

The Geography of Linguistic Diversity and the Provision of Public Goods*

Klaus Desmet
SMU, NBER and CEPR

Joseph Flavian Gomes
UCLouvain

Ignacio Ortuño-Ortín
UC3M

January 2019

Abstract

This paper analyzes the importance of local interaction between individuals of different ethnolinguistic groups for the provision of public goods at the national level. The conceptual framework we develop suggests that a country's public goods (i) decrease in its overall *ethnolinguistic fractionalization*, and (ii) either increase or decrease in its *local-global ethnolinguistic complementarity*, a measure of how local interaction affects antagonism towards other groups in the society at large. After constructing a 5 km by 5 km dataset on language use for 223 countries, we empirically explore these theoretical predictions. While overall fractionalization worsens public goods outcomes, local interaction mitigates this negative association. Conditional on a country's overall diversity, public goods outcomes are maximized when there are a few large-sized groups and the diversity of each location mirrors that of the country as a whole.

Keywords: ethnolinguistic diversity, geographic diversity, public goods, local interaction, antagonism

JEL Codes: H4, O18, O57, R1, Z10, Z13

1 Introduction

Although living in ethnolinguistically heterogeneous societies can be challenging, within countries diversity is often viewed more positively in highly diverse areas than in relatively homogenous locations. For example, in sub-Saharan Africa more heterogeneous countries exhibit relatively lower degrees of inter-ethnic trust, but within countries inter-ethnic trust tends to be higher in highly heterogeneous localities (Robinson, 2013). As another example, while the pro-Brexit movement ran on a platform of regaining control over the inflow of migrants into the United Kingdom, the perception of there being too much diversity was especially strong in areas with few foreign-born residents, and much less so in cosmopolitan London (Lawton and Ackrill, 2016). These examples are consistent with the view that local interaction may mitigate the antagonism individuals

*Desmet: Department of Economics and Cox School of Business, Southern Methodist University, Dallas, TX, kdesmet@smu.edu, Gomes: IRES, LIDAM, Université catholique de Louvain, Louvain-la-Neuve, Belgium, joseph.gomes@uclouvain.be, Ortuño-Ortín: Department of Economics, Universidad Carlos III, Getafe (Madrid), Spain, ior-tuno@eco.uc3m.es. Gomes gratefully acknowledges financial support from the Fundación Ramón Areces, and Ortuño-Ortín thanks the Spanish Ministry of Economics and Competitiveness (grant ECO2013-42710-P and MDM 2014-0431). We thank Nathan Nunn for sharing the concordance data between ethnic groups in the Ethnographic Atlas and linguistic groups in the Ethnologue. We benefitted from discussions with Damian Clarke, Boris Gershman, Stelios Michalopoulos, Ömer Özak, Diego Puga, Frédéric Robert-Nicoud, Jan Stuhler, Elisabet Viladecans and Romain Wacziarg, as well as from presentations at Calcutta, Claremont, Essex, Heidelberg, Lancaster, Miami, the Barcelona Summer Forum on "Geography, Economics and Political Structure", the NBER Political Economy Meeting, the ACCID Conference Jadavpur University, the Midwest Economy Theory meeting, the CEPR/CURE workshop in Paris, the University of Oxford CSAE Conference, the 10th CESifo Workshop on Political Economy, the British Academy Workshop on "Interdisciplinary Perspectives on Peace Operations and Conflict Resolution".

feel towards other groups in society.¹ Although there is (quasi)-experimental micro-evidence supporting this view, to date no large-scale analysis, spanning the entire globe, has been conducted. This paper sets out to fill this void.

We start by proposing a simple measure of how the intensity of local interaction with other groups affects the antagonism an individual feels towards those other groups in the rest of society. The effect may be positive or negative, depending on whether local interaction improves or worsens people’s prejudice towards the out-group. Formally, we assume that the intensity of an individual’s local interaction with another group is increasing in the local share of that group; and how much this affects his antagonism towards that other group in the society at large depends on the aggregate share of that group. To illustrate this with a simple example, the local interaction of a Dutch-speaking Belgian with French-speaking Belgians depends on the local share of French speakers, but how much that interaction affects the overall antagonism she experiences towards that group depends on the aggregate share of French speakers. This yields a complementarity between the local group share and the aggregate group share in their effects on an individual’s antagonism. The intuition for this local-global complementarity is straightforward: the change in an individual’s opinion about another group can only have an important impact on her overall antagonism if that other group is sizable in the society at large.

Why do we care about inter-group interaction? Antagonism between groups is a leading explanation for why more diverse countries have worse political economy outcomes. This is why many empirical papers have explored the effect of ethnolinguistic fractionalization on different economic, social and political outcomes.² We suggest that these studies suffer from an omitted variable bias when not controlling for local interaction. To be more precise, our measurement framework yields an antagonism index that is a linear combination of two terms: the society’s *global ethnolinguistic fractionalization* and the society’s *local-global ethnolinguistic complementarity*. In the absence of local interaction, antagonism is measured by the degree of global ethnolinguistic fractionalization; in the presence of local interaction, that antagonism is either diminished or augmented by the degree of local-global ethnolinguistic complementarity. Drawing the correct inference about the effect of diversity on different political economy outcomes requires us to control for both terms. In this paper, we focus on a variety of public goods outcomes related to health, education and infrastructure, and empirically explore the effects of a country’s global fractionalization and its local-global complementarity.

In addition to proposing an antagonism index that captures how local interaction affects the degree of antagonism between groups at the national level, this paper makes three more contributions. First, we model the link between a society’s antagonism and its public goods. We start by postulating that an individual’s valuation of public goods depends on how much antagonism she feels towards others in society. We then

¹In the literature, this evidence is often interpreted as evidence in favor of contact theory in social psychology: the idea that frequent interaction reduces prejudices against the out-group (Allport, 1954). As we later discuss, this is of course not the only reason why local interaction can make people change their attitude towards other groups.

²See, e.g., Easterly and Levine (1997), Alesina et al. (2003), Alesina and La Ferrara (2000) and Alesina, Baqir and Hoxby (2004).

show that if public goods are financed by private contributions, their provision depends on a society’s average antagonism. Second, to bring the theory to the data, we develop a global database of language use at the local level. Third, we empirically analyze the effect of our new antagonism measure on public goods for virtually all countries of the world, with the goal of understanding whether local interaction improves or worsens outcomes.

Before empirically investigating the effect of the spatial distribution of ethnolinguistic diversity on the provision of public goods, we need detailed geographic information on the number of people that belong to each group. While we generically refer to the different groups as ethnolinguistic groups, for reasons of data availability in the empirics we focus on linguistic groups. The Ethnologue has information on the number of speakers per country of 6,905 unique languages (Lewis, Simons and Fennig, 2014). It also provides language maps that show the geographic distribution of languages in each country. Together with information on population at a fine geographic resolution from Landsat, we use an iterative proportional fitting algorithm to assign the number of speakers of different languages to each 5 km by 5 km cell in the world. We cross-validate our algorithm using census data for sub-Saharan Africa from Gershman and Rivera (2018). Once we have information on the speakers of different languages at a 5 km by 5 km resolution, we compute our measures of global fractionalization and the local-global complementarity for each country. The underlying assumption is that local interaction occurs within cells of 5 km by 5 km.³

We then explore the effect of global ethnolinguistic fractionalization and local-global ethnolinguistic complementarity on a wide variety of public goods outcomes in health, education and infrastructure at the country level. We find that global ethnolinguistic fractionalization is negatively associated with public good provision, whereas the opposite holds for local-global ethnolinguistic complementarity. This suggests that local interaction mitigates the antagonism felt towards other groups. The magnitudes of the effects are considerable. For example, a one standard deviation increase in local-global ethnolinguistic complementarity is associated with an increase in child survival by 0.75 per hundred live births. To put this figure into perspective, in its relation to child mortality, a one standard deviation increase in local-global ethnolinguistic complementarity is equivalent to a 48 percent increase in GDP per capita. As another example, a one standard deviation increase in local-global ethnolinguistic complementarity is associated with an increase in literacy by 6.2 percentage points. This corresponds to a standardized β of 30 percent, i.e., a one standard deviation increase in local-global ethnolinguistic complementarity is associated with an increase in the literacy rate by 30 percent of its standard deviation. Reverse causality is a potential concern: in societies with poor public goods individuals from the same linguistic group may prefer to geographically cluster to support each other. In Appendix D we propose several ways of addressing this concern.

Our paper contributes to four strands of the literature. First, a few papers have analyzed how the spatial distribution of diversity matters for different outcomes at the country level. In an unpublished working paper, Matuszeski and Schneider (2006) show that civil war is more likely in societies where ethnic groups

³Note that we do not distinguish between contact and exposure. For a discussion on the difference between both concepts, see, for example, Finseraas et al. (2016).

are more clustered. However, their algorithm to assign speakers of different languages to geographic cells ignores the existence of widespread languages, thus introducing a bias in many countries of the New World. Also related is the empirical paper by Alesina and Zhuravskaya (2011) who find that geographic segregation has a negative effect on the quality of government. In contrast, a recent paper by Hodler, Valsecchi and Vesperoni (2017) that uses a more general segregation index concludes that outcomes such as governance and economic development improve when ethnically more distant groups live more separated in space.⁴ Finally, in an empirical study of sixteen African economies, Robinson (2013) shows that local diversity increases inter-ethnic trust at the country level. Different from our work, these papers measure the impact of local interaction by either using a segregation index or a local fractionalization index. These indices differ from ours in that they do not account for the complementarity between local shares and aggregate shares. As we will show, to the extent that local interaction with other groups only affects overall antagonism if those other groups are sizable in the aggregate, this complementarity is key.

Second, some studies look at the relation between diversity and different political economy outcomes at the local level. In the context of local public goods, examples include Alesina, Baqir and Easterly (1999) at the level of cities in the U.S.; Dahlberg et al. (2012) at the level of Swedish municipalities; Munshi and Rosenzweig (2015) at the level of wards in India; and Algan et al. (2016) at the level of apartment blocks in France. Also noteworthy is the paper by Montalvo and Reynal-Querol (2016) who analyze the effect of local diversity on local economic growth at a spatial resolution of one degree by one degree. Rather than analyzing the effect of local diversity on local outcomes, our paper explores how local interaction affects outcomes at the national level.

Third, a number of experimental papers have explored whether contact tends to reduce prejudice, as advocated by contact theory (Allport, 1954). For example, Burns, Corno and La Ferrara (2015) show how living with a randomly assigned roommate from a different ethnicity at the University of Cape Town reduces prejudice, and Carrell et al. (2015) find that white males in the U.S. Air Force Academy are more accepting of blacks if they are exposed to a greater number of black peers. Similarly, Boisjoly et al. (2006) conclude that white students who are randomly assigned African American roommates are more likely to endorse affirmative action and to be more empathetic towards other groups. In an overview of this literature, Bertrand and Duflo (2016) conclude that the overwhelming evidence is in favor of contact theory.⁵ Rather than focusing on small-scale micro-evidence of the relation between interaction and prejudice, we conduct a large-scale analysis of the relation between local interaction and public goods provision at the national level across the globe. In doing so, we confirm the findings of the majority of this experimental literature.

Fourth, several papers have constructed datasets of language use at the local level. Like us, they

⁴Tajima, Samphantharak and Ostwald (2018) find a similar result for public goods in the context of Indonesia.

⁵This echoes the meta-analysis by Pettigrew and Tropp (2006) who find that 94 percent of 713 independent samples of 515 studies support contact theory. Counter-examples include Putnam (2007) who finds a negative relation between diversity and social capital in U.S. communities, Enos (2014) who documents more exclusionary attitudes towards the out-group when a small number of Spanish-speaking people are sent to commuter trains in homogeneously white communities, and Condra and Linardi (2018) who report that in post-conflict Afghanistan contact with the Pashtun group makes non-Pashtuns less altruistic towards the out-group.

combine linguistic maps with information on the number of speakers at the country level. Weidman et al. (2010) rely on the Atlas Narodov Mira to create the GREG (Geo-Referencing of Ethnic Groups) database, whereas Matuszeski and Schneider (2006) use maps from the Ethnologue with the same goal.⁶ A recurring issue is how to allocate the shares of different language speakers in areas where more than one language is spoken. Our paper tackles this problem by using an iterative proportional fitting algorithm, commonly used in statistics. Unlike the methods used in other papers that allocate language speakers to cells, this algorithm respects both the cell populations and the country-language population totals. Another approach is to use census and survey data. Gershman and Rivera (2018) follow this method to estimate local diversity measures for around 400 first-level administrative regions in 36 countries of sub-Saharan Africa. To cross-validate our algorithm, we compare our local diversity measures to the ones in Gershman and Rivera (2018). Focusing on the regions for which Gershman and Rivera (2018) use census data, we find correlations in local diversity of 0.80 at the regional level and 0.95 at the country level. These high correlations for sub-Saharan Africa give us confidence in our algorithm. Compared to using census or survey data, our approach has the important advantage that we can generate a consistent dataset of local language use for the entire world.

The rest of the paper is organized as follows. Section 2 develops a conceptual framework of antagonism and shows how it relates to the provision of public goods. Section 3 discusses the data, and provides details of the algorithm to assign languages and populations to all 5 km by 5 km cells in the world. Section 4 empirically explores the relation between local interaction and public goods provision. Section 5 concludes.

