

A Framework for Evaluating Border Configurations: Applications to Africa

Sasha Petrov[†]

December 13, 2024

Updated frequently: [\[Click here for the most recent version\]](#)

Abstract

National border configurations significantly affect welfare: they govern trading opportunities and eligibility for public services. Empirical evidence suggests that postcolonial border design has harmed African long-term development through these two channels. This paper offers a spatial model of borders that evaluates welfare consequences through trade and public goods provision. This model features four key forces: the benefits of economic and fiscal integration, the costs of preference heterogeneity and span of control. To evaluate the inefficiencies of border configurations, I set up an optimal borders problem to balance the trade-off between these forces and then develop a decomposition method to solve the problem. I calibrate the parameters of the spatial model and use the proposed decomposition method to solve the optimal borders problem in the African context. With optimal borders, Africa could gain at least 28% in welfare. The primary shortcoming of current borders is their geographic position, not the number of countries.

[†]The University of Chicago; email: alexpetrov@uchicago.edu. I am grateful for the advice and support of my advisors Jonathan I. Dingel, James Robinson, Stéphane Bonhomme, and Esteban Rossi-Hansberg. This project also benefited greatly from conversations with Thomas Bourany, Tomás Domingues-Iino, Omkar Katta, Olivier Kooi, Hugo Lhuillier, Hugo Lopez, Tom Hierons, Shanon Hsu, Jordan Rosenthal-Kay, Camilla Schneier, Jeanne Sorin, and from the feedback of participants of International Trade Working Group, Public/Labor Advising Group, Development Lunch Workshop, Computational Working Group, Student Applied Micro Lunch, Mansueto Urban Doctoral Fellows meetings.

1 Introduction

National border configurations have fundamental welfare implications. Border configurations determine the number and sizes of countries, as well as the socioeconomic features of their populations. These country characteristics significantly impact two economic outcomes: trade and the provision of public goods. Trading is easier within countries than across them, while within-country gains from trade depend on the degree of dissimilarity in technologies across subnational regions. National governments of bigger countries typically can produce more public goods, but a higher degree of population heterogeneity reduces the average enjoyment of the kind of the provided public good. Given the complex interactions of these various forces, how does the spatial distribution of welfare respond to changes in the border configuration? Which border configuration would be optimal in terms of aggregate welfare?

These questions have particular relevance to the African continent. A unique feature of the contemporary African political landscape is that the national borders are still inherited mainly from the European colonial area and were established without representation of African people. Empirical micro-level evidence suggests that these postcolonial national borders have negatively affected African long-term development (e.g., [Michalopoulos and Papaioannou 2016](#); [Thomas 2018](#)). The tension between country and ethnic borders is evident, for example, in the pervasiveness of territorial disputes. The involvement of international organizations also speaks to the urgency of border issues in Africa. For example, the International Court of Justice has received 13 cases concerning boundary disputes in Africa, and the African Union maintains a dedicated Border Programme ([Oduntan, 2015](#)). National borders are also an important factor of trade in Africa. [Lebrand \(2021\)](#) estimates that crossing certain country borders impedes trade on the African continent. Thus, considering the impact on trade is essential in analyzing alternative border configurations in Africa.

In this paper, I propose a framework for evaluating the welfare consequences of border configurations and the sub-optimality of any given border configuration in spatial general equilibrium. In this framework, national borders govern bilateral trade costs and socioeconomic compositions of countries. I set up and solve an optimal borders problem that searches over all possible partitions into countries. I apply this framework to African national borders and consider counterfactual border configurations better than the postcolonial border design.

The analysis starts with building a spatial general-equilibrium model that micro-founds the well-established key forces in border formation ([Alesina and Spolaore, 1997](#); [Alesina et al., 2005a](#)): economic benefits of market and fiscal integration and costs of large and heterogeneous countries. These forces are reflected through two components of the model, which characterize trade and the provision of public goods, respectively. The model's environment consists of a network of locations, a realistic geography, and an exogenous partition of locations into countries. Geography characterizes features of routes between all pairs of locations. Every

location has a representative worker and a continuum of firms producing traded private goods varieties. Every country has a national government that provides a public good.

Workers are characterized by the Cobb-Douglas utility over private and public goods and the preferences over horizontal types of public goods. The characterization of the private goods sector of the economy relies on the workhorse Ricardian trade model ([Eaton and Kortum, 2002](#)). Producers of varieties draw productivity fundamentals from the Fréchet distribution with the location-specific shifter and use linear technology with the labor input. Then, the produced varieties are traded between locations subject to trade costs determined by the availability of transportation infrastructure and the need to cross national borders. The CES aggregator characterizes workers' preferences over bundles of varieties.

I embed public goods production into the trade model in the following way. National governments collect taxes using the exogenous and uniform proportional income tax rate. Then, the governments endogenously choose the quantity and the horizontal type of the public good to provide. Technologically, the public good is the CES aggregator of inputs that are the same varieties as those used for private consumption. The governments can procure these inputs from any location under their jurisdiction. Thus, the government's problem has a spatial dimension, as different locations can offer the same variety at different prices. In this framework, increased sizes of countries lead to higher tax bases and more locations. Both of these effects increase the quantity of the public good that the government can provide. Governments also choose the horizontal type of the public good, facing the heterogeneity in preferences between the locations. The resulting provided type maximizes the aggregate utility from public good consumption in the country.

In addition, the public good that locations enjoy is subject to two exogenous factors: the degree of rivalry in consumption and the amount of leakage in delivery due to the span-of-control cost. The rivalry in consumption puts a welfare cost on the country's population size. Conditional on the quantity of the produced public good, a bigger population implies higher congestion. For example, if the schooling infrastructure is provided nationally, a sufficient increase in the number of children requires new school buildings to accommodate everyone in classrooms. The span-of-control cost punishes the average geographic remoteness of locations in a country. The further away a location is from the capital of its country, assumed to be at the centroid of the country, the less effective provision of the public good is to that location.

After characterizing the equilibrium outcomes under exogenous borders, the next step is to analyze what happens if the border configuration changes. Changes in border configurations affect welfare through a) trade costs between locations, b) geographic sizes of countries, c) tax revenues of national governments, and d) distributions of fundamentals (productivity shifters and preferences over types of public goods) within countries. For example, under suitable values of parameters and fundamentals, the integration of two countries increases aggregate gains from trade and tax revenues, but decreases the average utility from public good con-

sumption due to increases in the average remoteness of locations and the level of preference heterogeneity. Although this case scenario helps with intuition about the trade-off forces in the framework, there is a multitude of scenarios whose welfare consequences are ambiguous. For example, changes in country sizes are less important when one country's location is annexed by a neighboring country. The aggregate effect on welfare then crucially depends on the changes in the compositions of the affected countries.

Relying on the developed framework, I set up an optimal borders problem to search for a border configuration that maximizes the utilitarian welfare of the entire territory under consideration. A crucial methodological choice is to represent a partition of the geography into countries as a set of bilateral integration-status variables. This approach allows me to rely on well-developed discrete optimization approaches while maintaining the richness of the economic framework. In contrast to existing border formation models, the problem's objective function here is a micro-founded welfare function rather than a linear heuristic. I follow the Mathematical Programming with Equilibrium Constraints approach and add the equilibrium conditions as constraints to the optimization problem.

The practical feasibility of the set-up problem is delivered by warm-starting with a mixed-integer linear approximation, an instance of the canonical Max- k -Cut problem. Economically, this approach relates to the network-based representation of coalitional games ([Chalkiadakis et al., 2012](#)). Max- k -Cut programs are widely applied in settings where clustering of agents is pursued (e.g., communications network design; see [Saad et al. 2009](#) for a review), which has facilitated the continued development of fast solution algorithms for them. Given the network representation of the geography, a partition into jurisdictions is equivalent to a clustering of locations. This approach has seen applications in political science in the context of gerrymandering ([Validi et al., 2022](#); [Validi and Buchanan, 2022](#)). Although the set-up Max- k -Cut approximation in my case is not equivalent to the full optimal borders problem, it is guaranteed to deliver a feasible border configuration that is weakly welfare-improving compared to the point of approximation.

I apply the developed framework and optimization method to national borders in Africa. For the trade costs, I employ the same calibration strategy as in [Desmet et al. \(2018\)](#) (based on the methodology from [Allen and Arkolakis 2014](#)) by using granular geospatial information on available transportation modes and finding the fastest paths between all pairs of locations. I take the value of the border effect on trade costs from [Conte \(2022\)](#), which estimates it specifically in the African context. Given the calibrated trade costs, I invert the equilibrium conditions of the framework to calibrate the technology fundamentals of all spatial units. To quantify public goods preference heterogeneity, I use the typical approach in the literature on ethnic heterogeneity, which parametrizes individual utility loss for public goods as a function of the country's ethnic composition (e.g., the ethnolinguistic fractionalization index).

The computational results point to a counterfactual border configuration that yields an almost

30% welfare gain compared to the status quo. I use the calibrated values of fundamentals as inputs to the optimal borders problem to explore the counterfactual with a better border configuration. The comparison between the solution and the status quo reveals several important patterns. First, aggregate welfare gains are approximately 28% in the counterfactual, with the main margin of status quo sub-optimality being the geographic position of borders, not the number of countries. Second, the optimal border configuration leads to significant redistributive consequences consistent with partial welfare convergence, benefiting inland territories at the cost of richer coastal areas. Nevertheless, around 85% of the African population benefits from better borders.

Most importantly, this paper makes a novel contribution to the literature on the economics of jurisdictional borders (going as far back as [Friedman 1977](#)). While the seminal papers (for an overview, see [Alesina et al. 2005b](#)) have successfully exposed the main forces in border formation, this paper presents a quantitative framework that micro-founds those forces and allows for arbitrary geographies. Recent papers ([Allen 2023](#); [Weese 2015](#); [Gancia et al. 2022](#); [Fernández-Villaverde et al. 2023](#)) have also moved away from the theoretical stylized approach. However, their border characterizations still rely on notions that deliver crucial linearity in the objective function. For example, [Allen \(2023\)](#) characterizes country borders that minimize the sum of transportation costs from the locations to the capitals of countries. This paper builds upon the general equilibrium and welfare notions that retain the interactive nature of choices made across space at the cost of introducing non-linearities.

