists a simple cycle, then there must exist a cycle of type (ii) as illustrated on Figure A.13.

Consider the simple cycle $p = (i_1, \dots, i_n)$. For convenience of notation, denote $i_0 = i_{n-1}$ and $i_{n+1} = i_2$. Denote $\tilde{Q}_{i_k, i_{k+1}} = Q_{i_k, i_{k+1}}^* - Q_{i_{k+1}, i_k}^*$ the net flow from i_k to i_{k+1} for $0 \leq k \leq n$, which can be either strictly positive or

strictly negative. We know from Lemma 1 that the net flows $\tilde{Q}_{i_k, i_{k+1}}$ cannot have the same sign, otherwise we would have a cycle with circular flow, violating the local optimality condition of (c^*, Q^*) . Hence, there must exist $1 \leq k \leq n$ such that $\tilde{Q}_{i_{k-1}, i_k} > 0$ and $\tilde{Q}_{i_k, i_{k+1}} < 0$. Node k is a location that receives goods from its two neighbors on the cycle, as illustrated by node c in case (iii) of Figure A.13. Set $a = 0$ and $b = i_k$. We know from Lemma 2 that there exists a path $p_1 = (j_1^1, \dots, j_{n_1}^1)$ such that $j_1^1 = a = 0$, $j_{n_1}^1 = i_{k-1}$ and $\tilde{Q}_{j_l^1, j_{l+1}^1} > 0$ for all $1 \leq l \leq n_1$. Similarly, there exists a path $p_2 = (j_1^2, \dots, j_{n_2}^2)$ such that $j_1^2 = a = 0$, $j_{n_2}^2 = i_{k+1}$ and $\tilde{Q}_{j_l^2, j_{l+1}^2} > 0$ for all $1 \leq l \leq n_2$. We now argue that the paths $\tilde{p}_1 = (j_1^1, \dots, j_{n_1}^1, b)$ and $\tilde{p}_2 = (j_1^2, \dots, j_{n_2}^2, b)$ are two distinct simple paths from a to b with strictly positive flows. By construction, we know that $\tilde{Q}_{j_{n_1}^1, i_k} > 0$ and $\tilde{Q}_{j_{n_2}^2, i_k} > 0$ so that the flows are strictly positive along both paths. We must only check that they are simple paths, i.e., that the nodes are not repeated. Let us treat the case of \tilde{p}_1 . The other one follows symmetrically. We must show in particular that there is no l with $1 \leq l \leq n_1$ such that $j_l^1 = b$. If this was the case, then $(b, j_{l+1}^1, \dots, j_{n_1}^1, b)$ would be a cycle with circular flows running in the same direction, which Lemma 1 rules out. Hence, \tilde{p}_1 is a simple path. \square

Lemma 4. For all $(a, b) \in \mathcal{I}^2$, $a \neq b$, if there are two simple paths p_1 and p_2 connecting a to b , i.e., $p_1 = (i_k^1)_{1 \leq k \leq n_1}$ and $p_2 = (i_k^2)_{1 \leq k \leq n_2}$ with $i_1^1 = i_1^2 = a$ and $i_{n_1}^1 = i_{n_2}^2 = b$, such that $Q_{i_k^l, i_{k+1}^l} > 0$ for $l \in \{1, 2\}$ and $1 \leq k \leq n_l$, then $p_1 = p_2$.

Proof. The objective of this lemma is to show that cycles of the type (ii) in Figure A.13 cannot exist. Assume by contradiction that such a cycle exists and that $p_1 \neq p_2$. Note that we can assume WLOG that $i_k^1 \neq i_l^2$ for all $1 < k < n_1$ and $1 < l < n_2$. To see this, let $\underline{k} = \min \{k | i_{k+1}^1 \neq i_{k+1}^2\}$ and $\bar{k}_1 = \min \{k > \underline{k} | \exists l > \underline{k}, i_k^1 = i_l^2\}$ and \bar{k}_2 be such that $i_{\bar{k}_1}^1 = i_{\bar{k}_2}^2$. By construction, the path $p'_1 = (i_{\underline{k}}^1, \dots, i_{\bar{k}_1}^1)$ and $p'_2 = (i_{\underline{k}}^2, \dots, i_{\bar{k}_2}^2)$ are two paths such that $i_{\underline{k}}^1 = i_{\underline{k}}^2$, $i_{\bar{k}_1}^1 = i_{\bar{k}_2}^2$, and $i_k^1 \neq i_l^2$ for all $\underline{k} < k < \bar{k}_1$ and $\underline{k} < l < \bar{k}_2$.