2 Measurement Framework

A country is geographically partitioned into K cells, indexed by ℓ or k . Each cell is small relative to the country. There are N individuals in the country, partitioned into M ethnolinguistic groups, indexed by i or j . Denote by $s_{\ell i}$ the share of people of location ℓ who belong to ethnolinguistic group i , by s_{ℓ} the share of individuals who live in cell ℓ , and by s_i the share of individuals who belong to ethnolinguistic group i . Each individual belongs to one ethnolinguistic group i and lives in one cell ℓ , so that $\sum_i s_{\ell i} = 1$, $\sum_{\ell} s_{\ell} = 1$ and $\sum_i s_i = 1$. An individual who speaks language i and resides in ℓ has preferences over private consumption c and public consumption G of the form

$$U_{\ell i}(c, G) = \ln c + v_{\ell i} \ln G, \tag{1}$$

where the valuation parameter $v_{\ell i}$ depends negatively on the antagonism the individual feels toward others in society. Public consumption G is common to the entire country, and is hence not cell-specific. Examples may include national education, nation-wide health policies and defense. Individuals only differ from each other by where they reside and which language they speak. In particular, all individuals have the same

⁶One important difference is that the Ethnologue has information on 6,905 distinct languages, whereas the GREG data has 929 groups.

income y , which we normalize to 1.⁷ In what follows we start by relating antagonism to local interaction, and then connect it to the valuation of public goods.⁸

2.1 Antagonism

Local interaction and overall antagonism. As in the framework of Esteban and Ray (1994), we derive a measure of antagonism in society based on the average antagonism of all individuals. Each individual feels a default antagonism of 1 towards people of other ethnolinguistic groups and of 0 towards people of her own group. An individual's default antagonism towards people of other groups in the society at large can be mitigated or reinforced by local interaction with these other groups. To be more precise, the antagonism an individual of group i and cell ℓ feels towards an individual of group j in the society at large is given by

$$1 + \beta s_{\ell j}, \quad (2)$$

where the share $s_{\ell j}$ captures the degree of interaction with people of group j in his own cell ℓ . We are agnostic about the sign of β : a positive value is consistent with interaction increasing antagonism, whereas a negative value means that local interaction with group j mitigates the antagonism an individual of group i feels towards an individual of group j in the society at large.⁹ A reduction in antagonism would be consistent with contact theory which argues that interaction reduces prejudice.¹⁰ While we take β to be a constant, a more comprehensive model would allow for the value of β to be situation-dependent.¹¹

Starting from (2), we can now compute the share-weighted antagonism felt by an individual of group i and cell ℓ towards all individuals of group j in society:

$$a_{\ell ij} = s_j(1 + \beta s_{\ell j}) = s_j + \beta s_j s_{\ell j}, \quad (3)$$

where the first term can be interpreted as the default antagonism in the absence of local interaction and the second term can be interpreted as the mitigating or reinforcing effect of local interaction. The average

⁷Some of the robustness tests in the empirical part of the paper will control for overall income inequality and for income inequality between ethnic groups.

⁸For alternative models of the relation between ethnic diversity and public goods, see, e.g., Miguel and Gugerty (2005) which focus on local public goods and the role of social sanctions, and Alesina and La Ferrara (2000) who look at the relation between diversity and group formation.

⁹An alternative assumption would be that interaction with *anyone* who is not from the own group affects antagonism towards *anyone* who is not from the own group. In that case, expression (2) would become $1 + \beta(1 - s_{\ell i})$. We briefly discuss this alternative in the empirical part.

¹⁰The psychological mechanisms underlying the contact hypothesis include learning about others and the ability to empathize with other groups. Going beyond contact theory, there may be other reasons why local interaction reduces antagonism. For example, if individuals locally live with other groups, there may be negative spillovers if those other groups are unhappy or do not have adequate access to healthcare. The realization that having bad relations with other groups is costly may lower antagonism towards the out-group. Under this interpretation attitudes towards other groups become more inclusive for materialistic reasons. For a similar example in a different context, see Lizzeri and Persico (2004).

¹¹As argued by Allport (1954), contact reduces prejudice only if groups collaborate as equals in the pursuit of a common goal. The nature of inter-group interaction may also depend on historical contingencies and institutions. For example, Sambanis and Shayo (2013) argue that effective nation building efforts and institutions that limit the contestability of resources make it less likely for different groups to engage in violent conflict.

antagonism of all individuals of cell ℓ towards all other individuals in society is then:

$$a_\ell = \sum_i s_{\ell i} \left(\sum_{j \neq i} s_j (1 + \beta s_{\ell j}) \right).$$

Taking the population-weighted average across all cells yields a measure of the average antagonism in society:

$$A = \sum_\ell s_\ell \left(\sum_i s_{\ell i} \sum_{j \neq i} s_j (1 + \beta s_{\ell j}) \right),$$

which can be re-written as

$$\begin{aligned} A &= \sum_i \sum_{j \neq i} s_i s_j + \beta \sum_\ell s_\ell \sum_i \sum_{j \neq i} s_{\ell i} s_{\ell j} s_j \\ &= ELF_{glob} + \beta \sum_\ell s_\ell \sum_i \sum_{j \neq i} s_{\ell i} s_{\ell j} s_j \\ &= ELF_{glob} + \beta LGC, \end{aligned} \tag{4}$$

where the first term is the probability that two randomly drawn individuals from the society at large belong to different groups, and the second term is the probability that when an individual 1 is randomly matched with an individual 2 of his own cell and an individual 3 from the society at large, 2 and 3 belong to the same group and 1 belongs to a different group.

The first term is of course nothing else than the well-known ethnolinguistic fractionalization index. Since it measures the overall fractionalization of a country, we refer to this term as *global ethnolinguistic fractionalization* and denote it by ELF_{glob} . The second term can be thought of as the average effect of local interaction on the overall antagonism experienced towards others in the society at large. We refer to this term as *local-global ethnolinguistic complementarity* and denote it by LGC . We use this terminology because of the complementarity between the local group shares, $s_{\ell j}$, and the aggregate group shares, s_j , in the second term. The intuition for this complementarity is easy to understand: how much an individual interacts with other groups depends on the local shares; and how much that interaction affects her antagonism towards those other groups in the society at large depends on the aggregate shares. That is, interacting with another group locally can only have an important effect on antagonism if that other group is sizable in the aggregate. As we will see, this complementarity is key when comparing our index to other indices. Note furthermore that if we were to ignore the effect of local interaction, then β would be zero, and overall antagonism would simply be equal to ELF_{glob} . This has been the standard assumption in much of the literature on the effect of ethnolinguistic fractionalization on different political economy outcomes.

Relation to other indices. We now briefly discuss the relation of our index to measures of segregation, polarization and local fractionalization.

Relation with segregation

Segregation measures the extent to which cells differ from each other, whereas local-global complementarity

measures how local interaction affects antagonism in the society at large.¹² To see the difference between both indices, take two countries, one with language shares $(1/3, 1/3, 1/3)$ and the other with language shares $(0.5, 1/6, 1/6, 1/6)$, so that ELF_{glob} is the same in both countries. There is perfect geographic mixing, so that the linguistic composition of each cell mirrors that of its country. As a result, segregation is zero in both countries. However, it is easy to show that the country with shares $(1/3, 1/3, 1/3)$ has a higher local-global complementarity than the country with shares $(0.5, 1/6, 1/6, 1/6)$. The intuition is straightforward: interacting with fewer, but bigger, groups has a greater impact on antagonism than interacting with more, but smaller, groups. Because of the complementarity between local group shares and aggregate group shares, local-global complementarity is bigger in societies with fewer, larger groups. In highly fractionalized societies, with many small groups, the effect of local interaction on overall antagonism is limited, even if societies are spatially mixed. As we will see in the data, this difference between segregation and local-global complementarity is not just theoretically relevant. Controlling for global fractionalization, the correlation between the residuals of both indices stands at -0.63. As expected, the correlation is negative. However, it is far from perfect, implying both concepts are empirically different.

Relation with polarization

The functional form of local-global complementarity is related to that of polarization, though their micro-foundations are very different. The comparison between both indices is sharpest when focusing on a society that is spatially perfectly mixed. In that case $s_{\ell i} = s_i$ for all ℓ and all i , yielding a local-global complementarity effect equal to $\sum_i \sum_{j \neq i} s_i s_j^2$. Esteban and Ray (1994) micro-found an index of polarization by positing that the antagonism an individual experiences towards others increases in the size of his own group. The idea is that an individual feels greater antagonism if he identifies more strongly with his own group, and this happens when the size of his group is bigger. As a result, instead of the standard fractionalization index, $\sum_i \sum_{j \neq i} s_i s_j$, which ignores the role of identification, they obtain a polarization index, $\sum_i \sum_{j \neq i} s_i^2 s_j$.¹³

As can be seen, in the case of perfect spatial mixing, local-global complementarity is identical to polarization, with one subtle difference: in the polarization index, the quadratic term shows up in the own-group share, reflecting antagonism increasing in the identification with the own group, whereas in the local-global complementarity index, the quadratic term shows up in the other-group share. This suggests a possible interpretation of our index in terms of identification: local interaction with another group leads to identification with that other group, which then affects the antagonism felt towards that group in the society at large. Under that interpretation, identification with the own group underlies the polarization index, whereas identification with the other group underlies local-global complementarity. In addition to the different micro-foundations, one other difference between the two indices is important: whereas polarization

¹²For a survey on segregation indices, see Massey and Denton (1988), and for an application to the quality of government, see Alesina and Zhuravskaya, (2011).

¹³The original index by Esteban and Ray (1994) is slightly more general. This is the specific measure of polarization used subsequently by Reynal-Querol (2002).

is commonly viewed as being detrimental to economic, political and social outcomes, we are agnostic about local interaction having a benign or a detrimental effect. Moreover, when groups are not perfectly mixed across a country's geography, the equivalence in the functional forms of polarization and local-global complementarity breaks down.

Relation with local fractionalization

Our measure of average antagonism is a linear combination of global fractionalization and local-global complementarity. Other papers have taken the view that what might matter is *global fractionalization* and average *local fractionalization* (Matuszeski and Schneider, 2006; Robinson, 2013). In our framework this would amount to ignoring the complementarity between local and global group shares by supposing that the antagonism an individual of group i and cell ℓ feels towards people of group j in the society at large is $a_{\ell ij} = s_j + \beta s_{\ell j}$ rather than $a_{\ell ij} = s_j(1 + \beta s_{\ell j}) = s_j + \beta s_j s_{\ell j}$.

In spite of its appealing simplicity, this alternative index based on local fractionalization does not appropriately capture the basic premise that local interaction can only impact overall antagonism if there is antagonism in the first place. To see this, consider a situation where $s_j \approx 0$ and $s_{\ell j} \approx 1$, and assume $\beta < 0$. Since no one outside cell ℓ is from group j , we would expect the antagonism experienced by individuals of group i and cell ℓ towards individuals from group j to be essentially zero. Consistent with this, using our index, the tension felt by individuals of group i and cell ℓ towards the out-group would be zero. However, the alternative index based on local fractionalization would imply that the antagonism experienced by individuals of group i and cell ℓ towards individuals of group j is β . That is, local fractionalization reduces antagonism even if there is no antagonism to start with. This happens because local fractionalization enters additively into the overall expression of antagonism. In contrast, in our corresponding measure of antagonism (3), the local-global complementarity index enters multiplicatively in the antagonism felt toward the out-group, so that the change in antagonism is always a fraction of overall antagonism.

2.2 Public Goods Provision

As shown in equation (1), we assume that the valuation $v_{\ell i}$ that an individual from ethnolinguistic group i and cell ℓ attaches to the public good G is a negative function of the antagonism he feels. The average antagonism experienced by an individual of group i living in cell ℓ towards the rest of society is

$$a_{\ell i} = \sum_{j \neq i} s_j (1 + \beta s_{\ell j}).$$

Note that since β can be negative, $a_{\ell i}$ need not be positive. Hence, we postulate

$$v_{\ell i} = \frac{\kappa_1}{\kappa_2 + a_{\ell i}}, \tag{5}$$

where $\kappa_1 \geq 1$ and $\kappa_2 \geq 0$ is large enough so that $v_{\ell i} > 0$.

Suppose that G is determined by private contributions.¹⁴ In particular, consider a simultaneous

¹⁴Later we will briefly discuss that a similar result can be derived in a model where G is determined by a democratic vote.

private contribution game where the equilibrium concept is Nash. We denote by $g_{\ell i}$ the contribution of an agent living in ℓ of group i . The total level of the public good is then

$$G = \sum_{\ell} \sum_i g_{\ell i}.$$

An agent with valuation $v_{\ell i}$ chooses his contribution $g_{\ell i}$ by solving

$$\begin{aligned} \max_{g_{\ell i}} & \ln(1 - g_{\ell i}) + v_{\ell i} \ln(G_{-\ell i} + g_{\ell i}) \\ \text{s.t.} & \quad 1 \geq g_{\ell i} \geq 0 \end{aligned}$$

where $G_{-\ell i}$ is the contribution of the rest of agents.

Proposition 1 *Suppose that in the Nash equilibrium of the contribution game all agents contribute a strictly positive amount. Then we have that G is a decreasing function of total average antagonism A . In particular,*

$$G = \frac{\kappa_1 N}{\kappa_1 + N(\kappa_2 + A)}. \quad (6)$$

Proof. See Appendix A. ■

For our empirical estimation it will be useful to linearly approximate (6). After dividing numerator and denominator by $\kappa_1 N$ and assuming that $1/N \approx 0$, we can write

$$G \approx \frac{\kappa_1}{\kappa_2 + A} \quad (7)$$

Assuming that κ_2 is sufficiently large and that A is sufficiently small, we can take a first-order Taylor approximation of (7) which yields

$$G \approx \frac{\kappa_1}{\kappa_2} - \frac{\kappa_1}{\kappa_2^2} A = \frac{\kappa_1}{\kappa_2} - \frac{\kappa_1}{\kappa_2^2} ELF_{glob} - \frac{\kappa_1}{\kappa_2^2} \beta LGC. \quad (8)$$

This will serve as our estimating equation in the empirical part. From (8) we can conclude that the provision of public goods depends negatively on global fractionalization and either positively or negatively on local-global complementarity (positively if $\beta < 0$ and negatively if $\beta > 0$). This theory therefore implies that one should distinguish between global fractionalization and local-global complementarity when empirically exploring the relation between diversity and public goods provision.¹⁵

3 Data

In this section we describe the different data we use in our empirical analysis. After briefly discussing the data on public goods, we mainly focus on how we construct our novel database of language use at the local level for the entire globe.

¹⁵If instead of a private contributions mechanism, a democratic vote decides the public good, the exact same result would hold if the mean agent coincides with the median agent. The result would still hold qualitatively if the median agent and the mean agent are not too different. In this context the mean agent refers to the agent with the mean valuation of the public good.