The computational results in this paper speak directly to the literature on the prominent interplay of national and ethnic borders in Africa. This empirical literature is limited recovering underlying mechanisms and considering counterfactuals in general equilibrium. The structural approach enables me to move away from the restriction of the partial equilibrium analysis and highlight the general equilibrium repercussions of border effects.

The spatial framework developed in this paper contributes twofold to the trade literature that studies border effects on the spatial distribution of economic outcomes. Existing papers that introduced public goods provision in a spatial framework (e.g., [Fajgelbaum et al. 2019](#); [Jannin and Sotura 2020](#)) impose equivalence between the spatial unit and the jurisdiction, while a key feature of this paper’s framework is that multiple spatial units can form one jurisdiction. [Fajgelbaum et al. \(2023\)](#) consider voting by agents residing in different locations, although the policy alternative is only binary. The typical focus on determining the equilibrium level of public spending is redirected in this paper to the problem of variability in *types* of public goods.

This paper’s focus on optimization extends the literature on the effects of economic unions by moving from the exogeneity of borders to introducing them into the choice set. In this sense, my approach finds a new balance between the analytical appeal of stylized theoretical models of border formation and the quantitative relevance of modern spatial models used in the welfare

analysis of trade and economic unions. [Caliendo and Parro \(2015\)](#) and [Caliendo et al. \(2021\)](#) develop and estimate a flexible spatial framework for evaluating the welfare effects of specific integration initiatives in NAFTA and the EU, respectively. Similarly to previously mentioned papers, they model locations and jurisdictions as identical units. The optimization challenge in this paper is in line with [Fajgelbaum and Schaal \(2020\)](#) that optimizes over possible road networks, with the link-level choice variables being continuous.

The rest of the paper is structured as follows. Section 2 develops the general equilibrium framework under exogenous borders that characterizes trade patterns and the provision of public goods. Section 3 sets up the optimal borders problem and explains the strategy to solve it computationally. Section 4 describes the calibration of the parameters and fundamentals of the framework. Section 5 applies the developed tools to the African context and discuss the results. The last section concludes with a discussion of a research agenda that this paper opens.

2 Framework with Exogenous Borders

This section develops a spatial quantitative framework that captures national border effects on welfare through trade and public goods provision. National borders impose additional trade costs for transporting goods and define for each location the government that is providing public goods. This framework enables welfare assessment of border configurations.

2.1 Environment




Figure 1: Illustration of agents in the model in a simple environment, in which there are 6 locations (circles) and 2 jurisdictions (bold boxes). Different arrows represent different types of flows, and dashed lines – trade routes through which goods can be transported.

There is a mass of workers L^ℓ in each location $\ell \in \mathcal{L}$, who are immobile and inelastically

supply labor to firms in their location. Countries partition the space into mutually exclusive subsets:

$$\mathcal{L} = \cup_y \mathcal{C}_y, \quad y \in \{1, \dots, \mathcal{Y}\}.$$

Each location ℓ belongs to some country $\mathcal{C}(\ell)$. Throughout this section, the assignment of locations to countries is fixed. Thus, the objective of this section is to deliver the indirect utility function that maps arbitrary border configurations to welfare levels of all locations. With an eye to considering alternative border configurations in the following sections, location outcomes whose values can change depending on the border configuration will have a country subscript $\mathcal{C}(\ell)$ besides a location superscript ℓ .

I model border configurations using binary variables $d = \{d_{\ell\ell'}\}_{\ell, \ell'}$ that represent whether locations ℓ and ℓ' are in the same country or not:

$$d_{\ell\ell'} = \mathbf{1} \{ \exists \mathcal{C} \text{ s.t. } \{\ell, \ell'\} \subseteq \mathcal{C} \}. \quad (1)$$




Figure 2: Illustration of how dyadic variables $d_{\ell\ell'}$ reflect border configurations.

2.2 Workers

Workers in location ℓ that is in country $\mathcal{C}(\ell)$ are characterized by the Cobb-Douglas utility over private consumption C and public goods $\mathcal{G}$:

$$U_{\mathcal{C}(\ell)}^\ell(C) = (\mathcal{G}_{\mathcal{C}(\ell)}^\ell)^\alpha C^{1-\alpha},$$

where $\mathcal{G}_{\mathcal{C}(\ell)}^\ell$ is the composite, location-specific value of the nationally provided public good; and C is the measure of the CES consumption bundle comprised of traded varieties $\nu \in [0, 1]$:

$$C^\ell = \left(\int_0^1 (c^\ell(\nu))^{\frac{\sigma-1}{\sigma}} d\nu \right)^{\frac{\sigma}{\sigma-1}}, \quad (2)$$

where σ is the elasticity of substitution between varieties. After paying the proportional income tax at rate t , households spend the remaining income on the consumption bundle whose unit

price is P^ℓ :

$$C^\ell = [1 - t] \frac{w^\ell}{P^\ell}. \quad (3)$$

The composite value of the national public good $\mathcal{G}_{\mathcal{C}(\ell)}^\ell$ has four components: the quantity of the produced public good and three location-specific factors that discount locations' utility from it. I model them in the following way:

$$\mathcal{G}_{\mathcal{C}(\ell)}^\ell = \frac{G_{\mathcal{C}(\ell)}}{(L_{\mathcal{C}(\ell)})^\xi} \times \delta_{\mathcal{C}(\ell)}^\ell \times \mu_{\mathcal{C}(\ell)}^\ell, \quad (4)$$

where $G_{\mathcal{C}(\ell)}$ is the quantity of public goods produced by the government, $L_{\mathcal{C}(\ell)}$ is the population size of country $\mathcal{C}(\ell)$, ξ is the degree of rivalry in the consumption of public goods, $\delta_{\mathcal{C}(\ell)}^\ell \in [0; 1]$ is the utility discount due to the preference heterogeneity over horizontally differentiated types of the public good, and $\mu_{\mathcal{C}(\ell)}^\ell \in [0; 1]$ represents the cost from span of control.

Firstly, the model allows for an arbitrary degree of rivalry in consumption of the public good, captured by the parameter $\xi \in [0, 1]$. Notice that when $\xi = 0$ everyone has same access to the provided quantity, while $\xi = 1$ amounts to per capita transfers of the total tax revenue. Thus, ξ is a key parameter determining the benefit from fiscal integration.

Secondly, the public good is characterized by horizontally differentiated types. Preferences of locations over types are represented in an indirect way following [Esteban and Ray \(2011\)](#): Each location prefers some type and does not like any other location's ideal type to a certain extent. This is captured in a cardinal way through location-pair fundamentals $\delta_{\ell\ell'} \in [0, 1]$:

$$\delta_{\mathcal{C}(\ell)}^\ell (a_{\mathcal{C}(\ell)}) = \left(\sum_{\ell' \in \mathcal{C}(\ell)} a_{\mathcal{C}(\ell)}^{\ell'} (\delta_{\ell\ell'})^{\zeta_0} \right)^{\zeta_1}, \quad \zeta_0, \zeta_1 \geq 0, \quad a_{\mathcal{C}(\ell)}^{\ell'} \in [0, 1], \quad \sum_{\ell' \in \mathcal{C}(\ell)} a_{\mathcal{C}(\ell)}^{\ell'} = 1. \quad (5)$$

As an assumption, $\delta_{\ell\ell} = 1$, implying that locations fully enjoy the provided public good if it is of their ideal type. In case locations do not gain any utility from public goods that are ideal for any other location, we would have $\delta_{\ell\ell'} = 0 \quad \forall \ell' \neq \ell$. In general, different locations are allowed to have same ideal types or share some similarity in them. The government can choose to provide any convex combination of ideal types of locations under its jurisdiction, reflected by $a_{\mathcal{C}(\ell)}^{\ell'}$. Parameters ζ_0, ζ_1 control the marginal effect of heterogeneity on the disutility from the provided public good. A useful benchmark is the case of $\zeta_0 = \zeta_1 = 1$ together with $a_{\mathcal{C}(\ell)}^{\ell'} = \frac{L^{\ell'}}{L_{\mathcal{C}(\ell)}}$ and $\delta_{\ell\ell'} = 0 \quad \forall \ell' \neq \ell$, because then, $\delta_{\mathcal{C}(\ell)}^\ell$ is directly related to location's population share in the country:

$$\delta_{\mathcal{C}(\ell)}^\ell = \frac{L^\ell}{L_{\mathcal{C}(\ell)}}. \quad (6)$$

Maintaining these assumptions on $a_{\mathcal{C}(\ell)}^{\ell'}$ and $\delta_{\ell\ell'}$, higher values of ζ_1 introduce decreasing marginal

effect when moving from perfect homogeneity to the minority status. Overall, this part of the public good utility is a key mechanism in how border configurations affect welfare not only through country sizes but also country compositions.

Lastly, the span-of-control cost $\mu_{\mathcal{C}(\ell)}^\ell$ represents potential leakage in the delivery of public goods to geographically remote locations of the country. Similarly to the heterogeneity cost, it is a function of bilateral characteristics, in this case distance measures $\gamma_{\ell\ell'}$:

$$\mu_{\mathcal{C}(\ell)}^\ell (b_{\mathcal{C}(\ell)}) = \left(\sum_{\ell' \in \mathcal{C}(\ell)} b_{\mathcal{C}(\ell)}^{\ell'} (\gamma_{\ell\ell'})^{\kappa_0} \right)^{-\kappa_1}, \quad \kappa_0, \kappa_1 \geq 0, \quad b_{\mathcal{C}(\ell)}^{\ell'} \in [0, 1], \quad \sum_{\ell' \in \mathcal{C}(\ell)} b_{\mathcal{C}(\ell)}^{\ell'} = 1. \quad (7)$$

The distance measure $\gamma_{\ell\ell'}$ reflects availability of transport infrastructure between ℓ to ℓ' (it will also be part of trade costs for transporting private goods). $b_{\mathcal{C}(\ell)}^{\ell'}$ can be interpreted as the position of the country's capital. For example, if the country's capital is set in ℓ' , then $b_{\mathcal{C}(\ell)}^{\ell'} = 1$ and the span-of-control cost for ℓ will be a function of only $\gamma_{\ell\ell'}$. In general, the capital's position can be more favorable to any of the locations in the country. Parameters κ_0, κ_1 control the marginal effect of geographic remoteness and, thus, are related to the scale of $\gamma_{\ell\ell'}$. When either $\kappa_0 = 0$ or $\kappa_1 = 0$, there is no effect to being more or less remote. Conditional on the value of κ_0 and the scale of $\gamma_{\ell\ell'}$, as κ_1 gets high enough, the effect becomes negligible and the utility from public goods is driven closer to 0. Conditional on the value of κ_1 , as κ_0 gets closer to 0, the effect of remoteness increases. Overall, the span-of-control cost directly penalizes the geographic area of a country.