Figure A.14: Redirecting the flows to one branch

We are now going to show that $p_1 \neq p_2$ leads to a contradiction. The idea behind the proof is illustrated in Figure A.14. We are going to show that if there exists two distinct simple paths with positive flows going from a to b , then it would be strictly preferable to redirect the flows from one branch to the other due to the non-concavity of the Lagrangian, violating the local optimality of (c^*, Q^*) . Consider the allocation $Q^\varepsilon = \{Q_{jk}^\varepsilon\}$ for $\varepsilon \in \mathbb{R}$ such that

$$Q_{jk}^\varepsilon = \begin{cases} Q_{jk}^* + \varepsilon & \text{if } \exists l \text{ such that } j = i_l^1 \text{ and } k = i_{l+1}^1 \\ Q_{jk}^* - \varepsilon & \text{if } \exists l \text{ such that } j = i_l^2 \text{ and } k = i_{l+1}^2 \\ Q_{jk}^* & \text{elsewhere.} \end{cases}$$

In other words, Q_{jk}^ε corresponds to the pattern of flows Q_{jk}^* where a volume ε of the flows going through path 2 are redirected through path 1. By construction, Q_{jk}^ε is feasible (we are redirecting a fraction of flows that were running through locations on path 2 but not serving any of these locations). In particular, it leaves the value of the Lagrangian (A.19) unchanged except through the term $\left[\sum_{j,k} \hat{\delta}_{jk}^{\frac{\gamma}{\gamma+1}} \left(P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}} \right]^{\gamma+1}$ where $\hat{\delta}_{jk} = \delta_{jk}^I / \delta_{jk}^T$.

Consider the derivative of the Lagrangian with respect to ε :

$$\frac{\partial \mathcal{L}}{\partial \varepsilon} = -(1 + \beta) \left[\sum_{j,k} \hat{\delta}_{jk}^{\frac{\gamma}{\gamma+1}} \left(P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}} \right]^\gamma \left[\sum_{1 \leq k \leq n_1 - 1} \hat{\delta}_{i_k^1 i_{k+1}^1}^{\frac{\gamma}{\gamma+1}} P_{i_k^1}^{\frac{1}{\gamma+1}} Q_{i_k^1 i_{k+1}^1}^{\frac{1+\beta}{1+\gamma} - 1} - \sum_{1 \leq k \leq n_2 - 1} \hat{\delta}_{i_k^2 i_{k+1}^2}^{\frac{\gamma}{\gamma+1}} P_{i_k^2}^{\frac{1}{\gamma+1}} Q_{i_k^2 i_{k+1}^2}^{\frac{1+\beta}{1+\gamma} - 1} \right]$$

which satisfies $\frac{\partial \mathcal{L}}{\partial \varepsilon} = 0$ by assumption (local optimum). Let us examine the second order condition:

$$\begin{aligned} \frac{\partial^2 \mathcal{L}}{\partial \varepsilon^2} &= -(1 + \beta) \left(\frac{1 + \beta}{1 + \gamma} - 1 \right) \left[\sum_{j,k} \hat{\delta}_{jk}^{\frac{\gamma}{\gamma+1}} \left(P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}} \right]^\gamma \times \\ &\quad \left[\sum_{1 \leq k \leq n_1 - 1} \hat{\delta}_{i_k^1 i_{k+1}^1}^{\frac{\gamma}{\gamma+1}} P_{i_k^1}^{\frac{1}{\gamma+1}} Q_{i_k^1 i_{k+1}^1}^{\frac{1+\beta}{1+\gamma} - 2} + \sum_{1 \leq k \leq n_2 - 1} \hat{\delta}_{i_k^2 i_{k+1}^2}^{\frac{\gamma}{\gamma+1}} P_{i_k^2}^{\frac{1}{\gamma+1}} Q_{i_k^2 i_{k+1}^2}^{\frac{1+\beta}{1+\gamma} - 2} \right] \\ &\quad - (1 + \beta)^2 \frac{\gamma}{\gamma + 1} \left[\sum_{j,k} \hat{\delta}_{jk}^{\frac{\gamma}{\gamma+1}} \left(P_j Q_{jk}^{1+\beta} \right)^{\frac{1}{\gamma+1}} \right]^{\gamma-1} \times \\ &\quad \left[\underbrace{\sum_{1 \leq k \leq n_1 - 1} \hat{\delta}_{i_k^1 i_{k+1}^1}^{\frac{\gamma}{\gamma+1}} P_{i_k^1}^{\frac{1}{\gamma+1}} Q_{i_k^1 i_{k+1}^1}^{\frac{1+\beta}{1+\gamma} - 1} - \sum_{1 \leq k \leq n_2 - 1} \hat{\delta}_{i_k^2 i_{k+1}^2}^{\frac{\gamma}{\gamma+1}} P_{i_k^2}^{\frac{1}{\gamma+1}} Q_{i_k^2 i_{k+1}^2}^{\frac{1+\beta}{1+\gamma} - 1}}_{=0} \right]^2. \end{aligned}$$

Hence, we see that $\frac{\partial^2 \mathcal{L}}{\partial \varepsilon^2} \geq 0$ when $\gamma > \beta$. Therefore, the point under consideration cannot be a local maximum. A tiny deviation in either direction for ε would increase welfare, thereby yielding a contradiction. \square