3.1 Public Goods

When analyzing the effect of diversity on the provision of public goods, we do not focus on a particular public good. Instead, we look at many different measures, related to health, education and infrastructure. In doing so, we build on previous work by La Porta et al. (1999), Alesina et al. (2003) and Desmet, Ortuno-Ortin and Wacziarg (2012). The exact variables, their sources and summary statistics are given in Appendix B and Table B.1. Our analysis is done at the country level.

3.2 Spatial Distribution of Languages

To compute the local-global ethnolinguistic complementarity index for all countries of the world, we need to know how many people speak each language at the local level. To that end, we start by splitting up the world in grid cells. Since our index depends on the degree of personal interaction in people’s daily lives, we need grid cells of a size that captures this daily interaction. In our baseline analysis we take a grid with a resolution of 5 km by 5 km to be a reasonable size. In what follows we explain how we allocate the speakers of the world’s different languages to the individual grid cells.

We use two main data sources. The first data source is the World Language Mapping System, the digitized version of the 17th edition of the Ethnologue. This gives us a polygon shapefile, where most of the 6,905 languages are represented as polygons across space. Since in some areas more than one language is spoken, these polygons may overlap. In addition, when certain widely spoken languages in a country cannot be assigned to any particular geographic region, Ethnologue classifies them as widespread languages. This is equivalent to having a polygon that consists of the entire country.¹⁶ A few languages are assigned to specific points, rather than to polygons, and a few others have unknown locations. For the point languages, we create circular polygons around the points, with a radius that is proportional to the number of speakers of those languages.¹⁷ As for the languages with unknown locations, we treat them as widespread. Some areas, such as the sparsely populated Sahara Desert, have no information on languages. In those cases we assign the language of the nearest cell that has information on language. Since we use grid cells of 5 km by 5 km, we rasterize the data using a resolution of 2’30’’ by 2’30’’. In addition to the linguistic polygons, the Ethnologue also provides the number of people that speak the different languages by country.

The second data source is Landscan which provides population at a resolution of 30’’ by 30’’. Here as well, we rasterize the data using a resolution of 2’30’’ by 2’30’’. To make the language data from the Ethnologue consistent with the population data from Landscan, we normalize the language data for each country so that the sum of a country’s different language speakers equals the country’s total population according to the World Bank.

These two data sources yield three pieces of information: the number of people per grid cell; the number of speakers of each language per country; and whether a language is spoken or not in a given grid

¹⁶Excluding widespread languages, there is a maximum of seven overlapping polygons.

¹⁷We ignore point languages that account for less than 0.5% of the population.

cell. What it does not tell us is how many people speak each language in each cell. Hence, using these three pieces of information, we need to allocate language speakers to grid cells. To that end, we use an iterative proportional fitting algorithm, commonly used in statistics, which we now describe in further detail.¹⁸ We later cross-validate the outcome of our algorithm with actual census data from Gershman and Rivera (2018).

The iterative proportional fitting algorithm is a way of allocating language speakers to cells, such that the total population per cell and the total population per language correspond to their actual values. Consider a country that has M linguistic groups and is split up into K cells. Using the two data sources, we construct three matrices that correspond to the three pieces of information we referred to above: $\mathcal{N}$ is a $K \times 1$ matrix of which the elements give the total population of each cell; $\mathcal{L}$ is a $1 \times M$ matrix of which the elements give the number of speakers of each language in the country; $\mathcal{B}$ is a $K \times M$ binary matrix of which the elements take a value 1 if the language corresponding to the column is spoken in the cell corresponding to the row (and a value 0 otherwise). The iterative proportional fitting algorithm then goes through the following steps.

1. *Step 0.* Define $\mathcal{T}^{(0)} = \mathcal{B}$.

2. *Step 1.* For each location ℓ , assign a share $\mathcal{T}^{(2n-2)}(\ell, i) / \sum_j \mathcal{T}^{(2n-2)}(\ell, j)$ to language i . Hence,

$$\mathcal{T}^{(2n-1)}(\ell, i) = \frac{\mathcal{T}^{(2n-2)}(\ell, i)}{\sum_j \mathcal{T}^{(2n-2)}(\ell, j)} \mathcal{N}(\ell, 1), \quad (9)$$

where $n = 1, 2, \dots$ refers to the times the algorithm has iterated through Step 1 and Step 2. To provide some intuition, the first time the algorithm goes through Step 1, the cell's population gets divided equally between the different languages that are spoken there. If, for example, 5 languages are spoken in a cell, then each language gets assigned 20 percent of that cell's population. This allocation always ensures that $\sum_j \mathcal{T}^{(2n-1)}(\ell, j) = \mathcal{N}(\ell, 1)$, i.e., the sum of people allocated to each cell is equal to the actual population of each cell. That is, the allocation satisfies the marginals on the cell populations.

3. *Step 2.* For each language i , assign a share $\mathcal{T}^{(2n-1)}(\ell, i) / \sum_k \mathcal{T}^{(2n-1)}(k, i)$ to cell ℓ . Hence,

$$\mathcal{T}^{(2n)}(\ell, i) = \frac{\mathcal{T}^{(2n-1)}(\ell, i)}{\sum_k \mathcal{T}^{(2n-1)}(k, i)} \mathcal{L}(1, i) \quad (10)$$

This allocation always ensures that $\sum_k \mathcal{T}^{(2n)}(k, i) = \mathcal{L}(1, i)$, i.e., the sum of population allocated to a language is equal to the actual total number of speakers of that language. That is, the allocation satisfies the marginals on the language populations.

4. *Step 3.* Go through Step 1 and Step 2 until $\mathcal{T}^{(2n-1)}(\ell, i)$ converge to $\mathcal{T}^{(2n)}(\ell, i)$ for all ℓ and i .

This iterative proportional fitting algorithm therefore provides us with an allocation of language speakers by cell, $\mathcal{T}^{(2n)}(\ell, i)$. If the three matrices $\mathcal{L}$, $\mathcal{N}$ and $\mathcal{B}$ are fully consistent with each other, then the

¹⁸See, e.g., Deming and Stephan (1940) and Bishop, Fienberg and Holland (1975).

iterative proportional fitting algorithm is guaranteed to converge. However, there may be small inconsistencies between the data sources. As a simple example, it is possible that the polygon assigned to language i has a population that is smaller than the total population that speaks language i . This inconsistency could in principle be due to three reasons: the local population data from Landscan may contain imprecisions; the country-level language shares from the Ethnologue might have errors; or the language polygons from the Ethnologue may not be a completely accurate reflection of where the different languages are spoken. How we deal with these minor inconsistencies requires taking a stance on the most likely source of error. We take the view that local population and language shares are relatively easy to estimate, whereas language polygons are unlikely to be completely precise. For example, although most Catalan speakers in Spain live in the East of the country, a small percentage of them live in other parts of the country. However, since Catalan is not a “widespread” language in the sense that is not widely spoken across the entire territory, the Ethnologue assigns a specific polygon to Catalan. Given the binary nature of such a geographic polygon, it hence assumes that all Catalan speakers reside within the polygon. Since this is obviously an approximation, we replace the 0 values in the binary matrix $\mathcal{B}$ by 0.000001. This amounts to allowing some speakers of language i to live outside their corresponding language polygon. With this correction, the iterative proportional fitting algorithm is once again guaranteed to converge (Fienberg, 1970).

3.3 Global Fractionalization and Local-Global Complementarity Measures


With the spatial distribution of languages at a resolution of 5 km by 5 km in hand, we can now compute, for each country, our measures of global ethnolinguistic fractionalization and local-global ethnolinguistic complementarity. For comparison purposes, we also compute measures of segregation and local fractionalization.

When measuring linguistic fractionalization, it is not always obvious which linguistic groups should be used as primitives. For example, should Walloon and Picard, two variations of French, be considered as two separate language groups or should they be aggregated into French? To address this issue, Desmet, Ortuño-Ortín and Wacziarg (2012) use the language tree of the Ethnologue to compute measures of linguistic fractionalization at different levels of aggregation. There are 15 possible levels of aggregation, going from the most aggregate at level 1, where only the big language families, such as Indo-European and Niger-Congo are considered to be different groups, to the most disaggregate at level 15, where Walloon and Picard are taken to be different groups. As Desmet, Ortuño-Ortín and Wacziarg (2012) argue, coarse divisions, obtained at high levels of aggregation, can be thought of as cleavages that go back far in history, whereas finer divisions, obtained at low levels of aggregation, are due to more recent cleavages. They show that certain political economy outcomes, such as conflict, are better explained by measures of linguistic fractionalization at high levels of aggregation, indicating that they have to do with deep cleavages. In contrast, other outcomes, such as economic growth, are better explained by measures of linguistic fractionalization at low levels of aggregation, suggesting that they depend on more shallow cleavages.

In the case of public goods, Desmet, Ortuño-Ortín and Wacziarg (2012) find that intermediate levels of aggregation are most relevant. Hence, in our benchmark empirical analysis of public goods, we aggregate

languages to level 5, though we also check the robustness of our results to using both higher and lower levels of aggregation. As an illustration, level 5 of aggregation implies that Spanish, Catalan and Portuguese are aggregated into the same group, but French and Italian are not. Similarly, Hindi and Urdu are considered to be in the same group. As another example, in Tanzania 104 out of the 129 languages are aggregated into the same group (Niger-Congo/Atlantic-Congo/Volta-Congo/Benue-Congo/Bantoid/Southern).

Figure 1. Global ELF by Country – Level 5



Using aggregation level 5, Figure 1 and Figure 2 show global ethnolinguistic fractionalization and local-global ethnolinguistic complementarity by country.¹⁹ As is immediately obvious, there are many differences between both indices. Compare, for example, Chad and the Central African Republic. Global ethnolinguistic fractionalization is much higher in Chad, but local-global ethnolinguistic complementarity is much higher in the Central African Republic. As another example, Malaysia and Indonesia have similar levels of global ELF, but local-global ethnolinguistic complementarity is much higher in Malaysia. Figure 3 also displays local-global ethnolinguistic complementarity, but now for each grid cell separately. Consistent with the previous examples, there are more areas of high local-global ethnolinguistic complementarity in the Central African Republic and Malaysia than in Chad and Indonesia.²⁰

The left-hand panel of Figure 4 shows a scatterplot of both indices. What stands out is that countries with very high degrees of global fractionalization tend to have low local-global ethnolinguistic complementarity. The most fractionalized countries, such as Papua New Guinea (PNG), Mali (MLI) and

¹⁹Table B.1 give summary statistics of the different indices and Table B.2 reports correlations between them.

²⁰Appendix Figures C.1 through C.3 display the same maps, but for level 15 of aggregation. Compared to level 5, some of the countries in the central and the southern part of Africa now display much higher levels of diversity. For example, Zambia, where nearly everyone speaks a language of the Niger-Congo family, is much more diverse at level 15 than at level 2. Appendix Figures C.4 through C.6 also show maps for level 2 of aggregation.

Figure 2. Local-Global Ethnolinguistic Complementarity by Country – Level 5






Figure 3. Local-Global Ethnolinguistic Complementarity at 5 km by 5 km Resolution – Level 5



Nigeria (NGA), have so many groups that local-global ethnolinguistic complementarity is limited, independently of how much the groups are locally mixed. Another relevant finding is the important heterogeneity in the degree of local-global ethnolinguistic complementarity between countries with intermediate levels of global fractionalization. Compare, for example, Guatemala (GTM) and Mauritius (MUS). In spite of overall diversity being essentially identical in Guatemala (0.53) and Mauritius (0.52), local interaction is lower in Guatemala. The reason is that in Guatemala many of the speakers of indigenous languages live in the central and northwestern highlands, having limited contact with Spanish speakers. In contrast, according to Chiba (2006), Mauritians “switch languages according to the occasion in the way other people change clothes”. As a result, local-global ethnolinguistic complementarity is much higher in Mauritius (0.20) than in Guatemala (0.06).

Figure 4. Global Fractionalization, Local-Global Complementarity, and Segregation



In the theoretical section we already addressed the relation between local-global complementarity and segregation. The right-hand side panel of Figure 4 plots their empirical relation. For segregation, we compute the segregation index of Alesina and Zhuravskaya (2011) using our data at the 5 km by 5 km level. Comparing both indices is cleaner when controlling for global fractionalization, so the scatterplot represents residuals of regressions on global fractionalization. As expected, local-global ethnolinguistic complementarity tends to be greater in less segregated societies. However, the correlation is far from perfect, standing at -0.63 . This suggests that there are important differences between both indices. Mali (MLI) is a case in point: there is a fair amount of local mixing, so that segregation is relatively low. However, because the country is made up of many different groups, the degree of local-global complementarity is limited.

3.4 Cross-Validation

To assess the quality of our iterative proportional fitting algorithm that allocates language speakers to grid cells, we compare local fractionalization measures based on our allocation to ones obtained by Gershman and Rivera (2018). Rather than using maps, they rely on national population censuses and regional household

surveys to infer the linguistic composition of almost 400 first-level administrative regions in 36 sub-Saharan African countries. They then use the language tree from the Ethnologue to compute local fractionalization measures for these regions at different levels of linguistic aggregation. To compare our measures to the ones in Gershman and Rivera (2018), we start by aggregating our 5 km by 5 km language allocation up to the same first-level administrative regions, and then calculate for each of these regions a measure of local fractionalization.

At linguistic aggregation level 5, we find a correlation between our measure of local fractionalization and the one in Gershman and Rivera (2018) of 0.77 at the level of first-level administrative regions. The local fractionalization measures in Gershman and Rivera (2018) that are based on census data are arguably of higher quality than those that are based on survey data. If we limit ourselves to first-level administrative regions for which local fractionalization comes from census data, the correlation with our index increases to 0.80. Since our empirical analysis uses population-weighted local indices at the country level, it is useful to compare country-level local fractionalization indices. When doing so, the correlation between our measure and the one based on Gershman and Rivera (2018) further increases to 0.95 for countries with census data. It is encouraging to find such high correlations, in spite of the differences in the underlying data on linguistic groups.²¹ This cross-validation gives us confidence in the algorithm we use to allocate language speakers to cells.