2.3 Firms

There is a continuum of firms in every location that produce varieties $\nu \in [0, 1]$ using the Cobb-Douglas production function with labor $L^\ell(\nu)$ and intermediate input $Q^\ell(\nu)$:

$$q^\ell(\nu) = z^\ell(\nu) (L^\ell(\nu))^\rho (Q^\ell(\nu))^{1-\rho}, \quad \rho \in [0, 1],$$

where ρ is the share of labor in firms' costs, and $z(\nu)$ is a variety-specific productivity shifter that is drawn from an extreme value distribution:

$$z^\ell(\nu) \sim \text{Fréchet}(A^\ell, \theta).$$

Location-specific scale parameter A^ℓ reflects the level of productivity, and the dispersion parameter θ governs the variability of productivity shifters. The intermediate input $Q^\ell(\nu)$ is the same CES bundle of traded varieties as the consumption bundle (2):

$$Q^\ell(\nu) = \left(\int_0^1 (q^\ell(\nu'))^{\frac{\sigma-1}{\sigma}} d\nu' \right)^{\frac{\sigma}{\sigma-1}}.$$

Given Hicks-neutrality of the productivity shifter, all firms in ℓ face the same unit cost of the factor and input bundle:

$$c^\ell = (w^\ell)^\rho (P^\ell)^{1-\rho}.$$

Firms in ℓ can sell their output to any location ℓ' in the economy but they face the iceberg trade cost $\tau_{\ell\ell'}$. Therefore, a producer of variety ν in ℓ sets the following price for location j :

$$p_{\ell\ell'}(\nu) = \frac{c^\ell}{z^\ell(\nu)} \tau_{\ell\ell'}. \quad (8)$$

I model trade costs as depending on two components – the availability of transportation infrastructure $\gamma_{\ell\ell'}$ and the border-induced institutional factor:

$$\tau_{\ell\ell'} = \gamma_{\ell\ell'} (1 - \beta d_{\ell\ell'}). \quad (9)$$

If locations ℓ and ℓ' are in the same country, trade between them is not affected by having to cross customs, dealing with differences in legal frameworks and quality standards. This is captured by the discount factor β .¹ As a result, border configurations affect the spatial structure of trade costs by determining which locations can trade more easily with each other. Because gains from trade for locations increase if trade costs with relatively more productive trading partners decrease, this is another way in which border configurations affect welfare through country compositions.

Buyers of varieties in any location choose to import from locations that offer the lowest price. Following [Eaton and Kortum \(2002\)](#) which relies on properties of the Fréchet distribution, I obtain a closed-form characterization of the price index:

$$P^\ell = \gamma \sum_{\ell'} \left[A^{\ell'} \left(c^{\ell'} \tau_{\ell'\ell} \right)^{-\theta} \right]^{-1/\theta}, \quad (10)$$

where the scaling factor γ depends on the elasticity of substitution between varieties σ and the dispersion of productivities θ : $\gamma = \Gamma((\theta + 1 - \sigma)/\theta)^{1/(1-\sigma)}$.

Another useful consequence of the Fréchet distribution assumption is the closed form characterization of the share of expenditures by workers and firms in ℓ on imports from ℓ' , denoted $\pi_{\ell\ell'}$:

$$\pi_{\ell\ell'} = \frac{A^{\ell'} \left(c^{\ell'} \tau_{\ell'\ell} \right)^{-\theta}}{(\gamma^{-1} P^\ell)^{-\theta}}. \quad (11)$$

¹Notice that this formulation is equivalent to a more common parametrization of the border effect as a tariff-like premium on the trade cost τ_b : $\tau_{\ell\ell'} = \gamma_{\ell\ell'} (1 + \tau_b) \left(1 - \frac{\tau_b}{1 + \tau_b} d_{\ell\ell'} \right)$.

2.4 Governments

Each country has its own government that provides a public good to locations under its jurisdiction. Its objective is to maximize the utility of the citizens derived from the public good. The government makes two decisions: How much of the public good to provide and what horizontally differentiated type to assign to it. A useful running example to keep in mind is national provision of schools. To build schools, governments need to procure inputs such as cement and bricks. They also decide on the curricula that might disfavor certain population groups, for example, by making learning one language mandatory.

Geographic remoteness To focus on the consequences of preference heterogeneity on welfare effects of border configurations, I exclude the choice of the capital's locations from government's problem. Guided by the principle that capitals tend to be located at the population-weighted centroids of countries (Allen, 2023), I assume that locations experience leakage in the delivery of public good according to their population-weighted geographic remoteness in the country.

Assumption 1 *The capital of country $\mathcal{C}$ favors locations under its jurisdiction proportionally to their population share:*

$$b_{\mathcal{C}(\ell)}^\ell = \frac{L^\ell}{L_{\mathcal{C}(\ell)}}. \quad (12)$$

Effectively, this introduces a mechanical penalty on the geographic area of a country by imposing welfare losses for locations on the country's outskirts.

Choosing quantity The government of country $\mathcal{C}$ assembles the public good with a CES technology, using varieties procured from locations within its jurisdiction:

$$G_{\mathcal{C}} = \left(\int_0^1 \left(\sum_{\ell \in \mathcal{C}} g_\ell(v) \right)^{\frac{\sigma-1}{\sigma}} dv \right)^{\frac{\sigma}{\sigma-1}}. \quad (13)$$

Firms set prices to the government according to:

$$p_{\ell g}(\nu) = \frac{c^\ell}{z^\ell(\nu)}.$$

Similarly to price indices of locations (10), properties of the Fréchet distribution imply the following expression for the unit cost $P_{\mathcal{C}}$ of the input bundle for the government:

$$P_{\mathcal{C}} = \gamma \left[\sum_{\ell \in \mathcal{C}} A^{\ell} (c^{\ell})^{-\theta} \right]^{-1/\theta}. \quad (14)$$

The amount that the government can spend on procurement is determined by the tax revenues coming from the income tax collected at the exogenous and uniform-across-countries rate t :

$$P_{\mathcal{C}} G_{\mathcal{C}} \leq \sum_{\ell \in \mathcal{C}} t w^{\ell} L^{\ell}. \quad (15)$$

Choosing type The approach to model the effect of heterogeneous preferences on the type of provided public good follows in nature [Esteban and Ray \(2011\)](#). As previewed in the discussion of the heterogeneity cost (5), the government can choose any convex combination of the ideal types of locations under its jurisdiction. Formally, this amounts to choosing values of the coefficients $a_{\mathcal{C}} = \{a^{\ell}\}_{\ell \in \mathcal{C}}$ that lie in the unit simplex of dimension $|\mathcal{C}| - 1$.

$$\begin{aligned} \max_{G, a_{\mathcal{C}}} \quad & \sum_{\ell \in \mathcal{C}} \frac{L^{\ell}}{L_{\mathcal{C}(\ell)}} \left(\frac{G}{(L_{\mathcal{C}(\ell)})^{\xi}} \delta_{\mathcal{C}(\ell)}^{\ell}(a_{\mathcal{C}}) \right)^{\alpha} \\ \text{s.t.} \quad & (13), (14), (15), (5) \\ & \sum_{\ell \in \mathcal{C}} a^{\ell} = 1, \quad a^{\ell} \in [0, 1] \quad \forall \ell. \end{aligned}$$

Because the government is procuring varieties from the cheapest sources and due to properties of the Fréchet distribution, the share of the tax revenue that the government spends on procurement from location $\ell \in \mathcal{C}$ is:

$$\pi_{\ell G} = \frac{A^{\ell} (c^{\ell})^{-\theta}}{\sum_{\ell' \in \mathcal{C}(\ell)} A^{\ell'} (c^{\ell'})^{-\theta}}. \quad (16)$$

Proposition 1 *If $\zeta_0 = 1$, then the government of country $\mathcal{C}$ picks coefficients $\{\bar{a}_{\mathcal{C}}^{\ell}\}_{\ell \in \mathcal{C}}$ according to:*

$$\bar{a}_{\mathcal{C}(\ell)}^{\ell} = \sum_{\ell'} \left(\frac{L^{\ell'}}{L_{\mathcal{C}(\ell)}} \right)^{\frac{1}{\alpha \zeta_1}} \delta_{\ell \ell'} \bigg/ \sum_{\tilde{\ell}, \tilde{\ell}'} \left(\frac{L^{\tilde{\ell}}}{L_{\mathcal{C}(\tilde{\ell})}} \right)^{\frac{1}{\alpha \zeta_1}} \delta_{\tilde{\ell} \tilde{\ell}'} \quad \forall \ell \in \mathcal{C}. \quad (17)$$

Notice that in the simplified scenario of complete incompatibility between types of locations (i.e., $\delta_{\ell \ell'} = 0 \quad \forall \ell' \neq \ell$), the solution implies equivalence between location's disutility from heterogeneity and its population share in the country (as in 6).

2.5 Trade flows and procurement

A crucial feature of the framework is that the provision of public goods interacts with the equilibrium determination of wages and prices. This is because the sales of any location ℓ comprise not only exports to all other locations but also supplies to the national government:

$$X^\ell = \sum_{\ell'} \pi_{\ell\ell'} \left[\underbrace{(1-t)w^{\ell'}L^{\ell'}}_{\text{Private demand}} + \underbrace{\frac{1-\rho}{\rho}w^{\ell'}L^{\ell'}}_{\substack{\text{Firms' demand} \\ \text{for intermediate inputs}}} \right] + \underbrace{\pi_{\ell G}t \sum_{\ell' \in \mathcal{C}(\ell)} w^{\ell'}L^{\ell'}}_{\substack{\text{Government's} \\ \text{procurement demand}}} . \quad (18)$$

Importantly, tax payments of any location ℓ are not bound to match exactly the procurement payment from the government. As a result, trade flows are not necessarily balanced at the location level:

$$\text{Expenditures}^\ell - \text{Sales}^\ell = P^\ell C^\ell - \underbrace{\sum_{\ell'} \pi_{\ell\ell'} (P^{\ell'} C^{\ell'} + P^{\ell'} Q^{\ell'})}_{\text{Private market imbalance}} - \underbrace{X_{G_{\mathcal{C}(\ell)}}^\ell}_{\substack{\text{Procurement} \\ \text{order}}} . \quad (19)$$

2.6 General equilibrium

Definition 1 (General equilibrium) *Given a border configuration d , parameter values $\Theta = \{\alpha, \sigma, \rho, \xi, \zeta, \kappa, \theta, \beta, t\}$, values of fundamentals $\mathbf{F} = \{L, \gamma, \delta, A\}$, the equilibrium consists of outcomes $\mathbf{x} = \{C^\ell, G^\ell, P^\ell, P_{\mathcal{C}(\ell)}, w^\ell, \delta^\ell, \mu^\ell\}_\ell$ such that:*

- *workers spend their post-tax income on private consumption;*
- *firms set prices in all locations according to (8);*
- *price indices follow (10);*
- *government budget constraint (15) holds;*
- *government's unit cost of procurement is (14);*
- *government makes procurement decisions according to (16);*
- *utility discounts due to geographic remoteness are set according to (12);*
- *utility discounts due to preference heterogeneity are determined by (17);*
- *wages in every location equate aggregate labor demand in the location and exogenous labor supply according to (18).*

2.7 Welfare effects of border changes

To summarize the mechanisms through which changes in a given border configuration affect welfare of location ℓ , I decompose the log welfare change in (20). For this expression, an outcome is written as a function of a border configuration d if it is subject to general equilibrium forces involving the network of all locations (for example, wage $w^\ell(d)$). If an outcome depends only on the set of fundamentals of locations in the same country as ℓ , it is reflected in the country subscript $\mathcal{C}(\ell; d)$, which is conditioned on the border configuration. Notice that tax revenues and the unit cost of public good are affected by both of these factors: they depend on equilibrium wages and on the set of locations in the country.

$$\begin{aligned} \Delta \ln \mathcal{W}^\ell = & \alpha \left(\left[\ln \sum_{\ell' \in \mathcal{C}(\ell; d')} w^{\ell'}(d') - \ln \sum_{\ell' \in \mathcal{C}(\ell; d)} w^{\ell'}(d) \right] - [\ln P_{\mathcal{C}(\ell; d')}(d') - \ln P_{\mathcal{C}(\ell; d)}(d)] \right. \\ & \left. - \xi [\ln L_{\mathcal{C}(\ell; d')} - \ln L_{\mathcal{C}(\ell; d)}] + [\ln \delta_{\mathcal{C}(\ell; d')}^\ell - \ln \delta_{\mathcal{C}(\ell; d)}^\ell] + [\ln \mu_{\mathcal{C}(\ell; d')}^\ell - \ln \mu_{\mathcal{C}(\ell; d)}^\ell] \right) \quad (20) \\ & + (1 - \alpha) \left([\ln w^\ell(d') - \ln w^\ell(d)] - [\ln P^\ell(d') - \ln P^\ell(d)] \right). \end{aligned}$$

The direction of causes and effects is summarized in the diagram in Figure 3. It highlights the main fundamentals through which border changes operate in the framework: structure of trade costs and sizes and compositions of countries. It also shows the predicted effects on outcomes due to a particular type of border change: integration of 2 countries.


The only two outcomes whose direction of change is ambiguous are wages and costs from heterogeneity. Still, it is guaranteed that the real wage of all locations increases for all locations even if the relative nominal wage goes down for some of them. The same is not true for the cost of heterogeneity, as any location can turn out to be either a winner or a loser depending on how the distribution of preferences changes after integration. For example, if no location in either of the integrated countries has the same type as any of the locations in the other country, then everyone's disutility will increase after integration.

Given that, the scenario considered in Figure 3 illustrates one type of trade-off that border changes can introduce. On the positive side, integration brings a) economic benefits by reducing trade costs and increasing private consumption, and b) fiscal benefits by increasing the quantity of public goods that the government can provide. On the negative side, bigger countries are harder to govern, especially hurting remote locations, and potential higher heterogeneity makes the type of the provided public good less favorable on average.

Importantly, the effects of border changes differ in terms of whether they depend on the whole compositions of countries or just certain simple statistics of them. The economy-of-scale benefit

is driven purely by the total population size of a country, irrespective of the distribution of locations' population sizes. In contrast, the trade effects depend crucially on the distribution of productivity fundamentals within a country. Integrating a highly productive location into a country consisting of similarly unproductive locations, leads to a relatively homogeneous and big welfare response. At the same time, if productivities are highly heterogeneous in the target country, then the welfare response is less pronounced and also highly heterogeneous.

Figure 3: Mechanisms of how a country expansion affects welfare of its locations.



3 Optimal Borders

To assess costs of border configurations, I set up an Optimal Borders problem whose solution provides a benchmark. Optimal borders maximize utilitarian welfare of all countries subject to equilibrium conditions of the framework from the previous section. I achieve tractability of the problem by employing a Max-k-Cut approximation as a warm-start to the full non-linear problem.

3.1 Set-up

The crucial choice for the optimization strategy is modelling border configurations as sets of location-pair binary variables $d_{\ell\ell'}$. This enables a network perspective on the analysis of border configurations, as country partitions of the geography are equivalent to clusterings of the network of locations. As a result, the problem of finding a border configuration that satisfies certain criteria is akin to problems of finding appropriate clusterings of networks. The computational solving strategy largely builds upon relating the canonical problem of splitting

a given network into k clusters (Max- k -Cut) under the linear objective to the optimal borders problem.




Figure 4: Illustration of how dyadic variables $d_{\ell\ell'}$ reflect border configurations.

Denote $\mathbf{x}$ the set of endogenous outcomes in the developed framework, and Θ – the set of parameters. Partition the set of equilibrium conditions into two subsets: one for partial equilibrium equations $\tilde{\mathbf{g}}$, another for general equilibrium equations $\bar{\mathbf{g}}$ that describe the feedback effects. This separation will facilitate the decomposition strategy, explained further in this section. In case of the developed framework, only the condition pinning down wages (18) is of general equilibrium type, as it does not let express wages as an explicit function of d . Finally, define the indirect welfare function that maps border configurations d to equilibrium welfare levels in all locations $\mathcal{W}^\ell(d)$, and the indirect function for aggregate welfare $\mathcal{W}(d)$.

The Optimal Borders problem is an instance of MPEC. The social planner knows the fundamentals of the economy and values of all parameters. But it only optimizes over border configurations d , being constrained by equilibrium conditions of the economic framework (Definition 1). Because in the set-up framework agents consider the border configuration as given, the approach here does not facilitate analysis in terms of correcting externalities, as is typical in studies of optimal policy design. In this case, the purpose of solving the Optimal Borders problem is to provide a benchmark against which any given border configuration can be assessed, providing an understanding of why it is sub-optimal.

Definition 2 (Optimal Borders problem)

$$\max_d \quad \sum_{\ell} L^{\ell} \mathcal{W}^{\ell} \quad (21)$$

$$s.t. \quad \mathcal{W}^{\ell} = U^{\ell}(\mathbf{x}; d, \Theta) \quad (22)$$

$$\tilde{\mathbf{g}}(\mathbf{x}, d; \Theta) = 0 \quad (23)$$

$$\bar{\mathbf{g}}(\mathbf{x}, d; \Theta) = 0 \quad (24)$$

$$d \text{ represents a partition.} \quad (25)$$

There are 3 groups of reasons for hardness of this problem. Firstly, the space of all feasible country partitions is both discrete and large, making the problem inherit some common challenges of integer programming. Even if all of the equilibrium constraints are linear, the problem is $\mathcal{NP}$ -complete, which implies that generally, the Optimal Borders problem is $\mathcal{NP}$ -complete as well.

Two other challenges are posed by the economic nature of the problem. Firstly, equilibrium conditions can be non-linear in d . For example, the price index of any location (10) as a function of d is a composition of multiple functions, two of which are power functions. On top of that, they are non-convex, making it hard to ensure that any computed solution is the unique global maximum. Another manifestation of non-linearity is interactive nature of the effect of changing $d_{\ell\ell'}$ for multiple pairs of (ℓ, ℓ') . Again, for the price index (10), it is clear that the positive effect of integrating ℓ and ℓ' on prices in ℓ is dampened if at the same time some ℓ'' is integrated with it.

Secondly, the general equilibrium nature of the framework implies that certain outcomes can only be defined as an implicit function of d . This means that it is necessary to solve a system of non-linear equations for the social planner to evaluate the objective function for any d . In the context of the framework, such a system is formed by (11), (16), and (18), which jointly determine wages, prices, and costs of public goods.

The strategy for computationally solving the problem can be thought of as a decomposition of it in a way that tackles these challenges one at a time, rather than simultaneously.²

3.2 Max- k -Cut Approximation

The first step in the approximation is to split the full vector of outcomes $\mathbf{x}$ into two sub-components $\mathbf{x} = \{\tilde{\mathbf{x}}, \bar{\mathbf{x}}\}$ that I will call “partial equilibrium outcomes” and “general equilibrium outcomes” respectively. The purpose of this splitting is to separate out outcomes that can be explicitly defined as a function of d , conditional on values of the $\bar{\mathbf{x}}$. In the case of the framework

²Decomposition of mixed-integer problems into continuous and integer sub-problems has a long tradition in integer programming (e.g., see a textbook exposition of Benders decomposition in [Conforti et al. 2014](#)).

outlined in the previous section, only wages satisfy the definition of a “general equilibrium outcome”: $\bar{\mathbf{x}} = \{w^\ell\}_\ell$.