4 Empirical Analysis

In this section we test our theory by exploring the impact of global ethnolinguistic fractionalization and local-global ethnolinguistic complementarity on a country's provision of public goods. Equation (8) from the theory provides the following estimating equation:

$$G_c = \gamma_0 X_c + \gamma_1 ELF_{glob,c} + \gamma_2 LGC_c + \varepsilon_c, \quad (11)$$

where G_c is the level of public goods in country c , $ELF_{glob,c}$ is the global fractionalization measure of country c , LGC_c is the local-global complementarity of country c , X_c are a set of additional controls, ε_c is an error term, $\gamma_0 = \kappa_1/\kappa_2$, $\gamma_1 = -\kappa_1/\kappa_2^2$, and $\gamma_2 = -(\kappa_1/\kappa_2^2)\beta$. Our main coefficient of interest will be γ_2 : if its sign is consistent with an improvement in public goods, we will interpret this as implying that local interaction reduces antagonism towards other groups, making individuals more willing to contribute to public goods. As mentioned in the previous section, in the benchmark we use grid cells of 5 km by 5 km to compute LGC_c and a linguistic aggregation level of 5 to compute both $ELF_{glob,c}$ and LGC_c .

²¹The Ethnologue is much more detailed in the number of linguistic groups than Gershman and Rivera (2018). For example, whereas Mali has 54 unique linguistic groups in the Ethnologue, it only has 9 unique groups in the Gershman and Rivera (2018) data.

4.1 Local-Global Ethnolinguistic Complementarity and Public Goods

In our baseline analysis we focus on child survival. It measures the survival rate under age 5 (per 100 live births), and captures well the effectiveness of public goods provision. We then extend our analysis to include a variety of additional measures. In particular, we focus on two more health outcomes (hospital beds per 1000 people, rate of measles immunization), two education outcomes (literacy rate, log of average years of schooling), and two infrastructure measures (percentage of households with access to improved sanitation and km of roads per 1,000 people).²²

Child survival. Table 1 analyzes the relation between global ethnolinguistic fractionalization, local-global ethnolinguistic complementarity and child survival, using OLS. The first two columns follow the standard approach of the literature and only control for a country’s global fractionalization. The first specification includes regional dummies and latitude as covariates, whereas the second specification is identical to the one in Desmet, Ortuño-Ortín and Wacziarg (2012) and also controls for legal origin and GDP per capita. Consistent with previous papers, we find that an increase in a country’s level of fractionalization is associated with worse outcomes. The coefficients on global fractionalization are statistically significant at the 1% level. The next two columns replace global fractionalization by local-global complementarity. The coefficients switch signs, suggesting that a higher local-global ethnolinguistic complementarity is associated with more child survival, but the coefficients are not statistically significant at the 10% level.

Motivated by the conceptual framework, in the last two columns we include simultaneously global fractionalization and local-global complementarity. In the rest of the paper, we will use the specification in column (6) as our baseline.²³ As before, countries with higher levels of global fractionalization continue to have lower rates of child survival. In fact, the coefficients are larger in magnitude. Focusing on column (6), we find that a one standard deviation increase in global fractionalization is associated with a decrease in child survival by 1.13 per hundred live births. More interestingly, the local-global complementarity coefficients are now statistically significant at the 1% level. Their sign is positive, indicating that local-global ethnolinguistic complementarity is associated with an improvement in child survival. The economic magnitude of the effect is substantial. Using once again column (6) as our preferred specification, a one standard deviation increase in local-global ethnolinguistic complementarity is associated with increasing child survival by 0.75 per hundred live births. To put this figure into perspective, a one standard deviation increase in local-global ethnolinguistic complementarity has the same effect on child survival as a 48% increase in GDP per capita.

These findings are consistent with local interaction mitigating antagonism towards the out-group.

²²In addition to the seven outcomes already mentioned, we include 13 additional outcomes in the robustness section.

²³It is debatable whether or not to include GDP per capita as a regressor. While it is an important determinant of public goods outcomes, it is also affected by ethnolinguistic diversity. Following Alesina and Zhuravskaya (2011) and most of the specifications in Alesina et al. (2003) and Desmet et al. (2012), on balance we prefer to include GDP per capita in our baseline. For completeness, Tables 1–3 in the Online Appendix report all our empirical results without including GDP per capita. If anything, this makes the results stronger, in terms of both the statistical and economic significance of local-global complementarity.

Table 1. Child Survival: Global ELF and Local-Global Ethnolinguistic Complementarity

	Child Survival					
	(1)	(2)	(3)	(4)	(5)	(6)
Global ELF	-3.922*** (1.055)	-2.399*** (0.871)			-6.970*** (1.323)	-4.257*** (1.072)
Local-Global Complem.			3.593 (3.887)	1.730 (3.061)	21.447*** (4.871)	12.486*** (3.891)
Absolute Latitude	0.071*** (0.018)	0.014 (0.020)	0.087*** (0.017)	0.013 (0.020)	0.065*** (0.017)	0.012 (0.020)
Latin America & Carib.	1.237* (0.696)	0.536 (0.538)	1.778** (0.689)	0.758 (0.534)	1.019 (0.651)	0.424 (0.550)
Sub-Saharan Africa	-7.787*** (0.940)	-5.945*** (0.835)	-7.865*** (0.970)	-5.942*** (0.855)	-7.477*** (0.864)	-5.919*** (0.788)
East and S.E. Asia	0.101 (0.999)	0.349 (0.816)	0.167 (1.012)	0.336 (0.825)	0.723 (0.951)	0.714 (0.810)
Log GDP per Capita		1.676*** (0.196)		1.792*** (0.200)		1.559*** (0.185)
French Legal Origin		1.798*** (0.599)		1.393** (0.561)		1.324** (0.521)
German Legal Origin		2.397*** (0.472)		2.274*** (0.476)		1.960*** (0.419)
UK legal origin		2.061*** (0.591)		1.583*** (0.571)		1.623*** (0.540)
Constant	95.308*** (0.909)	80.886*** (1.909)	93.554*** (0.885)	79.568*** (1.877)	95.107*** (0.852)	82.156*** (1.746)
Observations	178	171	178	171	178	171
R^2	0.685	0.815	0.655	0.804	0.719	0.825

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the child survival rate per 100 live births. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation and are based on the authors' calculations. The variable definitions and data sources for each of the variables are provided in Appendix B.

A higher degree of linguistic diversity at the country level implies greater antagonism, so that people value public goods less, but that negative effect is weakened if people interact with other groups in their daily lives. Next, we analyze whether this result generalizes to a wider variety of public goods.

Other public good outcomes. Table 2 reports the results for an additional six public goods outcomes; the specification is identical to the one in column (6) of Table 1. As in the case of child survival, an increase in a country's overall ethnolinguistic fractionalization tends to be associated with worse outcomes, whereas an increase in local-global ethnolinguistic complementarity tends to be associated with better outcomes. The coefficients on global fractionalization are statistically significant at the 5% level in four out of the six outcomes, whereas the coefficients on local-global complementarity are always statistically significant at the 5% level, with the exception of road density. In terms of magnitude, a one standard deviation increase in local-global complementarity is associated with an increase in the rate of measles immunization by 4.0 percentage points. The corresponding standardized β is 26%, meaning that a one standard deviation increase in local-global complementarity is associated with an increase in the measles immunization rate by 26% of its standard deviation. In the case of literacy, the standardized β is 30%. These results are further evidence in favor of local interaction mitigating antagonism towards the out-group in the society at large.

In an additional robustness check, we further extend our analysis to all public goods outcomes with wide country coverage in the World Development Indicators. Table 4 in the Online Appendix shows

that local-global complementarity is statistically significant at the 5% level for 10 out of the 13 additional outcomes. We also construct an overall measure of public goods by taking the first principal component of all 20 outcomes (which accounts for 63 percent of the variance). Here as well, our main result is confirmed: local-global ethnolinguistic complementarity has a robust and statistically significant association with public goods outcomes.

Table 2. Other Public Goods Outcomes: Global ELF and Local-Global Ethnolinguistic Complementarity

	(1) Measles Immunization	(2) Hospital Beds	(3) Literacy Rate	(4) Schooling	(5) Improved Sanitation	(6) Road Density
Global ELF	-25.034*** (4.563)	-1.246** (0.612)	-31.697*** (5.945)	-0.244* (0.137)	-25.065*** (6.177)	1.146 (1.759)
Local-Global Compl.	66.700*** (15.851)	7.037** (3.042)	102.057*** (22.085)	1.369*** (0.493)	81.892*** (24.820)	5.387 (12.411)
Absolute Latitude	0.197** (0.087)	0.101*** (0.019)	0.301** (0.117)	0.008*** (0.003)	0.136 (0.161)	0.173*** (0.065)
Latin America & Carib.	4.500* (2.672)	0.311 (0.450)	6.256* (3.434)	0.211** (0.089)	-4.755 (4.781)	2.332 (1.770)
Sub-Saharan Africa	-9.335*** (2.973)	0.119 (0.470)	-8.128* (4.163)	-0.125 (0.109)	-25.303*** (5.206)	2.958* (1.779)
East and S.E. Asia	2.434 (3.777)	0.983 (0.903)	14.255*** (4.778)	0.085 (0.102)	-6.498 (6.484)	-1.298 (2.124)
Log GDP per Capita	1.384* (0.831)	0.257* (0.155)	4.348*** (1.137)	0.117*** (0.022)	10.189*** (1.279)	1.801*** (0.452)
French Legal Origin	2.192 (2.949)	0.491 (1.448)	-3.917 (3.495)	0.077 (0.080)	13.410*** (4.079)	-14.003* (7.885)
German Legal Origin	3.260 (2.765)	2.323 (1.483)	0.000 (.)	0.235*** (0.064)	13.573*** (3.585)	-11.608 (7.914)
UK legal origin	7.458** (3.302)	0.529 (1.454)	-0.171 (4.142)	0.251*** (0.082)	10.584*** (4.038)	-10.932 (8.094)
Constant	66.595*** (8.819)	-2.049 (1.962)	46.127*** (10.918)	0.647** (0.256)	-17.277 (14.158)	-0.657 (9.531)
Observations	171	173	138	136	171	172
R^2	0.572	0.597	0.665	0.676	0.783	0.459

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The column headings give the dependent variables for each of the columns. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation and are based on the authors' calculations. The variable definitions and data sources for each of the variables are provided in Appendix B.

4.2 Further Robustness

In what follows we explore the robustness of our findings to a variety of concerns. In the interest of space, in some cases we will focus on child survival as our main outcome of interest. Results for all other outcomes can be found in the Online Appendix.²⁴

Decentralization. In our conceptual framework the geography consists of two levels: the local grid cells, where individuals interact in their daily lives, and the country, where individuals from the different local cells decide the level of public goods that everyone in the country has access to. In reality, in some countries there may be heterogeneity in the access to certain public goods, and the financing of public goods may be decentralized. To see whether this changes our results, we use data on decentralization from Treisman (2008). Table 3, column (1), takes the benchmark specification and controls for having a federal structure, as well as

²⁴In particular, Tables 5 – 10 in the Online Appendix show the equivalent results to Table 3 for all other outcomes.

for the interaction of a federal structure with both global fractionalization and local-global complementarity. Being a federal state has no significant effect on the provision of public goods. The rest of the results are largely unchanged. In particular, global fractionalization continues to lower child survival, whereas local-global complementarity continues to have a benign effect, with both coefficients being statistically significant at the 1% level.

Table 3. Child Survival: Decentralization, Linguistic Aggregation, Income & Ethnic Inequality, Segregation

	(1) Decentralization	(2) Level 15	(3) Level 2	(4) 10 × 10 km	(5) Gini	(6) Linguistic Inequality	(7) Linguistic Segregation
Global ELF	-4.206*** (1.414)	-3.015*** (0.924)	-3.789** (1.541)	-4.307*** (1.060)	-4.811*** (1.129)	-4.018*** (1.210)	-3.549*** (1.317)
Local-Global Compl.	15.133*** (5.492)	8.322** (3.876)	9.096** (4.524)	12.814*** (3.825)	16.507*** (5.319)	11.334*** (4.083)	10.994** (4.675)
Decentralization	0.334 (0.559)						
Decentral. × Global ELF	-0.042 (2.825)						
Decentral. × Local-Global	-8.024 (14.020)						
Gini					0.004 (0.044)		
Linguistic Inequality						-1.053 (0.838)	
Linguistic Segregation							-0.929 (1.329)
Constant	81.843*** (2.040)	82.719*** (1.840)	80.750*** (1.825)	82.217*** (1.740)	81.723*** (2.190)	84.101*** (1.953)	82.506*** (1.782)
Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	149	171	171	171	124	165	161
R ²	0.850	0.819	0.812	0.826	0.842	0.827	0.833

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the child survival rate per 100 live births. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation, except in columns 2 and 3 where they are measured at, respectively, levels 15 and 2 of aggregation. In column 1 decentralization is equal to 1 if country is a federal state, and 0 otherwise. In column 4, we use the Local-Global Ethnolinguistic Complementarity variable calculated at the 10 km × 10 km spatial resolution, instead of 5 km × 5 km. In column 5, we control for the income Gini coefficient, in column 6 we control for income inequality between linguistic groups, and in column 7 we control for linguistic segregation. All regressions control for the baseline controls: Log GDP per capita, legal origins dummies, regional dummies and absolute latitude. The variable definitions and data sources for each of the variables are provided in Appendix B.

Different levels of linguistic and geographic aggregation. Desmet, Ortuño-Ortín and Wacziarg (2012) find that diversity measured at an intermediate level of linguistic aggregation is most significant for the provision of public goods. This is why our baseline analysis aggregates languages up to level 5 (out of a maximum 15). To see whether our results change for lower and higher levels of linguistic aggregation, we recompute our measures of global fractionalization and local-global complementarity for levels 15 and 2. Recall that at level 15 closely-related dialects, such as Walloon and Picard, are considered to be different languages, whereas at level 2 all Romance languages, such as French and Spanish, pertain to the same group. Columns (2) and (3) in Table 3 report our baseline regression, using level 15 and level 2. The results are virtually identical; we only see a small drop in the statistical significance of local-global complementarity, from 1% to 5%.