The nature of the approximation is to consider 1-step deviations from the baseline border configuration. Denote $\mathbf{d}_{\ell\ell'}$ a $\mathcal{L} \times \mathcal{L}$ matrix with all zeros except for 1 in (ℓ, ℓ') and (ℓ', ℓ) . The key object is the location-pair variable $\mathcal{W}_{\ell\ell'}^+$ that reflects the change in aggregate welfare due to setting to 1 the integration status just between ℓ and ℓ' , keeping the rest of values of d^0 intact. It is defined to be non-zero only if the integration status between these locations is 0 in the baseline border configuration. $\mathcal{W}_{\ell\ell'}^-$ is defined symmetrically.

Thus, the social planner in this problem is searching for the set of these 1-step deviations from the baseline border configuration that leads to the highest increase in aggregate welfare. It is clear that an arbitrary set of such deviations might lead to a logically infeasible border configuration. The last constraint on d imposes that such cases cannot be the solution of the problem.

Intuitively, this approach can be thought of as an analogue of the first-order Taylor approximation to the full non-linear problem, where infinitesimal marginal changes considered in continuous cases are replaced with discrete 1-step deviations. Similarly to continuous cases, this leads to ignoring interactive effects of border changes.

Definition 3 (Max-k-Cut approximation) *Given a feasible border configuration d^0 and an equilibrium generated by it $\mathbf{x}^0 \in \{\mathbf{x} : \mathbf{g}(\mathbf{x}, d^0; \Theta) = 0\}$, the Max-k-Cut approximation to the Optimal Borders problem around d^0 is defined by:*

$$\mathcal{W}^*(d^0) \equiv \max_d \sum_{\ell, \ell' \geq \ell} (1 - d_{\ell\ell'}) [\mathcal{W}_{\ell\ell'}^- - \mathcal{W}_{\ell\ell'}^+] \quad (26)$$

$$\begin{aligned} \text{s.t. } \mathcal{W}_{\ell\ell'}^+ &= (1 - d_{\ell\ell'}^0) \left[\sum_{\tilde{\ell}} L^{\tilde{\ell}} U^{\tilde{\ell}} (\tilde{\mathbf{x}}_{\ell\ell'}^+, \bar{\mathbf{x}}^0 | d^0 + \mathbf{d}_{\ell\ell'}, \Theta) \right] \\ \mathcal{W}_{\ell\ell'}^- &= d_{\ell\ell'}^0 \left[\sum_{\tilde{\ell}} L^{\tilde{\ell}} U^{\tilde{\ell}} (\tilde{\mathbf{x}}_{\ell\ell'}^-, \bar{\mathbf{x}}^0 | d^0 - \mathbf{d}_{\ell\ell'}, \Theta) \right] \\ \tilde{\mathbf{g}}(\tilde{\mathbf{x}}_{\ell\ell'}^+, \bar{\mathbf{x}}^0, d^0 + \mathbf{d}_{\ell\ell'} | \Theta) &= 0 \\ \tilde{\mathbf{g}}(\tilde{\mathbf{x}}_{\ell\ell'}^-, \bar{\mathbf{x}}^0, d^0 - \mathbf{d}_{\ell\ell'} | \Theta) &= 0 \\ d &\text{ represents a partition.} \end{aligned} \quad (27)$$

Proposition 2 *The solution to the [Max-k-Cut approximation](#) problem around d^0 yields a border configuration that is welfare-improving relative to d^0 . Formally:*

$$\mathcal{W}(d'^*) \geq \mathcal{W}(d^0).$$

where d'^* is a solution to the [Max-k-Cut approximation](#) around d^0 .

Algorithm 1: Algorithm to solve the [Optimal Borders problem](#)

Data: $\Theta, \mathbf{F}$
Result: $d^* \in \arg \max_d \mathcal{W}(d)$ s.t. (22) – (25)
 Choose d_0 ;
 Solve for $\mathbf{x}^0 : \mathbf{g}(\mathbf{x}^0, d^0 | \Theta) = 0$;
 $i \leftarrow 0$;
 Choose $i_{max}, \varepsilon \geq 0$;
 Solve for $d_{i+1} \in \arg \max_d \mathcal{W}^*(d; d_0)$ s.t. (27);
 Solve for $\mathbf{x}^{i+1} : \mathbf{g}(\mathbf{x}^{i+1}, d^{i+1} | \Theta) = 0$;
while $||\mathcal{W}(d_{i+1}) - \mathcal{W}(d_i)|| \geq \varepsilon$ *and* $i \leq i_{max}$ **do**
 $i \leftarrow i + 1$;
 Solve for $d_{i+1} \in \arg \max_d \mathcal{W}^*(d; d_i)$ s.t. (27);
 Solve for $\mathbf{x}^{i+1} : \mathbf{g}(\mathbf{x}^{i+1}, d^{i+1} | \Theta) = 0$;
end
 $d^* \leftarrow d_{i+1}$;

As neither uniqueness of the global maximum nor convergence of Algorithm 1 to it are guaranteed, I resort to a robustness strategy. Relying on the need to specify an initial d^0 , I simulate a set of them and report how variable the outputs are.

Definition of the Optimal Borders problem in the context of the framework:

$$\begin{aligned}
 & \max_d \max_{\mathbf{x}} \sum_{\ell} L^{\ell} \mathcal{W}^{\ell} \\
 & \text{s.t. } \mathcal{W}^{\ell} = (\mathcal{G}^{\ell})^{\alpha} (C^{\ell})^{1-\alpha} \\
 & \quad (3), (4), (9) - (11), (12), (14) - (18).
 \end{aligned}$$

3.3 Characterizing optimal borders in special cases

3.3.1 Maximizing gains from trade

This subsection sheds light on patterns of country compositions in terms of fundamentals that are characteristic of optimal border configurations. To that end, I consider a simplified setting with no trade costs and set up a problem, in which the social planner is maximizing aggregate gains from trade through choosing the whole set of productivities of all locations. Notice that this problem omits any consideration of border configurations. This way, the exercise can be intuitively thought of as learning what an optimal country looks like in terms of its composition.

$$\begin{aligned}
\sup_A \quad & \mathcal{W}(A) \equiv - \sum_{\ell} \theta \ln \pi_{\ell\ell} \\
\text{s.t.} \quad & \sum_{\ell} \pi_{\ell\ell} = 1 \\
& \pi_{\ell\ell} = \frac{A^{\ell} (w^{\ell})^{-\theta}}{\sum_{\ell'} A^{\ell'} (w^{\ell'})^{-\theta}}.
\end{aligned}$$

Proposition 3 $\mathcal{W}(A)$ is monotonically increasing in $[\max_{\ell} A^{\ell} - \min_{\ell} A^{\ell}]$.

4 Model quantification

This section brings the developed framework and optimization method to the African context. I calibrate framework parameters using either external calibration or the inversion of the equilibrium conditions given observed outcomes. Solving the Optimal Borders problem with calibrated values reveals significant a opportunity cost of the postcolonial borders and the relative importance of the position of borders compared to the number of countries.

4.1 Data

To calibrate the parameters Θ and fundamentals $\mathbf{F}$ of the model, I need both granular data on socio-economic and geographic variables, and the spatial distribution of ethnic groups. For economic and population data, I turn to the G-ECON dataset that measures gross output and population size at the level of 1°-by-1° cells.^{3,4} For spatial data on ethnic groups, I rely on the widely used Murdock map (Murdock, 1959). Due to a high level of detail in its classification of ethnic groups, the Murdock map features more than 800 groups, which contributes to the intractability of the search for optimal borders over such large territory. Therefore, I combine the Murdock map with the more aggregate classification in the Ethnic Power Relations (EPR) dataset that documents ethnic conflicts (Vogt et al., 2015; Wucherpfennig et al., 2012).

4.1.1 Spatial unit

As the number of spatial units is a significant factor of the tractability of the Optimal Borders problem, it is infeasible to use one degree grid cells as units of analysis, given the available

³The website of the G-ECON project: <https://gecon.yale.edu>.

⁴As measuring economic activity at such a granular level in the developing-country context can be challenging, two countries that are large geographically and population-wise are missing in the sample: Libya and Zimbabwe.

solvers. Therefore, I take two aggregation steps to form a sample of units that finds a balance between achieving tractability of the Optimal Borders problem and preserving the relevance of feasible border configurations. Firstly, I aggregate sets of neighboring grid cells into one unit if they are inhabited by the same ethnic group according to the Murdock map. However, as mentioned before, the resulting number of ethnic groups (around 800) is still intractable.

To aggregate ethnicity polygons in the Murdock map into larger spatial units that still reflect some notion of shared ethnic identity, I rely on the matching from the Murdock map to the EPR classification done by [Michalopoulos and Papaioannou \(2016\)](#). It provides a correspondence between names of ethnic groups in the Murdock map and the ones in the EPR dataset according to various, manually checked criteria. Thus, I merge sets of polygons in the Murdock map into one if they are matched to the same EPR group by [Michalopoulos and Papaioannou \(2016\)](#). In contrast to other commonly used ethnicity classifications (such as the Ethnologue), this approach identifies ethnic groups as entities that have shown united political interests. This matches an intuitive notion of political actors that could be relevant for real-world changes in border configurations. I later exploit this feature of the unit definition to match outputs of my model to historical secessionist attempts.

In order to keep the current border configuration as a feasible solution for the Optimal Borders problem, I define spatial units at the ethnic group by country level. That is, if an ethnic group is split by a current national border (for example, the Mbundu people live in both Angola and the Democratic Republic of the Congo), it leads to two spatial units in my sample. As a result, there are 363 units, which are displayed in Figure 5. Notice that although it seems from the map that there is a big dispersion in sizes of ethnic groups, that is not the case in terms of population sizes, as certain ethnic groups are spread out in areas with low population density.




Figure 5: Spatial units of analysis in the sample. Different colors represent different ethnic groups as classified by the EPR database.

4.1.2 Graph definition

Every spatial unit is represented by a node in the constructed graph. This graph is fully connected, with edges being potentially cut to represent the integration relation. While in theory it is possible to map jurisdictional partitionings onto planar graphs, the used algorithm to solve the Optimal Border problem mechanically delivers non-contiguous countries in that case. In the Max- k -Cut approximation of the problem, weights of the graph edges correspond to marginal effect values of changing the integration status between the according end locations.

4.2 Calibration of fundamentals

4.2.1 Trade costs

Following [Desmet et al. \(2018\)](#), the calibration of trade costs is an application of the framework of [Allen and Arkolakis \(2014\)](#), in which goods are shipped through the least costly transportation routes. In order to conform with the theoretical set-up of [Allen and Arkolakis \(2014\)](#), I represent the geography in a granular way by the grid of 1°-by-1° cells, indexed by $r \in \mathcal{R}$. Importantly, I discretize the whole planet, rather than just the African continent, to allow for possibilities like traders from Lagos using maritime routes to ship goods to Cape Town. Using data from Natural Earth, I characterize each cell by a set of available transportation modes.⁵ They include railroads, two types of roads (major and other), and water. These transportation modes have different costs, reflected by parameters estimated in [Allen and Arkolakis \(2014\)](#). As a result, each grid cell r has a cost value $\gamma(r)$ of trespassing it along a transportation route, parametrized by

$$\log \gamma(r) = \sum_m \sum_{s \in \mathcal{S}_m} \log \gamma_m^s m^s(r) + \sum_m \log \bar{\gamma}_m [1 - m(r)],$$

where m indexes transportation modes, s indexes possible sub-types of the considered mode, γ_m^s is the mode-sub-type-specific cost parameter, $\bar{\gamma}_m$ is the cost parameter for the case when m is not present, and $m(r)$ is the indicator function that takes value 1 if transportation mode m is available in grid cell r . A transportation route between cells r_o and r_d is an ordered sequence of cells, where every next cell is neighboring the previous one. The resulting trade cost between r_o and r_d denoted $\Gamma(r_o, r_d)$ is a cumulative function of costs of trespassing grid cells that lie on the least costly path between r_o and r_d . The continuous version of this problem is captured by

$$\Gamma(r_o, r_d) = \left[\inf_{u(r_o, r_d)} \int_{u(r_o, r_d)} \gamma(r) dr \right]^v,$$

⁵The website of the Natural Earth project: <https://www.naturalearthdata.com>.

where v is the elasticity of converting transportation costs into trade costs, and $u(r_o, r_d)$ is some route between r_o and r_d .

The cell-level cost of transporting $\gamma(r)$ is calibrated using the data that is more granular than the cell size. It is set equal to the average bilateral transporting cost across all pairs of sub-cells.

In order to aggregate bilateral cell-level trade costs into bilateral location-level trade costs, I follow a similar strategy as in the case of cell-level transportation cost by averaging trade costs between all pairs of cells from the two considered locations.

4.2.2 Productivity fundamentals

Values of productivity fundamentals A^ℓ are set such that, given the observed income levels w^ℓ , calibrated trade costs $\hat{\gamma}$, and the calibrated values of the parameters $\hat{\Theta}$, the general equilibrium conditions that depend on A^ℓ hold. Such conditions include (18) together with (10), (11), (14), (16). Formulating these conditions as a function of A , the inversion problem amounts to solving the following system:

$$\mathbf{G}(A; w, L, \hat{\tau}, \hat{\Theta}) = 0. \quad (28)$$

The important data variation for recovering productivity fundamentals comes from income levels w^ℓ . However, they are not necessarily correlated as the incomes of locations are determined not only by their technologies but also by their market access.

Results of inverting the general equilibrium conditions to obtain values of productivity fundamentals A^ℓ are displayed in Figure 6b, alongside observed income levels.

4.2.3 Preference heterogeneity

To calculate utility discounts due to heterogeneity $\delta_{C(\ell)}^\ell$ as a function of arbitrary border configurations, it is necessary to quantify bilateral preference distance fundamentals $\delta_{\ell\ell'}$. Following the literature that associates ethnic heterogeneity with preference heterogeneity when explaining low quality of provided public goods (e.g., [Alesina and Zhuravskaya 2011](#)), I conceptualize $\delta_{\ell\ell'}$ as a measure of ‘distance’ between ethnic groups in ℓ and ℓ' .⁶ This is another way in which associating locations with ethnic groups facilitates the calibration.

As a baseline, I impose a simple parametrization in which every ethnic group only enjoys their ideal type of public good and completely distastes ideal types of other ethnic groups: $\delta_{\ell\ell'} = 0 \forall \ell \neq \ell'$. On its own, the strength of this assumption is reduced by how aggregated the chosen level of ethnicity classification is. For example, it is more reasonable to believe that

⁶The term ‘distance’ is a slight abuse of terminology as $\delta_{\ell\ell'}$ takes value 1 if ℓ does not experience any disutility discount from the ideal public good of ℓ' and, thus, is similar to ℓ' in this sense.




Figure 6: Results of inverting the general equilibrium conditions to obtain productivity fundamentals A .

the Muslim Arabs and the Christian Copts have bigger disagreements over the preferred public good than, say, the Tuareg and the Beydan, who are both recognized by EPR to be in the same group as the Muslim Arabs. In other words, I set $\delta_{\ell\ell'}$ between the Tuareg and the Beydan to 1, but the one between the Muslim Arabs and the Christian Copts to 0.

As discussed in subsection 2.4, such parametrization leads to $\delta_{\mathcal{C}(\ell)}^\ell$ being equal to the population share of ℓ in $\mathcal{C}(\ell)$. Then, the population-weighted average level of ethnic remoteness in country $\mathcal{C}(\ell)$ becomes exactly the ethnic fractionalization index, widely used in the literature studying ethnic heterogeneity (e.g., [Alesina and Ferrara 2005](#); [Alesina et al. 2019](#); [Montalvo and Reynal-Querol 2021](#)). Therefore, one can think of the calibration approach taken here as an implementation of the endogenous ethnic fractionalization index whose value can be computed for an arbitrary composition of countries.

If, alternatively, $\delta_{\ell\ell'}$ took values strictly between 0 and 1, we would expect a spatial gradient in the values of $\delta_{\ell\ell'}$ as ℓ' gets geographically more distant from ℓ . The reason for such gradient is that ethnic groups which are closer to each other in terms of identity and cultural traits are likelier to reside in relatively closer geographic areas. The effect of this gradient on the optimal sizes of countries is ambiguous. Directly, integrating with neighboring, albeit different, ethnic groups leads to a smaller increase in the disutility discount for the public goods, which strengthens the force to make countries larger. At the same time, this gradient would dampen the economic integration motif to make countries larger. With a higher potential to exploit gains from fiscal integration, the social planner cares less about increasing technological heterogeneity within countries, which sometimes can only be achieved by bigger country sizes.

Table 1: Summary of sources for calibrated parameters.

Parameter	Description	Source	Value
α	Utility	Desmet et al. (2022)	.3
<i>Trade</i>			
β	Border effect	Conte (2022)	.2
$\gamma_{\ell j}$	Non-institutional component	Allen and Arkolakis (2014)	
ν	Elasticity of trade costs wrt transportation cost	Desmet et al. (2018)	.363
θ	Trade elasticity	Conte (2022)	6.63
<i>Public good</i>			
t	Tax rate	National income accounts	.2
ξ	Degree of rivalry	Jannin and Sotura (2020) ¹	.7
δ	Ethnic distances	EPR classification	
N	Number of units	EPR groups $\cap$ Countries	363

¹ Jannin and Sotura (2020) do not settle on the most preferred value of this parameter (denoted κ in their case), as their estimates are not robust across specifications. At the same time, the literature does not offer alternative estimates of this parameter. For example, Fajgelbaum et al. (2019) altogether avoid the issue by only considering the extreme cases of no rivalry and full rivalry.

5 Applying the framework to Africa

5.1 Analysis of secessionist attempts

As the first application of the developed framework, I turn to the question of secessionist movements and test whether the my calibrated framework can speak to the historically observed secessionist behaviors. I use EPR dataset’s records of rebellious actions taken by ethnic groups as an indicator of discontent with their position in the country and, thus, a step towards potential secession from the country. I encode the existence of such a record for an ethnic group as a binary variable, as EPR does not provide a quantification of the degree to which the rebellious activity was prominent. Because this indicator variable on secessionist attempts was not used in the calibration of the model, this analysis can be also seen as an external validity test of my calibration.

Employing the calibration of my framework, I construct variables dubbed “secession outcomes” to associate with secessionist attempts. It varies at the level of the spatial unit (i.e., ethnicity-by-country unit) and results from computing particular counterfactuals within the framework. For every spatial unit ℓ , I consider the counterfactual in which it secedes from its current country and forms a new, separate one. As a result, a new set of national boundaries arise that surround the considered ethnic group. Formally, this amounts to setting $d_{\ell\ell'} = 0 \forall \ell' \neq \ell$. Thus, I construct 363 counterfactual border configurations – one for every secession scenario.


For each such scenario, I recompute the general equilibrium outcomes of the framework using the calibrated values of parameters and fundamentals. I interpret these as predictions of what would happen in case any considered ethnic group decided to secede from its current country

(for example, what would happen if the Tutsi people seceded from Ethiopia?) I particularly focus on the change in the overall welfare of the ethnic group in case it decides to secede:


$$\Delta \mathcal{W}^\ell = \mathcal{W}^\ell(\ell \text{ secedes from } \mathcal{C}(\ell)) - \mathcal{W}^\ell(\text{status quo}).$$

This change comprehends both the trade and the public goods motifs behind secessions and, thus, should reflect the trade-off between gaining full control over the provision of public goods and losing easier trade access with the rest of the status quo country.


Figure 7: Effects of counterfactual secessions for all ethnic groups.




(a) Marginal distribution of secession values relative to the baseline.



(b) Marginal distribution of secession values in US Dollars.