Our regressions so far are based on a measure of local-global complementarity that takes 5 km by 5

km cells to be a reasonable area for an individual's daily interaction. We next analyze whether our results are sensitive to that particular choice. Column (4) in Table 3 reports our findings when computing local-global complementarity based on 10 km by 10 km cells. Essentially nothing changes: when in the benchmark a one standard deviation increase in local-global complementarity is associated with an increase in child survival by 0.75 per hundred live births, it now raises it to 0.78 per hundred live births.

Table 4. All Outcomes: Comprehensive Specification

	(1) Child Survival	(2) Measles Immunization	(3) Hospital Beds	(4) Literacy Rate	(5) Schooling	(6) Improved Sanitation	(7) Road Density
Global ELF	-3.927*** (1.041)	-26.074*** (4.992)	-1.038 (0.642)	-31.271*** (6.626)	-0.276** (0.127)	-29.751*** (6.726)	2.044 (2.015)
Local-Global Compl.	13.812*** (3.783)	61.983*** (20.909)	8.237** (3.725)	113.923*** (25.651)	1.451*** (0.520)	98.518*** (31.794)	-1.128 (15.601)
Log GDP per Capita	1.410*** (0.225)	0.897 (1.025)	0.245 (0.176)	6.451*** (1.527)	0.153*** (0.022)	9.254*** (1.481)	1.982*** (0.624)
Absolute Latitude	0.003 (0.026)	0.080 (0.115)	0.082*** (0.025)	0.090 (0.166)	0.000 (0.003)	-0.038 (0.170)	0.138 (0.107)
Soil Fertility	1.613** (0.694)	3.587 (2.977)	0.940 (0.595)	9.472** (4.610)	0.315*** (0.078)	2.197 (4.196)	2.118 (2.431)
Roughness	-1.033 (3.032)	-12.824 (12.409)	-0.047 (2.042)	-1.743 (16.160)	-0.588** (0.288)	26.825 (18.806)	-24.410*** (7.714)
Elevation	0.276 (0.934)	1.584 (3.446)	-0.608 (0.609)	4.930 (4.739)	0.203** (0.088)	-6.554 (5.632)	5.481*** (1.799)
Island	-0.286 (0.864)	-7.659* (4.023)	1.038 (0.787)	2.186 (6.011)	0.042 (0.083)	-7.494 (5.559)	3.468 (2.245)
Land Locked	-1.371** (0.622)	-1.311 (2.548)	1.072** (0.474)	1.973 (3.737)	-0.082 (0.069)	2.885 (3.423)	-2.661* (1.410)
Log Population	-0.117 (0.135)	-0.889 (0.640)	-0.064 (0.124)	-0.118 (0.871)	-0.024 (0.016)	-0.410 (1.041)	-1.244*** (0.414)
Constant	87.473*** (3.081)	98.263*** (15.741)	0.941 (3.732)	39.611** (19.651)	1.395*** (0.377)	11.148 (23.928)	21.100* (12.029)
Other Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	147	147	147	121	126	145	147
R^2	0.868	0.632	0.677	0.709	0.739	0.821	0.518

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The column headings give the dependent variables for each of the columns. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation and are based on the authors' calculations. All regressions include the other baseline controls: legal origins dummies and regional dummies. The variable definitions and data sources for each of the variables are provided in Appendix B.

Income inequality and linguistic inequality. The degree of inequality in a country may be related to its linguistic diversity, and may affect public goods outcomes. Table 3, column (5), introduces the income Gini coefficient as an additional control. We do not find a direct effect of income inequality on child survival, and the coefficients on global fractionalization and local-global complementarity hardly change.

Another concern is that income inequality between linguistic groups may be driving the results. Different groups may differ in their willingness to contribute to public goods because their incomes differ. To explore whether this is the case, we use the ethnic inequality variable of Alesina, Michalopoulos and Papaioannou (2016). They construct several measures of ethnic inequality. To be consistent with our measures of diversity, we use the one based on linguistic groups from the Ethnologue, aggregated to level 5. Column (6) in Table 3 reports the results when we control for this measure of linguistic inequality. Once again, the coefficients on global fractionalization and local-global complementarity are very similar to our

baseline regression.

Linguistic segregation. As discussed in the conceptual framework, though our local-global complementarity measure is related to the degree of local mixing, it differs from segregation in an important way. In particular, local-global complementarity tends to be high when there is spatial mixing *and* there are relatively few large-sized groups, whereas the degree of segregation only depends on the degree of local mixing. To further explore the relative importance of both indices, we use our data to calculate the linguistic segregation measure of Alesina and Zhuravskaya (2011),²⁵ and include it jointly with local-global complementarity in our regression. Table 3, column (7), reports the results. Our coefficients of interest do not change, while the segregation variable itself is not statistically significant. From this we conclude that local-global complementarity captures a relevant determinant of public goods provision.

Local fractionalization. Local-global complementarity is also related to local fractionalization. Our main reason for preferring local-global complementarity is conceptual: as we have argued, it captures how local interaction affects the antagonism experienced towards others in a way that local fractionalization does not. However, given the high correlation of 0.92 between the two measures, adding both in the same regression does not provide much insight into which one we should prefer empirically. Tables 11 and 12 in the Online Appendix explore different ways to deal with this multicollinearity issue. Though these suggest that local-global complementarity has more predictive power than local fractionalization, these results should be interpreted with caution.²⁶

Regional variation and geography. We now analyze how robust our findings are to the inclusion of different geographic controls and to regional variation. In Table 4 we present a comprehensive specification that controls for geography, such as roughness of terrain, mean elevation, soil fertility and being landlocked, as well as all baseline regressors. As before, higher global fractionalization is associated with worse outcomes, whereas greater local-global complementarity is associated with improved outcomes. The coefficients on global fractionalization are statistically significant at the 5% level in five out of the seven outcomes, whereas the coefficients on local-global complementarity are always statistically significant at the 5% level, with the exception of road density. This confirms our findings of Table 2. In the Online Appendix we further establish that our results are not driven by particular regions.²⁷

²⁵To be precise, we do not use the data of Alesina and Zhuravskaya (2011), but rather use their segregation index applied to our language data at the 5 km by 5 km grid cell level.

²⁶In addition to segregation and local fractionalization, we also explore the explanatory power of an alternative index that assumes antagonism towards another group no longer depends on the interaction with that specific group but on the interaction with all groups but the own group (see footnote 9). Table 13 in the Online Appendix shows this alternative to have less explanatory power than local-global complementarity.

²⁷To that end, Table 14 in the Online Appendix drops, one at a time, sub-Saharan Africa, East and Southeast Asia and Latin America & the Caribbean. In the same table we also show that replacing legal origins by colonial origins, or our three regional dummies by the full set of six World Bank regional dummies does not affect our main findings.

Complementary interaction. The contact hypothesis argues that interaction is particularly effective in reducing prejudice when different groups need to collaborate to reach a common goal (Allport, 1954). This is consistent with Jha (2013) and Becker and Pascali (2018) who find that local interethnic relations are better if groups are complementary and specialize in different activities.²⁸ To explore whether local interaction has a more benign effect when groups are complementary, we use data on traditional occupations by language group.²⁹ For each country, we compute a Krugman-style index of occupational specialization by linguistic group (Krugman, 1991). In Table 15 of the Online Appendix we interact this specialization index with local-global ethnolinguistic complementarity. As expected, local interaction has a stronger positive association with public goods outcomes in countries where different linguistic groups have traditionally been more complementary to each other.

Decade by decade. Another question is whether our results hold in different time periods. Contingent on data availability, we run regressions for six different decades, starting in the 1960s until the 2010s. Focusing on child survival, for which we have data going back to the 1960s, Panel A of Table 16 in the Online Appendix shows that local-global complementarity is statistically significant at the 10% level in all decades, except for the 1970s. One issue we face is the loss of observations as we go back in time, primarily because of missing data for GDP per capita in some countries. We can increase the number of observations in the earlier decades by relying on GDP per capita data from the Penn World Tables. When doing so, local-global complementarity is statistically significant for all decades at the 5% level (Panel B of Table 16 in Online Appendix).³⁰

Motorized vehicles. In countries with more motorized vehicles, we would expect local interaction to play a lesser role. To see whether this is the case, Table 17 in the Online Appendix interacts the motorization rate (vehicles per 1,000 inhabitants) with local-global ethnolinguistic complementarity. Consistent with our intuition, we get statistically significant signs on the interaction term for 5 out of the 7 outcomes. This implies that a higher motorization rate weakens the effect of localized interaction.

4.3 Causality

In our analysis we have been cautious not to give a causal interpretation to our results. Given the possibility of spatial sorting, a potential endogeneity concern is whether worse provision of public goods may give individuals of the same group an incentive to geographically cluster. For example, evidence by Greif (1993) and others have suggested that in the absence of an effective state, ethnic groups are sometimes able to provide security and contract enforcement to individuals of their own groups. If so, in states with relatively

²⁸See also Lowe (2018) for an application to mixed-caste vs. homogeneous-caste cricket teams.

²⁹The data on traditional occupations by ethnic group come mostly from the Ethnographic Atlas (Murdock, 1967). To get global coverage, Giuliano and Nunn (2018) extended these data from 1,265 to 1,310 ethnic groups. They also matched ethnic groups to languages in the Ethnologue, hence providing historical occupation data at the linguistic group level.

³⁰Panel C of the same table also shows that these results are robust to dropping GDP per capita from the specification.

poor public goods provision, people of the same linguistic or ethnic group may choose to live together. This would imply a positive effect of public goods provision on the degree of spatial mixing.

Does this link between public goods provision and spatial mixing matter for our interpretation of the effect of local-global ethnolinguistic complementarity? We address this question in three ways. First, recall that local-global ethnolinguistic complementarity is not a measure of the degree of geographic mixing of different ethnolinguistic groups. Since there is a fair amount of variation in local-global complementarity for countries with similar levels of spatial mixing, this weakens the potential endogeneity concern. Second, we consider the empirical evidence on whether spatial sorting by ethnolinguistic groups is likely to be affected by public goods provision. Gershman and Rivera (2018) use census data for a number of African countries, and find that diversity at the local level is very stable over time. In particular, the correlation in local diversity over two to three decades exceeds 0.95. More importantly for our purpose, their results indicate that the small changes in subnational diversity are not related to regional economic performance. This also alleviates the endogeneity concern. Third, in Appendix D we design an instrumental variable strategy to further address the issue of causality.

While these three approaches clearly mitigate endogeneity concerns, they do not completely eliminate them. We therefore prefer to err on the side of caution, and conclude that our results are suggestive of local-global ethnolinguistic complementarity improving public goods outcomes.

5 Concluding Remarks

In this paper we proposed a conceptual framework of antagonism which suggests that both a society’s overall ethnolinguistic diversity and its local-global ethnolinguistic complementarity should affect the provision of public goods. Theoretically, we showed that local-global complementarity tends to be high when there are a few large-sized groups that are spatially mixed. This makes local-global complementarity different from both segregation and local diversity. Empirically, we found that overall ethnolinguistic diversity worsens the provision of public goods, whereas the opposite is true for local-global complementarity. This finding is consistent with local interaction mitigating the antagonism experienced toward other groups in the society at large.

Another key contribution of this paper is the construction of a worldwide dataset on local language use. To that end, we combined data on local population, country-level language use, and local language maps. We then applied an iterative proportional fitting algorithm to allocate the speakers of 6,905 different languages to all 5 km by 5 km cells in the world. By using data that span virtually all countries of the world, our results provide important macro support to experimental evidence in favor of local interaction reducing prejudice towards the out-group. This database should also be useful for researchers interested in analyzing the effect of local diversity on a variety of political economy outcomes, such as development and conflict.

References

- [1] Alesina, A., Baqir, R., and Easterly, W., 1999. "Public Goods and Ethnic Divisions," *Quarterly Journal of Economics*, 114, 1243-1284.
- [2] Alesina, A., Baqir, R., and Hoxby, C., 2004. "Political Jurisdictions in Heterogeneous Communities," *Journal of Political Economy*, 112, 348-396.
- [3] Alesina, A., Devleeschauwer, A., Easterly, W., Kurlat, S., and Wacziarg, R., 2003. "Fractionalization," *Journal of Economic Growth*, 8, 155-194.
- [4] Alesina, A. and La Ferrara, E., 2000. "Participation in Heterogeneous Communities," *Quarterly Journal of Economics*, 115, 847-904.
- [5] Alesina, A., Michalopoulos, S., and Papaioannou, E., 2016. "Ethnic Inequality," *Journal of Political Economy*, 124, 428-488.
- [6] Alesina, A. and Zhuravskaya, E., 2011. "Segregation and the Quality of Government in a Cross-Section of Countries," *American Economic Review*, 101, 1872-1911.
- [7] Algan, Y., Hémet, C., and Laitin, D., 2016. "The Social Effects of Ethnic Diversity at the Local Level: A Natural Experiment with Exogenous Residential Allocation," *Journal of Political Economy* 124, 696-733.
- [8] Allport, G., 1954. *The Nature of Prejudice*, Reading, MA: Addison-Wesley.
- [9] Ashraf, Q. and Galor, O., 2013. "The 'Out of Africa' Hypothesis, Human Genetic Diversity, and Comparative Economic Development," *American Economic Review*, 103, 1-46.
- [10] Barro, R. and Lee, J., 2010. "A New Data Set of Educational Attainment in the World, 1950-2010," NBER Working Paper #15902.
- [11] Becker, S. and Pascali, L., 2018. "Religion, Division of Labor and Conflict: Anti-Semitism in German Regions over 600 Years," unpublished manuscript.
- [12] Bertrand, M. and Duflo, E., 2016. "Field Experiments on Discrimination," NBER Working Paper #22014.
- [13] Bishop, Y. M. M., Fienberg, S. E., and Holland, P. W., 1975. *Discrete Multivariate Analysis: Theory and Practice*, MIT Press.
- [14] Boisjoly, J., Duncan, G.J., Kremer, M., Levy, D. M., and Eccles, J., 2006. "Empathy or Antipathy? The Impact of Diversity," *American Economic Review*, 96, 1890-1905.
- [15] Burns, J., Corno, L. and La Ferrara, E., 2015. "Interaction, Prejudice and Performance. Evidence from South Africa," unpublished working paper.