(c) Conditional distributions of secession effects relative to the baseline.



(d) Conditional distributions of secession effects in US Dollars.

The marginal distribution of the welfare change in case of seceding is plotted in Figures 7a (relative change compared to status quo) and 7b (2005 US Dollars equivalent of the welfare change). One immediate pattern of this distribution is that most of the welfare effects from seceding are positive, with the average of around \$500 mln in 2005 US Dollars equivalent. Under the assumption of credibility of the calibrated framework, it evokes two important conclusions. Firstly, it suggests the sub-optimality of the post-colonial national borders. If most ethnic groups would prefer to exit their current countries, European colonial powers did not settle on the border configuration that maximizes integration synergies between ethnic groups in the same country. Note that this does not imply that the configuration with only ethnic states

would be Pareto-improving or even increase aggregate welfare. Due to the general equilibrium nature of the framework, it is crucial for the obtained positive numbers that the rest of the status quo border configuration stays intact.

Secondly, reconciling these positive values with the fact that secessionist attempts are quite rare in practice suggests that attempting a secession is very costly to the seceder. Staging a rebellion requires a significant amount of arms that most likely needs to be imported from outside the country. Sufficient production or financial resources are hard to acquire for a sub-national ethnic region. Also, civil wars almost inevitably lead to human losses. Therefore, the rarity of secessions might speak to the value of life by potential participants and victims of war.

Still, there have been historical cases of ethnic secessionism in Africa, and they can be related to the predicted welfare gains from seceding in my framework. Figures 7c and 7d compare conditional distribution of secession effects on welfare for those who attempted a secession (seceders) and for those who did not. Clearly the conditional density for seceders is shifted to the right compared to the conditional density for non-seceders, which speaks in favor of the relevance of my calibrated framework. The ethnic groups that tried to secede had sufficiently higher prospective gains compared to the non-seceders to outweigh the costs of seceding.

Table 2: Comparisons of outcomes of counterfactual secessions between actual seceders and non-seceders.

Outcome	<i>Means</i>			<i>Test p-values</i>			
	Seceders	Non-Seceders	Diff	ANOVA	Kolmogorov-Smirnov test	Mann-Whitney U test	Kruskal-Wallis test
Log welfare change for seceder	0.351	0.296	0.055	0.078	0.054	0.126	0.125
Log change in real income	−0.035	−0.046	0.011	0.036	0.255	0.564	0.564
Welfare change for seceder in 2005 US dollars equivalent (millions)	888.953	446.781	442.172	0.026	0.0	0.0	0.0


Technically speaking, it is not straightforward to characterize the statistical nature of these distributions of outcomes that are obtained through counterfactual analysis within a general equilibrium structural framework. Still, as a way to verify that the difference in the obtained conditional distributions is not only visual, I apply a series of statistical tests designed for cases when the available sample is not drawn from the joint distribution of two compared outcomes. The p -values for these tests are reported in Table 2. For the US-dollars-equivalent measure of the secession effect, the differences are strongly significant according to these tests.

5.2 Optimal borders in Africa


5.2.1 Welfare implications

To understand whether and why post-colonial borders are bad, I consider the counterfactual obtained as the solution to the Optimal Borders problem with calibrated parameter values, and compare it to the status quo. The optimal borders for Africa differ substantially from the status quo ones. Here, I mainly focus on the differences in aggregate welfare measures and the associated changes in fundamentals, leaving an overview of all outcomes to the Appendix.

Aggregate welfare as a function of the number of countries First, I examine the relative importance of the number of countries compared to the compositions of countries. To that end, I solve the [Optimal Borders problem](#) with an additional constraint on the number of countries. Specifically, I sequentially set the allowed number of countries from 2 to 60 and obtain 59 solutions. In Figure 8a, I use blue dots to report the aggregate welfare under the optimal country composition for each considered number of countries (on the x -axis) and fit a quadratic curve (red line) to these values. Aggregate welfare levels only have a relative meaning, with the reference point being the status quo configuration represented by the orange dot.



(a) Global welfare under optimal borders, conditional on the pre-set number of countries.



(b) Distributive impact of optimal borders, conditional on the pre-set number of countries.

Figure 8: Aggregate welfare as a function of the number of countries, conditional on the optimal geographic position of borders.

The inverse- U shape of the indirect utility curve is consistent with existing theoretical frameworks that posit the trade-off between the cost of social tension and the benefit of heterogeneity. As the number of countries increases, the benefits of fewer tensions are nullified by lost economic opportunities that arise due to gains from trade and economies of scale in public goods provision. To further emphasize the importance of the composition of countries, I plot distributive consequences of re-arranging borders in Figure 8b as a function of the number of countries. Comparing boxplots across numbers of countries shows that even though the aggregate welfare varies significantly with the number of countries, most of the locations would gain

from the re-arrangement of borders regardless of the number of countries. That is, relative to the status quo, moving to almost any counterfactual number of countries would benefit most of the locations as long as the geographic position of borders is optimal.




Figure 9: The effect of the optimal border configuration on the spatial distribution of welfare

Quantitatively, in the counterfactual with optimal borders and the current number of countries, the aggregate welfare is 28% higher than in the status quo. Although this counterfactual does not lead to a Pareto improvement, at least 75% of regions (85% of African population) gain from transitioning to the optimal borders. This aggregate welfare gain is massive. To compare with another continent-wide counterfactual reform, [Graff \(2024\)](#) calculates that the optimal re-organization of the African road network country-by-country would yield a 1.8% aggregate welfare gain.

Besides the statistical uncertainty not accounted for in the analysis here, two economic reasons can explain such stark difference. First, the workers in [Graff \(2024\)](#) derive utility only from the consumption of private goods, while my framework includes the consumption of public goods. Due to this utility component, changes in the country's composition can have drastic discontinuous effects on welfare. For example, when some poor location gets integrated with a much richer location, the quantity of the provided public good in that location can increase dramatically due to non-complete rivalry in public goods consumption.⁷

Second, in terms of the trade forces, the marginal effect of improving roads on the trade costs is not high enough to compare with the discrete jump in trade costs due to a country border. Given the moderate level of baseline infrastructure in Africa, it is hard for the optimal re-organization to achieve a high increase in aggregate gains from trade. At the same time,

⁷Using the framework's notation, due to ξ being less than 1.

as discussed later in this section, the optimal borders ensure that locations with a higher potential for gains from trade due to a big difference in productivities are less likely to have a country border between them. The spatial distribution of productivity fundamentals in Africa is dispersed enough to allow the social planner to significantly increase within-country differences in productivities through re-organizing borders.

Strikingly, the actual number of countries (45) is close to the number that yields the maximum aggregate welfare (50), with a minor difference in the welfare value.⁸ This speaks to a long-standing discussion in the literature on the political economy of Africa around the number of countries inherited from the colonial era.




Figure 10: The solution to the approximation to the partial equilibrium Optimal Borders problem

Spatial redistributive implications of the optimal borders Figure 11 plots the welfare change at the location level against the status quo level of welfare. The negative association between these two variables indicates that the optimal borders lead to the partial convergence of welfare levels across space. In particular, regions at extremes of the status quo welfare distribution experience the most impact in the counterfactual with the optimal border configuration.

Figure 9 shows significant spatial heterogeneity in welfare impacts of optimal borders. There

⁸See Data Appendix for the explanation of why the number of countries in the sample is lower than the actual one.




Figure 11: Redistributive consequences of optimal borders.

are both areas that significantly gain and areas that significantly lose in the considered counterfactual. A good predictor of the predicted welfare change is the status quo welfare (see Figure 11), suggesting that the social planner is trying to reduce spatial inequality. Similarly, one can observe that inland areas are more likely to gain, at the cost of initially richer coastal regions. From this perspective, the social planner is re-configuring borders so that the high gains from trade associated with access to ports trickle down to remote regions.

Mechanisms In terms of the mechanisms, optimal borders lead to patterns of fundamentals that are consistent with the logic formally derived in Section 3. Firstly, it is evident that countries in the optimal configuration are on average less ethnically diverse (see Figure 12). This matches the common perception that post-colonial border design was negligent of the historical distribution of ethnic groups in Africa.




Figure 12: Social planner is reducing ethnic heterogeneity across countries, making the distribution of fractionalization indices have higher frequency around 0.

On top of that, the pattern of spatial correlation between ethnic and technological heterogeneity in the optimal borders case is consistent with the logic that the cost of higher ethnic heterogeneity should be compensated with higher gains from trade. The latter follow higher technological heterogeneity, and the social planner strengthens the association between ethnic and technological heterogeneity (see Figure 13).




Figure 13: Social planner is improving the trade-off between economic and public goods forces: With optimal borders, ethnically heterogeneous countries are likelier to have higher technological heterogeneity – a positive factor for higher gains from trade.

5.2.2 The Case of South Africa

To understand why particular border stretches are different under the optimal configuration, I zoom into the region comprising South Africa and its neighboring countries, and decompose how individual spatial units are affected by a counterfactual change in borders.




Figure 14: Optimal border configuration for the Southern Region.


In this sub-section I define as South the following set of countries: South Africa, Namibia, Mauritius, Lesotho, and Botswana; the rest of African countries in my sample are denoted as North. To explain why social planner is splitting the South into ethnically homogenous countries (see Figure 14), I explore differences in fundamentals between the North and the South (see Figure 15) and tie their patterns to the logic of my framework.

Firstly, it is important to notice that the social planner makes more people live in large countries (see Figure 15a). This leads to the question of why the North is getting large countries at the cost of small countries in the South under optimal borders.


Differences in ethnic heterogeneity and spatial productivity patterns indicate that benefits to larger country sizes are higher in the North than in the South.

Southern regions are less ethnically heterogeneous than the Northern region (see Figure 15b), which technically yields room for making countries more homogeneous in the South. On the


Figure 15: Patterns of differences in fundamentals between the North and the South.




(a) Optimal borders make more people live in larger countries.



(b) Ethnic heterogeneity is higher in the North.



(c) Baseline disutility from ethnic heterogeneity is higher in the North.