- [16] Carrell, S. E. , Hoekstra, M. and West, J. E., 2015. “The Impact of Intergroup Contact on Racial Attitudes and Revealed Preferences, NBER Working Paper #20940.
- [17] Central Intelligence Agency (2001). *The World Factbook 2001*, Washington, DC: Central Intelligence Agency.
- [18] Chiba, E., 2006. “English Usage in Mauritius,” <http://homes.chass.utoronto.ca/~cpercy/courses/6362-chiba.htm> (accessed September 1, 2014).
- [19] Clarke, D., 2014. PLAUSEXOG: Stata module to implement Conley et al’s plausibly exogenous bounds. Statistical Software Components.
- [20] Condra, L. N. and Linardi, S., 2018. “Casual Contact and Ethnic Bias: Experimental Evidence from Afghanistan.” *Journal of Politics*, forthcoming.
- [21] Conley, T. G., Hansen, C. B., and Rossi, P. E. (2012). Plausibly exogenous. *Review of Economics and Statistics*, 94(1), 260-272.
- [22] Cornes, R. and Hartley, R., 2007. “Aggregative Public Good Games,” *Journal of Public Economic Theory*, 9, 201-219.
- [23] Dahlberg, M., Edmark, K., and Lundqvist, H., 2012. “Ethnic Diversity and Preferences for Redistribution,” *Journal of Political Economy*, 120, 41-76.
- [24] Deming, W. E. and Stephan, F. F., 1940. “On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known,” *Annals of Mathematical Statistics*, 11, 427-444.
- [25] Desmet, K., Ortuño-Ortín, I., and Wacziarg, R., 2012. “The Political Economy of Linguistic Cleavages,” *Journal of Development Economics*, 97, 322-338.
- [26] Easterly, W. and Levine, R., 1997. “Africa’s Growth Tragedy: Policies and Ethnic Divisions,” *Quarterly Journal of Economics*, 112, 1203-1250.
- [27] Enos, R. D., 2014. “Causal Effect of Intergroup Contact on Exclusionary Attitudes,” *PNAS*, 111, 3699-3704.
- [28] Esteban, J. M. and Ray, D., 1994. “On the Measurement of Polarization,” *Econometrica*, 62, 819-851.
- [29] Fienberg, S. E., 1970. “An Iterative Procedure for Estimation in Contingency Tables,” *Annals of Mathematical Statistics*, 41, 907-917.
- [30] Finseraas, H., Hanson, T., Johnsen, A. A., Kotsadam, A., and Torsvik, G., 2016. “Trust, Ethnic Diversity, and Personal Contact: Experimental Field Evidence,” unpublished manuscript.
- [31] Gershman, B. and Rivera, D., 2018. “Subnational diversity in Sub-Saharan Africa: Insights from a new dataset.” *Journal of Development Economics*, 133, 231-263.

- [32] Giuliano, P. and Nunn, N., 2018. "Ancestral Characteristics of Modern Populations," *Economic History of Developing Regions*, 33, 1-17.
- [33] Greif, A., 1993. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition," *American Economic Review*, 83, 525-48.
- [34] Hodler, R., Valsecchi, M. and Vesperoni, A., 2017. "Ethnic Geography: Measurement and Evidence," CEPR Discussion Paper 12378.
- [35] Jha, S., 2013. "Trade, Institutions, and Ethnic Tolerance: Evidence from South Asia," *American Political Science Review*, 107, 806-832.
- [36] Karaivanov, A., 2009. "Heterogeneity, Returns to Scale, and Collective Action," *Canadian Journal of Economics*, 42, 771-807.
- [37] Krugman, P., 1991. *Geography and Trade*, MIT Press.
- [38] La Porta, R., Lopez-de-Silanes, F., Shleifer, A., and Vishny, R., 1999. "The Quality of Government," *Journal of Law, Economics, and Organization*, 15, 222-79.
- [39] La Porta, R., Lopez-de-Silanes, F., and Shleifer, A. 2008. "The Economic Consequences of Legal Origins," *Journal of Economic Literature*, 46, 285-332.
- [40] Lawton, C. and Ackrill, R. 2016. "Hard Evidence: How Areas with Low Immigration Voted Mainly for Brexit," *The Conversation*, <http://theconversation.com/hard-evidence-how-areas-with-low-immigration-voted-mainly-for-brexit-62138> (accessed 15 January 2017).
- [41] Lewis, P. M., Simons, G. F., and Fennig, C. D., 2014. *Ethnologue: Languages of the World*, 17th edition, Dallas: SIL International.
- [42] Lizzeri, A. and Persico, N., 2004. "Why did the Elites Extend the Suffrage? Democracy and the Scope of Government, with an Application to Britain's Age of Reform," *Quarterly Journal of Economics*, 119, 707-765.
- [43] Lowe, M., 2018. "Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration," unpublished manuscript.
- [44] Massey, D. S. and Denton, N. A., 1988. "The Dimensions of Residential Segregation," *Social Forces*, 67, 281-315.
- [45] Matuszeski, J. and Schneider, F., 2006. "Patterns of Ethnic Group Segregation and Civil Conflict," Harvard University, unpublished working paper.

- [46] Miguel, E. and Gugerty, M. K., 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya," *Journal of Public Economics*, 89, 2325-2368.
- [47] Montalvo, J. G. and Reynal-Querol, M., 2005. "Ethnic Polarization, Potential Conflict, and Civil Wars," *American Economic Review*, 95, 796-816.
- [48] Montalvo, J. G. and Reynal-Querol, M., 2016. "Ethnic Diversity and Growth: Revisiting the Evidence," UPF, mimeo.
- [49] Munshi, K. and Rosenzweig, M., 2015. "Insiders and Outsiders: Local Ethnic Politics and Public Goods Provision," NBER Working Paper # 21720.
- [50] Murdock, G. P., 1967. *Ethnographic Atlas: A Summary*, University of Pittsburgh Press, Pittsburgh, PA.
- [51] Pettigrew, T.F. and Tropp, L.R., 2006. "A Meta-Analytic Test of Intergroup Contact Theory," *Journal of Personality and Social Psychology*, 90, 751-783.
- [52] Putnam, R.D., 2007. "*E Pluribus Unum*: Diversity and Community in the Twenty-First Century," *Scandinavian Political Studies*, 30, 137-174.
- [53] Sambanis, N. and Shayo, M., 2013. "Social Identification and Ethnic Conflict," *American Political Science Review*, 107, 294-325.
- [54] Reynal-Querol, M., 2002. "Ethnicity, Political Systems, and Civil Wars," *Journal of Conflict Resolution*, 46, 29-54.
- [55] Robinson, A.L., 2013. "Ethnic Diversity, Segregation, and Ethnocentric Trust in Africa," Stanford University, unpublished working paper.
- [56] Tajima, Y., Samphantharak, K., and Ostwald, K., 2018. "Ethnic Segregation and Public Goods: Evidence from Indonesia," *American Political Science Review*, 112, 637-653.
- [57] Treisman, D., 2008. Decentralization Dataset. Downloaded on August 10, 2016 at <http://www.sscnet.ucla.edu/polisci/faculty/treisman/>.
- [58] Weidmann, N. B., Rod, J. K., and Cederman, L., 2010, "Representing Ethnic Groups in Space: A New Dataset," *Journal of Peace Research*, 47(4), 491-499.

A. Proof of Proposition 1

Proof. We use the “replacement function” approach of Cornes and Hartley (2007) to analyze the Nash equilibrium of a private contribution game. Given that agents have Cobb-Douglas utility functions and all of them have the same income, the proof is simple. In our case the replacement function for an agent i from cell ℓ with valuation of the public good $v_{\ell i}$ is

$$\max\{y - \frac{G}{v_{\ell i}}, 0\}$$

where $y = 1$ denotes the individual’s income and G denotes the total provision of public good (Karaivanov, 2009, example 3.1.1, page 780). The replacement function $\max\{y - \frac{G}{v_{\ell i}}, 0\}$ gives “the quantity that if subtracted from the total provision G , the player’s best reply response to the remaining quantity would exactly replace the quantity removed” (Cornes and Hartley, 2007, page 205). Assuming all individuals contribute a strictly positive amount, Karaivanov (2009) shows that the equilibrium total contribution G^* solves

$$\sum_{\ell} \sum_i N s_{\ell} s_{\ell i} (1 - \frac{G}{v_{\ell i}}) = G \quad (12)$$

which can be written as

$$\sum_{\ell} \sum_i N s_{\ell} s_{\ell i} - \sum_{\ell} \sum_i N s_{\ell} s_{\ell i} \left(\frac{\kappa_2 + a_{\ell i}}{\kappa_1} \right) G = G \quad (13)$$

or

$$N - NG \frac{\kappa_2}{\kappa_1} - NG \sum_{\ell} \sum_i s_{\ell} s_{\ell i} \frac{a_{\ell i}}{\kappa_1} = G,$$

so

$$N - \frac{NG}{\kappa_1} (\kappa_2 - \sum_{\ell} s_{\ell} a_{\ell}) = G.$$

Recall that average antagonism is

$$A = \sum_{\ell} s_{\ell} a_{\ell}$$

so the solution to (13) can be written as

$$G = \frac{\kappa_1 N}{\kappa_1 + N(\kappa_2 + A)} \quad (14)$$

Expression (14) shows that the equilibrium total contribution G is a decreasing function of average antagonism A . ■

B. Data Appendix

Child survival. Child survival rate per 100 live births, 1990-2010 average. *Source: World Development Indicators, World Bank.*

Hospital beds. Hospital beds per 1,000 people, 1990-2010 average. *Source: World Development Indicators, World Bank.*

Measles immunization. Percentage of children between the age of 12 and 23 months that have been immunized against measles, 1990-2010 average. *Source: World Development Indicators, World Bank.*

Improved sanitation. Percentage of population with access to improved sanitation facilities, 1990-2010 average. *Source: World Development Indicators, World Bank.*

Roads. Road network density, km per 1,000 people, 1990-2010 average. *Source: World Development Indicators, World Bank.*

Literacy. Percentage of people aged 15 and above who are literate, 1990-2010 average. *Source: World Development Indicators, World Bank.*

School attainment. Log of $1 +$ average years of schooling for people 25 years of age and above, 1990-2010 average. *Source: Barro R. and J.W. Lee v. 1.3, 04/13.*

Decentralization. A dummy variable indicating whether the country is a federal state or not. *Source: Treisman (2008).*

Log GDP per capita and log population. Both variables are the average for the period 1990-2010. *Source: World Development Indicators, World Bank.*

Legal origin. French, German or UK legal origin *Source: La Porta, Lopez-de-Silanes and Shleifer (2008).*

Geographic controls. Absolute latitude, log of soil fertility, roughness of terrain and mean elevation. *Source: Ashraf and Galor (2013).*

Gini coefficient. Income Gini coefficient, 1990-2010 average. *Source: World Development Indicators, World Bank.*

Linguistic inequality. Income inequality between ethnic groups, based on linguistic groups of Ethnologue aggregated to level 5. *Source: Alesina, Michalopoulos and Papaioannou (2016).*

Country boundary shapefile. ArcGIS shapefile with the political boundaries of all countries in the world. *Source: Seamless Digital Chart of the World Base Map Version 10.0, World GeoDatasets.*

Ethnolinguistic maps. For information on linguistic groups we use the digitized version of the 17th edition of Ethnologue which maps over 6,905 linguistic groups for the whole world. The data on different languages come from a variety of censuses and years and approximately correspond to the 1990s. *Source: World Language Mapping System Version 17, World GeoDatasets.*

Cell-level population data. The cell level population data comes from LandScan who provide global population distribution data at the resolution of approximately $1\text{km} \times 1\text{km}$ ($30'' \times 30''$) which represents an ambient population (average over 24 hours). *Source: <http://web.ornl.gov/sci/landscan/>*

Global fractionalization, local fractionalization, local-global complementarity, and linguistic segregation. *Source: Own calculations.*

Table B.1. Summary Statistics

	Obs	Mean	SD	Min	Max
1. Indices					
Local-Global Complementarity (level 2)	223	0.05	0.07	0.00	0.25
Local-Global Complementarity (level 5)	223	0.05	0.06	0.00	0.25
Local-Global Complementarity (level 15)	223	0.06	0.06	0.00	0.25
Local-Global Complementarity (level 5, 10 km)	223	0.05	0.07	0.00	0.25
Global ELF (level 2)	223	0.19	0.21	0.00	0.90
Global ELF (level 5)	223	0.25	0.26	0.00	0.97
Global ELF (level 15)	223	0.38	0.31	0.00	0.99
Local ELF (level 2)	223	0.12	0.15	0.00	0.71
Local ELF (level 5)	223	0.14	0.16	0.00	0.73
Local ELF (level 15)	223	0.20	0.17	0.00	0.73
Local ELF (level 5, 10 km)	223	0.14	0.16	0.00	0.73
Local-Global Complementarity (level 5, instrument)	223	0.06	0.06	0.00	0.25
2. Dependent variables					
Child survival	184	94.35	5.56	76.92	99.58
Measles immunization	183	81.72	15.48	27.00	99.00
Hospital beds	188	3.64	3.25	0.17	20.62
Literacy	147	79.922	20.39	19.032	99.998
School attainment	141	1.98	0.44	0.67	2.63
Improved sanitation	190	69.35	30.70	7.03	100.00
Roads	181	7.43	8.52	0.32	54.12

This table reports the number of observations, the mean, the standard deviation, the minimum and the maximum of the different indices. Local-Global Complementarity (level 5, instrument) refers to the instrument for Local-Global Complementarity discussed in Appendix D.