(d) Distribution of productivities is more dispersed in the North.

contrary, due to higher ethnic diversity in the North, it is technically harder to find a border configuration that makes an average country more homogenous without losing out significantly on other margins of welfare. The margin for increasing the utility from public goods has high potential in the South. This is reinforced by the fact that the distribution of real income is more uniform in the South.

In status quo, locations in the South have less disutility from heterogeneity in consumption from public goods (see Figure 15c). This implies that the margin for increasing the utility from public goods has high potential in the South. Furthermore, the fact that the distribution of real income is more uniform in the South ensures that the loss in the economy of scale in public goods production is less significant there. The opposite picture is true for the North. The baseline level of disutility due to ethnic heterogeneity is high, making marginal improvements on this dimensions not significant.

Another force in driving higher benefits from larger country sizes in the Northern regions is higher dispersion in productivities (see Figure 15d). Consistent with a key feature of the theoretical framework, integration of locations with big differences in productivities leads to

higher aggregate gains from trade. Thus, it is expected that optimal borders seek to exploit this potential by expanding sizes of countries in the North.

6 Conclusion

This paper developed a new theoretical framework to analyze jurisdictional borders. Moving beyond stylized models that offer abstract characterizations of border configurations, this paper's framework quantitatively assesses the welfare effects of an arbitrary change in border configuration in general equilibrium with a realistic geography. It still reflects the main forces that the theoretical literature has explored. The bigger population size of a country introduces fiscal savings in environments with non-fully-rival national public goods. It also implies a larger private market unaffected by the need to cross national borders. At the same time, geographically larger countries are harder to govern, differentially affecting regions on the country's outskirts. Alongside the size, the composition of a country is crucial for its welfare. Higher heterogeneity in terms of preferences over horizontal types of public goods introduces tension and makes it likelier that, on average, everyone is not content with the provided type of the public good. However, high heterogeneity in productivity fundamentals induces higher gains from trade, as the terms-of-trade effect for locations with lower productivities outweighs the losses from reallocating production to locations with higher productivities.

Keeping track of this multitude of forces, especially given that the space of theoretically feasible border configurations is huge and cannot be ordered in any obvious way, complicates the analytical characterization of welfare-maximizing borders. In this paper, I make a step in this direction by developing a feasible way to computationally obtain a welfare-optimizing border configuration despite the richness of the general equilibrium structure. It hinges on recasting the optimal borders problem as a network optimization problem in which border configurations are equated with clusterings of the network of locations. Then, decomposing the problem into two sub-problems, integer and non-linear, makes the computational solution feasible with a modern-class solver.

Multiple natural and exciting avenues for further research can build upon this paper's contribution. Firstly, this paper does not attempt to characterize any notion of the decentralized border configuration. As much as an analytical solution to the Optimal Borders problem is hard to get because of the integer nature of the optimized outcome, there is no obvious way to approach a characterization of an equilibrium border configuration formally. One promising direction draws on the coalitional game theory, which offers formal notions of stable partitions of a given set of agents. Although it has been mainly developed for settings with linear utilities and superadditive characteristic functions, there is growing literature studying coalitional games in settings that match the general equilibrium structure of this paper's framework.

One important component missing from this paper's framework seems to be the endogenous choice to engage in a military conflict. This extension is particularly relevant for the question of border formation as many wars are fought to annex new territories or secede from an oppressive federal government. Thus, it would be a step in micro-founding the decentralized determination of border configurations. Crucially, a framework with endogenous conflict would enable a quantitative assessment of realistic border changes. Adding an evaluation of the cost of war can substantially enhance the analysis of secessions offered in this paper. Such a framework would be able to predict whether a region will secede, not just whether it is more likely to secede, as it would quantify both the benefits and costs of seceding. More generally, the framework offered in this paper lends itself to feasible incorporations of forces developed in other frameworks with networks of agents. However, it will be challenging to establish whether the tractable solution strategy in this paper performs as well in extended frameworks.

References

- Alesina, A., Angeloni, I., and Etro, F. (2005a). International Unions. *American Economic Review*, 95(3):602–615.
- Alesina, A. and Ferrara, E. L. (2005). Ethnic Diversity and Economic Performance. *Journal of Economic Literature*, 43(3):762–800.
- Alesina, A., Gennaioli, C., and Lovo, S. (2019). Public Goods and Ethnic Diversity: Evidence from Deforestation in Indonesia. *Economica*, 86(341):32–66.
- Alesina, A. and Spolaore, E. (1997). On the Number and Size of Nations. *The Quarterly Journal of Economics*, 112(4):1027–1056.
- Alesina, A., Spolaore, E., and Wacziarg, R. (2005b). Chapter 23 - Trade, Growth and the Size of Countries. In Aghion, P. and Durlauf, S. N., editors, *Handbook of Economic Growth*, volume 1, pages 1499–1542. Elsevier.
- Alesina, A. and Zhuravskaya, E. (2011). Segregation and the Quality of Government in a Cross Section of Countries. *American Economic Review*, 101(5):1872–1911.
- Allen, T. (2023). The Topography of Nations.
- Allen, T. and Arkolakis, C. (2014). Trade and the Topography of the Spatial Economy. *The Quarterly Journal of Economics*, 129(3):1085–1140.
- Caliendo, L., Opromolla, L. D., Parro, F., and Sforza, A. (2021). Goods and Factor Market Integration: A Quantitative Assessment of the EU Enlargement. *Journal of Political Economy*, 129(12):3491–3545.
- Caliendo, L. and Parro, F. (2015). Estimates of the trade and welfare effects of NAFTA. *Review of Economic Studies*, 82:1–44.
- Chalkiadakis, G., Elkind, E., and Wooldridge, M. (2012). *Computational Aspects of Cooperative Game Theory*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Springer International Publishing, Cham.
- Conforti, M., Cornuéjols, G., and Zambelli, G. (2014). *Integer Programming*, volume 271 of *Graduate Texts in Mathematics*. Springer International Publishing, Cham.
- Conte, B. (2022). Climate Change and Migration: The Case of Africa. *SSRN Electronic Journal*.
- Desmet, K., Nagy, D. K., and Rossi-Hansberg, E. (2018). The Geography of Development. *Journal of Political Economy*, 126(3):903–983.

- Desmet, K., Ortuño-Ortín, I., and Özak, Ö. (2022). Is Secessionism Mostly About Income or Identity? A Global Analysis of 3,003 Subnational Regions.
- Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70(5):1741–1779.
- Esteban, J. and Ray, D. (2011). Linking Conflict to Inequality and Polarization. *American Economic Review*, 101(4):1345–1374.
- Fajgelbaum, P., Gaubert, C., Gorton, N., Morales, E., and Schaal, E. (2023). Political Preferences and Transport Infrastructure: Evidence from California’s High-Speed Rail. Technical Report w31438, National Bureau of Economic Research, Cambridge, MA.
- Fajgelbaum, P. D., Morales, E., Suárez Serrato, J. C., and Zidar, O. (2019). State Taxes and Spatial Misallocation. *The Review of Economic Studies*, 86(1):333–376.
- Fajgelbaum, P. D. and Schaal, E. (2020). Optimal Transport Networks in Spatial Equilibrium. *Econometrica*, 88(4):1411–1452.
- Fernández-Villaverde, J., Koyama, M., Lin, Y., and Sng, T.-H. (2023). The Fractured-Land Hypothesis. *The Quarterly Journal of Economics*, 138(2):1173–1231.
- Friedman, D. (1977). A Theory of the Size and Shape of Nations. *Journal of Political Economy*, 85(1):59–77.
- Gancia, G., Ponzetto, G. A. M., and Ventura, J. (2022). Globalization and Political Structure. *Journal of the European Economic Association*, 20(3):1276–1310.
- Graff, T. (2024). Spatial inefficiencies in Africa’s trade network. *Journal of Development Economics*, 171:103319.
- Jannin, N. and Sotura, Â. (2020). This Town Ain’t Big Enough? Quantifying Public Good Spillovers. *SSRN Electronic Journal*.
- Lebrand, M. (2021). Corridors without Borders in West Africa. *Policy Research Working Paper*, 9855.
- Michalopoulos, S. and Papaioannou, E. (2016). The Long-Run Effects of the Scramble for Africa. *American Economic Review*, 106(7):1802–1848.
- Montalvo, J. G. and Reynal-Querol, M. (2021). Ethnic diversity and growth: Revisiting the evidence. *Review of Economics and Statistics*, 103(3):521–532.
- Murdock, G. P. (1959). *Africa: Its Peoples and Their Culture History*. McGraw-Hill.
- Oduntan, G. (2015). Africa’s border disputes are set to rise – but there are ways to stop them. <http://theconversation.com/africas-border-disputes-are-set-to-rise-but-there-are-ways-to-stop-them-44264>.

- Saad, W., Han, Z., Debbah, M., Hjørungnes, A., and Basar, T. (2009). Coalitional game theory for communication networks. *IEEE Signal Processing Magazine*, 26(5):77–97.
- Thomas, C. G. (2018). Secession and Separatism in Modern Africa. In Shanguhya, M. S. and Falola, T., editors, *The Palgrave Handbook of African Colonial and Postcolonial History*, pages 729–757. Palgrave Macmillan US, New York.
- Validi, H. and Buchanan, A. (2022). Political districting to minimize cut edges. *Mathematical Programming Computation*, 14(4):623–672.
- Validi, H., Buchanan, A., and Lykhovyd, E. (2022). Imposing Contiguity Constraints in Political Districting Models. *Operations Research*, 70(2):867–892.
- Vogt, M., Bormann, N.-C., Rüegger, S., Cederman, L.-E., Hunziker, P., and Girardin, L. (2015). Integrating Data on Ethnicity, Geography, and Conflict: The Ethnic Power Relations Data Set Family. *Journal of Conflict Resolution*, 59(7):1327–1342.
- Weese, E. (2015). Political mergers as coalition formation: An analysis of the Heisei municipal amalgamations. *Quantitative Economics*, 6(2):257–307.
- Wucherpennig, J., Metternich, N. W., Cederman, L.-E., and Gleditsch, K. S. (2012). Ethnicity, the State, and the Duration of Civil War. *World Politics*, 64(1):79–115.