Table B.2. Correlations between Indices

	LGC (2)	LGC (5)	LGC (15)	LGC (5,10k)	G-ELF (2)	G-ELF (5)	G-ELF (15)	L-ELF (2)	L-ELF (5)	L-ELF (15)	L-ELF (5,10k)	LGC Inst (5)
Local-Global Complementarity (level 5)	0.95	1.00										
Local-Global Complementarity (level 15)	0.69	0.73	1.00									
Local-Global Complementarity (level 5, 10 km)	0.94	0.98	0.70	1.00								
Global ELF (level 2)	0.74	0.66	0.39	0.66	1.00							
Global ELF (level 5)	0.59	0.58	0.28	0.58	0.84	1.00						
Global ELF (level 15)	0.38	0.39	0.26	0.40	0.59	0.76	1.00					
Local ELF (level 2)	0.97	0.89	0.62	0.88	0.83	0.65	0.43	1.00				
Local ELF (level 5)	0.91	0.92	0.62	0.91	0.78	0.78	0.55	0.92	1.00			
Local ELF (level 15)	0.67	0.69	0.75	0.68	0.60	0.63	0.76	0.68	0.76	1.00		
Local ELF (level 5, 10 km)	0.89	0.89	0.58	0.92	0.79	0.78	0.56	0.91	0.98	0.73	1.00	
Local-Global Complementarity (level 5, instr)	0.79	0.83	0.59	0.82	0.67	0.66	0.44	0.75	0.79	0.60	0.77	1.00
Local ELF (level 5, instr)	0.68	0.68	0.41	0.67	0.80	0.90	0.64	0.74	0.83	0.65	0.81	0.85

LGC refers to local-global complementarity, G-ELF refers to global ELF, L-ELF refers to local ELF, and LGC Inst refers to the instrument for LGC discussed in Appendix D. First number in brackets corresponds to the linguistic level of aggregation, whereas the second number corresponds to the geographic size of cells, with a default cell size of 5 km by 5 km. For example, L-ELF (5,10k) refers to the local fractionalization index at linguistic aggregation level 5, using a cell size of 10 km by 10 km.

Table B.3. Global Fractionalization and Local-Global Complementarity by Country

Country	Global ELF	LGC	Country	Global ELF	LGC
Afghanistan	0.681	0.060	Denmark	0.022	0.006
Albania	0.562	0.079	Djibouti	0.492	0.077
Algeria	0.342	0.024	Dominica	0.313	0.155
American Samoa	0.066	0.033	Dominican Republic	0.035	0.017
Andorra	0.000	0.000	East Timor	0.635	0.152
Angola	0.011	0.005	Ecuador	0.182	0.043
Anguilla	0.141	0.070	Egypt	0.080	0.034
Antigua and Barbuda	0.500	0.250	El Salvador	0.004	0.002
Argentina	0.011	0.005	Equatorial Guinea	0.031	0.012
Armenia	0.119	0.029	Eritrea	0.334	0.041
Aruba	0.328	0.143	Estonia	0.383	0.192
Australia	0.033	0.000	Ethiopia	0.782	0.042
Austria	0.017	0.008	Fiji	0.506	0.248
Azerbaijan	0.410	0.065	Finland	0.097	0.042
Bahamas	0.295	0.147	France	0.090	0.030
Bahrain	0.000	0.000	French Polynesia	0.430	0.212
Bangladesh	0.018	0.008	Gabon	0.060	0.029
Barbados	0.092	0.046	Gambia	0.658	0.138
Belarus	0.000	0.000	Georgia	0.547	0.115
Belgium	0.537	0.041	Germany	0.011	0.006
Belize	0.755	0.046	Ghana	0.649	0.062
Benin	0.552	0.057	Greece	0.061	0.021
Bermuda	0.000	0.000	Grenada	0.064	0.030
Bhutan	0.746	0.062	Guadeloupe	0.033	0.016
Bolivia	0.538	0.115	Guam	0.499	0.250
Bosnia and Herzegovina	0.002	0.001	Guatemala	0.517	0.058
Botswana	0.084	0.028	Guinea	0.603	0.017
Brazil	0.034	0.007	Guinea-Bissau	0.773	0.092
British Virgin Islands	0.167	0.084	Guyana	0.055	0.011
Brunei	0.578	0.135	Haiti	0.000	0.000
Bulgaria	0.150	0.007	Honduras	0.044	0.011
Burkina Faso	0.439	0.032	Hungary	0.082	0.036
Burundi	0.000	0.000	Iceland	0.000	0.000
Cambodia	0.067	0.015	India	0.863	0.015
Cameroon	0.474	0.024	Indonesia	0.760	0.036
Canada	0.429	0.011	Iran	0.513	0.126
Cape Verde Islands	0.000	0.000	Iraq	0.436	0.102
Cayman Islands	0.000	0.000	Ireland	0.063	0.006
Central African Republic	0.433	0.113	Israel	0.521	0.154
Chad	0.848	0.044	Italy	0.255	0.083
Chile	0.042	0.017	Jamaica	0.000	0.000
China	0.084	0.014	Japan	0.036	0.008
Colombia	0.021	0.008	Jordan	0.031	0.016
Comoros	0.008	0.004	Kazakhstan	0.600	0.201
Congo	0.521	0.117	Kenya	0.595	0.020
Cook Islands	0.000	0.000	Kiribati	0.017	0.008
Costa Rica	0.040	0.009	Kuwait	0.040	0.020
Cote d'Ivoire	0.715	0.059	Kyrgyzstan	0.600	0.205
Croatia	0.090	0.037	Laos	0.489	0.018
Cuba	0.000	0.000	Latvia	0.453	0.224
Cyprus	0.244	0.040	Lebanon	0.146	0.069
Czech Republic	0.024	0.009	Lesotho	0.000	0.000
DRC	0.370	0.072	Liberia	0.715	0.082

LGC refers to local-global complementarity.

Country	Global ELF	LGC	Country	Global ELF	LGC
Libya	0.111	0.017	Russian Federation	0.184	0.018
Liechtenstein	0.000	0.000	Rwanda	0.001	0.000
Lithuania	0.342	0.148	St. Kitts and Nevis	0.010	0.005
Luxembourg	0.353	0.177	St. Lucia	0.020	0.010
Macedonia	0.472	0.042	St. Pierre and Miquelon	0.069	0.035
Madagascar	0.025	0.013	St. Vincent & Grenadines	0.007	0.004
Malawi	0.002	0.001	Samoa	0.002	0.001
Malaysia	0.626	0.129	San Marino	0.000	0.000
Maldives	0.000	0.000	Saudi Arabia	0.000	0.000
Mali	0.821	0.073	Senegal	0.282	0.044
Malta	0.076	0.038	Serbia	0.413	0.046
Marshall Islands	0.000	0.000	Seychelles	0.099	0.048
Martinique	0.041	0.021	Sierra Leone	0.746	0.038
Mauritania	0.221	0.026	Slovakia	0.192	0.025
Mauritius	0.529	0.202	Slovenia	0.012	0.006
Mayotte	0.437	0.208	Solomon Islands	0.238	0.066
Mexico	0.107	0.023	Somalia	0.057	0.007
Micronesia	0.075	0.037	South Africa	0.393	0.080
Moldova	0.549	0.183	South Korea	0.000	0.000
Monaco	0.000	0.000	South Sudan	0.717	0.039
Mongolia	0.156	0.026	Spain	0.027	0.011
Montenegro	0.365	0.150	Sri Lanka	0.319	0.062
Montserrat	0.026	0.013	Sudan	0.331	0.043
Morocco	0.435	0.013	Suriname	0.788	0.152
Mozambique	0.122	0.061	Swaziland	0.037	0.018
Myanmar	0.473	0.028	Sweden	0.048	0.020
Namibia	0.388	0.070	Switzerland	0.483	0.034
Nauru	0.189	0.095	Syria	0.197	0.021
Nepal	0.697	0.060	Taiwan	0.035	0.006
Netherlands	0.185	0.043	Tajikistan	0.430	0.146
Netherlands Antilles	0.455	0.062	Tanzania	0.070	0.010
New Caledonia	0.507	0.194	Thailand	0.180	0.022
New Zealand	0.081	0.035	Togo	0.571	0.080
Nicaragua	0.068	0.005	Tonga	0.000	0.000
Niger	0.633	0.076	Trinidad and Tobago	0.509	0.248
Nigeria	0.872	0.035	Tunisia	0.008	0.003
Niue	0.071	0.036	Turkey	0.375	0.060
Norfolk Island	0.311	0.155	Turkmenistan	0.176	0.025
North Korea	0.000	0.000	Turks and Caicos Islands	0.146	0.073
N. Mariana Islands	0.538	0.214	Uganda	0.673	0.136
Norway	0.011	0.004	Ukraine	0.090	0.039
Oman	0.298	0.075	United Arab Emirates	0.445	0.167
Pakistan	0.666	0.042	United Kingdom	0.025	0.006
Palau	0.063	0.032	United States	0.246	0.088
Palestinian WB & Gaza	0.000	0.000	U.S. Virgin Islands	0.454	0.227
Panama	0.265	0.032	Uruguay	0.004	0.002
Papua New Guinea	0.967	0.011	Uzbekistan	0.338	0.076
Paraguay	0.517	0.143	Vanuatu	0.179	0.082
Peru	0.315	0.048	Venezuela	0.024	0.010
Philippines	0.421	0.023	Viet Nam	0.181	0.039
Poland	0.042	0.018	Wallis and Futuna	0.018	0.009
Portugal	0.050	0.025	Yemen	0.014	0.000
Puerto Rico	0.049	0.024	Zambia	0.028	0.014
Qatar	0.485	0.242	Zimbabwe	0.032	0.016
Romania	0.165	0.012			

LGC refers to local-global complementarity.

C. Appendix Figures

Figure C.1. Global ELF by Country – Level 15




Figure C.2. Local-Global Complementarity by Country – Level 15




Figure C.3. Local-Global Complementarity at 5 km by 5 km Resolution – Level 15




Figure C.4. Global ELF by Country – Level 2




Figure C.5. Local-Global Complementarity by Country – Level 2





Figure C.6. Local-Global Complementarity at 5 km by 5 km Resolution – Level 2



D. Instrumental Variable Approach


Instrument for local-global complementarity. As explained in Section 4.3 of the main text, our endogeneity concern relates to the spatial distribution of a country’s language groups across its territory. That is, it does not pertain to the number and the population shares of the groups, but to their geographic distribution. We therefore want to predict $\mathcal{B}$, the $K \times M$ binary matrix of language use by cell, in a way that is independent of any spatial sorting. Once we have a predicted matrix of $\mathcal{B}$, we can apply our iterative proportional fitting algorithm to get an instrument for the local-global complementarity variable.

To create a predicted measure of local language use, we follow an approach similar to the one in Alesina and Zhuravskaya (2011) by relying on language use in neighboring countries. In particular, for each cell ℓ in country c , we determine the closest cell k in any of the neighboring countries of c . Any language that is spoken in k and that is also spoken in c is then assigned to ℓ . For languages that are spoken in c and that are not spoken in any of the closest cells in the neighboring countries, we assume that they are spoken in all cells of c . This methodology yields a $K \times M$ binary matrix $\hat{\mathcal{B}}$ with predicted values of language use. We then apply the same algorithm as the one described in Section 3.2, but using $\hat{\mathcal{B}}$ instead of $\mathcal{B}$. This yields a predicted measure of local-global complementarity. For this to be a valid instrument, the exclusion restriction requires that the spatial distribution of languages in the closest cells of the neighboring countries does not directly affect the public goods outcomes in the own country. We will discuss the plausibility of the exclusion restriction at the end of this appendix.

We now give an example of how neighboring countries are used to predict a country’s geographic distribution of languages. We focus on the case of Belgium, and go through the following steps.

1. For each cell ℓ in Belgium, we determine the closest cell k in any of Belgium’s neighboring countries of c . Figure D.1 illustrates this.

Figure D.1. Closest Neighbors of Each Cell in Belgium



5. To allocate the number of language speakers to each cell in Belgium, we use the same algorithm as the one described in Section 3.2, but using $\hat{\mathcal{B}}$ instead of $\mathcal{B}$.
6. We then use this predicted allocation to construct our instrument for local-global complementarity in Belgium.

Table D.1. Global ELF and Local-Global Ethnolinguistic Complementarity (IV)

	(1) Child Survival	(2) Measles Immunization	(3) Hospital Beds	(4) Literacy Rate	(5) Schooling	(6) Improved Sanitation	(7) Road Density
Global ELF	-6.445*** (1.343)	-29.824*** (5.317)	-0.400 (0.804)	-40.948*** (7.208)	-0.360** (0.167)	-30.904*** (7.433)	4.326** (2.003)
Local-Global Compl.	27.188*** (7.398)	98.887*** (27.932)	1.368 (4.253)	164.212*** (42.787)	2.107** (0.964)	120.844*** (41.819)	-15.960 (12.495)
Absolute Latitude	0.009 (0.020)	0.192** (0.086)	0.102*** (0.019)	0.249** (0.121)	0.008*** (0.003)	0.130 (0.159)	0.177*** (0.062)
Latin America & Carib.	0.291 (0.615)	4.209 (2.734)	0.345 (0.430)	5.263 (3.569)	0.212** (0.087)	-5.121 (4.830)	2.511 (1.687)
Sub-Saharan Africa	-5.887*** (0.745)	-9.267*** (2.876)	0.096 (0.458)	-8.990** (4.068)	-0.128 (0.103)	-25.216*** (5.038)	2.908* (1.691)
East and S.E. Asia	1.143 (0.789)	3.374 (3.634)	0.809 (0.860)	15.138*** (4.893)	0.107 (0.101)	-5.347 (6.378)	-1.924 (2.093)
Log GDP per Capita	1.421*** (0.190)	1.083 (0.835)	0.307** (0.154)	3.548*** (1.135)	0.111*** (0.022)	9.810*** (1.253)	1.998*** (0.447)
French Legal Origin	0.767 (0.599)	0.972 (2.962)	0.693 (1.415)	-3.942 (2.681)	0.050 (0.087)	11.930*** (4.184)	-13.203* (7.599)
German Legal Origin	1.445*** (0.504)	2.133 (2.798)	2.513* (1.438)	0.076 (4.260)	0.211*** (0.073)	12.200*** (3.764)	-10.865 (7.737)
UK legal origin	1.107* (0.591)	6.329** (3.204)	0.712 (1.426)	0.000 (.)	0.232*** (0.084)	9.222** (4.116)	-10.185 (7.787)
Constant	83.652*** (1.874)	69.868*** (8.960)	-2.579 (1.960)	53.067*** (11.373)	0.712*** (0.258)	-13.221 (14.085)	-2.798 (9.148)
First-Stage F-Statistic	60.885	60.885	61.604	41.292	36.568	63.386	60.943
Observations	171	171	173	138	136	171	172
R^2	0.810	0.562	0.589	0.646	0.670	0.780	0.445

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table gives the second stage of 2SLS IV regressions. The column headings give the dependent variables for each of the columns. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation and are based on the authors' calculations. The Local-Global Ethnolinguistic Complementarity variable has been instrumented using a predicted Local-Global Ethnolinguistic Complementarity variable based on languages spoken in neighboring countries. The variable definitions and data sources for each of the variables are provided in Appendix B.

IV results. Table D.1 reports the IV results for our baseline regression. For five of the seven public goods, local-global complementarity has the expected sign and is statistically significant at the 5% level. (For the other two outcomes, the effect is not statistically significant.) The F-statistics of the first stage are all larger than the Stock-Yogo critical values, so we can reject the hypothesis that the instrument is weak. In general we find the effect of local-global complementarity to be between 50% and 150% larger than in the OLS regressions. For example, while a one standard deviation increase in local-global complementarity increased child survival by 0.75 per hundred in the OLS regression, it now increases child survival by 1.64 per hundred. The corresponding standardized β is 29%.

Table D.2 reports the IV regressions for all outcomes, using the most comprehensive specification of Table 4, which includes a host of geographic controls, as well as legal origins, regional dummies, population

Table D.2. Global ELF and Local-Global Complementarity, Comprehensive Specification (IV)

	(1) Child Survival	(2) Measles Immunization	(3) Hospital Beds	(4) Literacy Rate	(5) Schooling	(6) Improved Sanitation	(7) Road Density
Global ELF	-6.643*** (1.373)	-32.883*** (6.292)	0.222 (1.037)	-47.484*** (9.072)	-0.344* (0.180)	-39.049*** (9.515)	8.439*** (3.016)
Local-Global Compl.	35.117*** (8.901)	115.402*** (43.620)	-1.651 (7.496)	235.940*** (67.851)	1.932 (1.316)	171.615** (68.807)	-51.297** (22.084)
Log GDP per Capita	1.300*** (0.232)	0.620 (1.017)	0.297* (0.166)	5.451*** (1.567)	0.151*** (0.022)	8.911*** (1.449)	2.243*** (0.627)
Absolute Latitude	-0.006 (0.026)	0.058 (0.111)	0.086*** (0.024)	-0.007 (0.183)	0.000 (0.003)	-0.069 (0.167)	0.159 (0.099)
Soil Fertility	1.596** (0.680)	3.543 (2.837)	0.948* (0.565)	9.151** (4.513)	0.316*** (0.072)	2.197 (4.108)	2.159 (2.253)
Roughness	1.252 (3.162)	-7.095 (13.156)	-1.107 (2.130)	13.487 (17.791)	-0.535* (0.308)	34.978* (19.036)	-29.791*** (8.222)
Elevation	-0.342 (0.931)	0.035 (3.494)	-0.321 (0.655)	0.934 (4.933)	0.190** (0.088)	-8.727 (5.333)	6.936*** (1.936)
Island	-0.288 (0.792)	-7.664** (3.736)	1.039 (0.745)	0.401 (6.184)	0.042 (0.077)	-7.517 (5.149)	3.472* (2.037)
Land Locked	-1.174* (0.637)	-0.819 (2.509)	0.981** (0.469)	2.704 (3.587)	-0.082 (0.064)	3.568 (3.464)	-3.123** (1.546)
Log Population	0.064 (0.159)	-0.435 (0.707)	-0.148 (0.136)	0.820 (0.916)	-0.020 (0.016)	0.231 (1.110)	-1.671*** (0.518)
Constant	86.129*** (3.150)	94.892*** (15.570)	1.565 (3.618)	33.377 (22.155)	1.375*** (0.345)	5.912 (23.192)	24.265** (11.738)
Other Baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F-Statistic	29.078	29.078	29.078	20.204	21.688	28.456	29.078
Observations	147	147	147	121	126	145	147
R ²	0.847	0.614	0.661	0.655	0.737	0.813	0.470

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table gives the second stage of 2SLS IV regressions. The column headings give the dependent variables for each of the columns. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation and are based on the authors' calculations. The Local-Global Ethnolinguistic Complementarity variable has been instrumented using a predicted Local-Global Ethnolinguistic Complementarity variable based on languages spoken in neighboring countries. All regressions include the other baseline controls: legal origins dummies and regional dummies. The variable definitions and data sources for each of the variables are provided in Appendix B.

and GDP per capita. With the exception of road density, for which we get the wrong sign, the results for local-global complementarity become stronger. A one-standard deviation increase in local-global complementarity increases child survival by an estimated 1.92 per hundred, with a corresponding standardized β of 33%. Our IV results are suggestive of a causal effect between higher local-global complementarity and improved public goods outcomes.

Plausibility of the exclusion restriction. We now discuss the plausibility of our identification strategy which relies on using predicted local-global complementarity as an instrument for observed local-global complementarity. For the exclusion restriction to hold, the predicted local-global complementarity must not have a direct effect on public goods over and above its indirect effect through the observed local-global complementarity. This exclusion restriction might be violated if public goods in the neighboring countries have a direct effect on public goods in the home country. To see why such a situation might be problematic for our identification strategy, recall that the spatial distribution of languages in the neighboring countries affects both the predicted local-global complementarity in the home country and public goods in the neighboring countries. The presence of cross-border spillovers in public goods might then lead to a direct relation between the predicted local-global complementarity and public goods in the home country. For example, if higher

immunization rates in the neighboring countries lowers the spread of infectious diseases in the broader region, then any increase in child survival in the neighboring countries would extend to the home country. If there is a strong correlation between the predicted local-global complementarity in the home country and the observed local-global complementarity in the neighboring countries, such spillovers could imply a violation of the exclusion restriction. To address this potential concern, we adopt a three-pronged strategy.

First, we clarify that we do not directly use the local-global complementarity of the neighboring countries to predict the local-global complementarity of the home country. Instead, for each cell in the home country we only use information on which languages are spoken in the closest cell of its neighboring countries to predict which languages are spoken in the home country cell. This procedure yields a predicted binary matrix of language use for all cells in the home country, which is then used to construct a predicted local-global complementarity. Hence, the predicted local-global complementarity of the home country may be very different from the local-global complementarity levels of its neighbors. In fact, the correlation between both measures is essentially zero (to be precise, between 0.03 and 0.06, depending on how we average the neighbors' local-global complementarity levels). With such a low correlation, it would be highly improbable that cross-border spillovers in public goods would invalidate the exclusion restriction.

Second, to be exhaustive, we can of course still investigate whether the neighbors' local-global complementarity, global fractionalization and public goods affect outcomes. When controlling for these additional variables in our otherwise identical IV specification of Table D.2, the results are unchanged.³¹ In particular, in Table D.3 local-global complementarity continues to have a significant and robust effect on the different public goods. The magnitudes of the effects are very similar to those found in the specification without the additional controls. This reinforces our prior conclusion that the possible presence of cross-border spillovers in public goods does not undermine our identification strategy. As an alternative specification, we could separately control for, on the one hand, public goods, and on the other hand, local-global complementarity and global fractionalization. Doing so does not substantially alter the magnitudes of the original effects of local-global complementarity.³²

Third, while in the previous paragraphs we have argued that our identification strategy is likely to hold, next we check for the sensitivity of our results to small violations of the exclusion restriction. To that end, we follow the method outlined in Conley et al. (2012). Consider the equation we are interested in estimating, but now allow the instrument to have a direct impact on the outcome. That is, $G_c = \gamma_0 X_c + \gamma_1 ELF_{glob,c} + \gamma_2 LGC_c + \gamma_3 LGC_{nc} + \varepsilon_c$, where LGC_{nc} is the instrument. The exclusion restriction assumes that $\gamma_3 = 0$. Conley et al. (2012) provide a procedure that allows for inference even if γ_3 is not exactly equal to zero. The idea is intuitive. If we were to know the true value of γ_3 , it would be straightforward

³¹Neighboring country variables are computed as population-weighted averages of the neighboring countries' variables. Results do not change when either taking simple averages or using the populations in the home country cells that are exposed to each neighboring country as weights.

³²Results are available upon request.

to estimate the above equation by 2SLS using LGC_{nc} as an instrument for LGC_c .³³ Of course, we do not know the true value of γ_3 . Conley et al. (2012) tackle this problem by making assumptions on the support of γ_3 , and then run 2SLS on different possible values of γ_3 . Following their *Local to Zero* approach, we assume that $\gamma_3 \sim N(0, \delta^2)$. Using their methodology, we can obtain a 95% confidence interval for our coefficient of interest, γ_2 .

Table D.3. Neighbors' Global ELF and Local-Global Complementarity, Comprehensive Specification (IV)

	(1) Child Survival	(2) Measles Immunization	(3) Hospital Beds	(4) Literacy Rate	(5) Schooling	(6) Improved Sanitation	(7) Road Density
Local-Global Compl.	31.431*** (8.195)	105.024** (46.363)	-0.886 (7.414)	211.004*** (60.452)	1.798 (1.280)	168.179** (68.696)	-41.570** (20.454)
Global ELF	-6.137*** (1.363)	-30.656*** (6.888)	-0.084 (1.105)	-40.548*** (9.041)	-0.346* (0.183)	-38.103*** (10.341)	6.089* (3.317)
Neighbors' Local-Global Compl.	5.308 (9.980)	54.277* (32.644)	-1.789 (6.297)	71.297 (57.609)	0.034 (1.085)	33.850 (63.486)	-24.927 (26.124)
Neighbors' Global ELF	2.612 (1.717)	5.872 (6.482)	0.320 (1.025)	-0.018 (10.154)	0.141 (0.168)	2.503 (9.931)	3.592 (4.998)
Neighbors' Public Goods	0.325*** (0.125)	0.080 (0.120)	0.352*** (0.124)	-0.273** (0.121)	0.090 (0.097)	-0.005 (0.101)	0.136 (0.152)
Controls Table D.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F-Statistic	31.922	27.839	27.586	20.045	22.119	34.871	27.949
Observations	146	146	146	121	125	144	146
R^2	0.864	0.634	0.688	0.700	0.745	0.813	0.444
Observations	146	146	146	121	125	144	146
fs	31.922	27.839	27.586	20.045	22.119	34.871	27.949
r2	0.864	0.634	0.688	0.700	0.745	0.813	0.444

Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table gives the second stage of 2SLS IV regressions. The column headings give the dependent variables for each of the columns. The Global ELF and the Local-Global Ethnolinguistic Complementarity variables are measured at level 5 of aggregation and are based on the authors' calculations. The Local-Global Ethnolinguistic Complementarity variable has been instrumented using a predicted Local-Global Ethnolinguistic Complementarity variable based on languages spoken in neighboring countries. The variable definitions and data sources for each of the variables are provided in Appendix B.

Table D.4. Plausibly Exogenous IV Estimation, Comprehensive Specification

	(1) Child Survival	(2) Measles Immunization	(3) Hospital Beds	(4) Literacy Rate	(5) Schooling	(6) Improved Sanitation	(7) Road Density
Panel A: IV							
Local-Global Compl.	35.117*** (8.901)	115.402*** (43.620)	-1.651 (7.496)	235.940*** (67.851)	1.932 (1.316)	171.615** (68.807)	-51.297** (22.084)
Panel B: 95% confidence intervals for Local-Global Complementarity under $\gamma_3 \sim N(0, \delta^2)$							
CI ($2\delta = 10\%$)	(16.3 54.0)	(26.8 204.0)	(-16.3 13.0)	(94.4 377.5)	(-0.7 4.5)	(32.4 310.9)	(-95.8 -6.8)
CI ($2\delta = 15\%$)	(14.7 55.6)	(23.0 207.8)	(-16.4 13.0)	(84.4 387.5)	(-0.7 4.6)	(27.1 316.2)	(-97.3 -5.3)
CI ($2\delta = 25\%$)	(10.3 60.0)	(12.0 218.8)	(-16.4 13.1)	(56.0 415.9)	(-0.8 4.7)	(11.3 331.9)	(-101.7 -0.9)
CI ($2\delta = 50\%$)	(-43.8 74.6)	(-29.1 259.9)	(-16.4 13.1)	(-40.6 512.5)	(-1.3 5.2)	(-48.0 391.2)	(-118.8 16.2)
Observations	147	147	147	121	126	145	147

Panel A: Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This panel presents the coefficients and standard errors for the Local-Global Ethnolinguistic Complementarity variable from the second stage of 2SLS IV regressions estimated in Table 8. The column headings give the dependent variables for each of the columns. The Local-Global Ethnolinguistic Complementarity variable has been instrumented using a predicted Local-Global Ethnolinguistic Complementarity variable based on languages spoken in neighboring countries. Please refer to the notes from Table D.2 for further details. Panel B: The 95% confidence intervals are computed allowing for violations in the exclusion restriction following the local to zero approach of Conley et al. (2012) and using the Stata ado program "plusexog" created by Clarke (2014). The row which says CI ($2\delta = x\%$) reports the 95% confidence interval for the Local-Global Complementarity coefficient when we allow the direct effect of the instrument to be up to $x\%$ of the marginal effect of Local-Global Complementarity from the original instrumental variable specification.

How much this violation of the exclusion restriction affects our results for local-global complemen-

³³This requires first subtracting $\gamma_3 LGC_{nc}$ from both sides of the above equation.

tarity depends of course on the magnitude of δ . Table D.4 presents the results when we allow the direct effect of the instrument to be up to 10% to 50% of the marginal effect of local-global complementarity on public goods from the IV estimation. To interpret our findings, consider, for example, the row which says CI ($2\delta = 15\%$). It reports the 95% confidence interval for the local-global complementarity coefficient when we allow the direct effect of the instrument to be up to 15% of the marginal effect of local-global complementarity from the original instrumental variable specification. When comparing the results with our original findings, reported in the upper panel of Table D.4, we notice that the confidence intervals do not change the significance of our original results. Even if we allowed for a rather strong violation of the exclusion restriction so that the direct effect of the instrument were up to 25% of the marginal effect of local-global complementarity, our findings continue to be statistically significant.