Journal of Development Economics 164 (2023) 103113

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

Regular Article

A global map of amenities: Public goods, ethnic divisions and decentralization☆

André Seidel *

Department of Business Administration, Western Norway University of Applied Sciences, Haugesund, 5528, Norway

ARTICLE INFO

Keywords:
Public goods
Amenities
Decentralization
Ethnic fractionalization
OpenStreetMap

ABSTRACT

I analyze the effects of ethnic divisions on the provision of public goods in countries throughout the world. Using OpenStreetMap data, I construct a new, global dataset that pinpoints locations of public amenities, such as schools, hospitals, and libraries. I devise two new proxies to address the variability in the completeness of such data. I find that more autonomous subnational regions with a high degree of ethnic fractionalization provide significantly fewer productive public goods. Therefore, decentralization can decrease the provision of local public goods in areas characterized by higher levels of ethnic heterogeneity than those that prevail nationwide.

1. Introduction

Does ethnic heterogeneity prevent local governments from supplying public goods? A growing body of literature indicates a negative relationship between public spending and ethnic heterogeneity in several countries. High levels of ethnic fractionalization—the likelihood that two individuals randomly selected from a given population do not belong to the same ethnicity—have, for example, been associated with lower subnational spending on education in the U.S. (Alesina et al., 1999). Understanding the mechanism behind this phenomenon becomes increasingly relevant as ethnic heterogeneity within countries slowly increases internationally.[1] In addition, the ethnic composition of policymakers responsible for public spending keeps changing substantially due to institutional changes such as decentralization, particularly in developing countries.[2]

A central mechanism advocated by the literature is that social heterogeneity leads to a failure in the collective action to supply public goods (e.g., Alesina et al. (1999)). This is based on the idea that policymakers disagree on public spending if the proposed political actions become too heterogenous, as each policymaker will be aiming for policies tailored to their respective interest groups. Higher social heterogeneity, therefore, translates into a lower likelihood to agree on public spending, even if there is a need for higher spending within the collective. Studying this mechanism is challenging, as the collective action failures of policymakers are difficult to observe directly. What is observable to some extent is the outcome of political action. If the correlation between ethnic heterogeneity and public good supply originates mainly from a policy failure, then it should occur only where the supply of public goods is in the hands of politicians in ethnically heterogeneous regions. This paper is the first to test this hypothesis using a novel worldwide dataset on the supply of public goods at the subnational level. This novel dataset allows me to capitalize on the variation in subnational ethnic fractionalization and cross-country heterogeneity in the location of spending power across subnational units. Therefore, and in contrast to previous studies in the field, my design permits accounting for national omitted variables[3] and provides a global picture on the relationship between ethnic heterogeneity and public good supply. Furthermore, my design allows narrowing down possible mechanisms driving the effect of

☆ I thank Christian Lessmann, David Weil, Ferdinand Rauch, Hans K. Hvide, Maxim Pinkovskiy, Melanie Krause, Roland Hodler, Rune Jansen Hagen and Stelios Michalopoulos as well as participants at conferences in Glasgow (75th IIPF), Berlin (Annual meeting of the Research Group on Development VfS), Amsterdam (9th EUEA) and Hamburg (2nd Workshop on Geodata and Economics) for their valuable input. Furthermore, I want to thank fellow researchers who participated in colloquiums at the NHH, University of Bergen, University of Hamburg, Leibniz University Hannover, Ifo Munich, BI Oslo and Brown, as well as participants at the virtual annual meeting of the VfS, for helpful comments and suggestions.

* Corresponding author
E-mail address: Andre.Seidel@hvl.no.

[1] Measuring changes in ethnic heterogeneity over time is difficult; however, the existing data indicate an increasing trend toward more ethnic fractionalization across countries over the last 100 years. See, for example, Dražanová (2019).

[2] For a survey exploring the adoption of decentralization worldwide, see OECD (2019).

[3] Using national level panel data to account for national omitted variables is problematic, as ethnic heterogeneity is a slow-moving variable.

https://doi.org/10.1016/j.jdeveco.2023.103113
Received 1 February 2022; Received in revised form 2 May 2023; Accepted 6 May 2023
Available online 25 May 2023

ethnic heterogeneity on public spending.

This paper demonstrates that ethnic heterogeneity has a negative correlation with the supply of regional public goods in subnational regions worldwide. However, this relationship only emerges to be significant in subnational regions that are part of decentralized countries or have power over spending on public goods. I do not find a similar pattern in regions where fiscal power is located at higher or lower levels of government than the one studied. Both findings are consistent with the mechanism theorized by Alesina et al. (1999) that the effect of ethnic heterogeneity on public spending is dominated by the failure of politicians to supply public goods. Furthermore, in contrast to the previous literature, my findings indicate that the correlation between ethnic heterogeneity and spending on public goods is conditioned by political institutions. When there is no local power over spending, regional ethnic heterogeneity does not seem to be linked to regional public spending.

My analysis relies on a novel self-assembled dataset that contains the geocoded locations of various amenities that are closely linked to some of the core public goods that are typically funded through government spending, such as schools, libraries, hospitals and police stations. I use the volunteered, crowd-sourced data collected by the OpenStreetMap (OSM) project.[4] To study regional spending, I aggregate information at the regional level by counting the number of specific amenities in first-level administrative regions of countries. The new dataset covers 3352 regions in 204 countries. The number of public amenities per region is a proxy of the quantity and quality of public speeding. Greater local availability of public amenities usually results in higher welfare, as the consumption of the associated public goods becomes less costly. The literature concerning school attainment indicates that a primary driver of school attendance is the distance to schools (e.g., Duflo (2001), Burde and Linden (2013), Kazianga et al. (2013) and Muralidharan and Prakash (2017)). The literature concerning other public goods, such as public safety (e.g., Blanes i Vidal and Kirchmaier (2018)) and emergency health care (e.g., Buchmueller et al. (2006) or Wilde (2013)), shows that the response time is a key issue. The distance to amenities typically decreases as the number of amenities in a region increases. The number of public amenities is also a simple but powerful proxy for the quantity of spending on local public goods. For example, in the case of the U.S., the number of primary and secondary schools per school district can explain between 68% and 74% of the variation in district-level educational spending.[5]

Using my new dataset, I estimate the relationship between regional ethnic fractionalization and the number of different public amenities observed in first-level subnational administrative regions across countries. The cross-country, cross-regional setup allows me to use the variation in ethnic fractionalization among subnational regions of a country while controlling for country-level fixed effects. I show that regions in decentralized countries with high levels of social heterogeneity have a significantly lower supply of schools, libraries, and hospitals. The effect is large; for example, an increase in ethnic fractionalization by one standard deviation implies a decrease in the number of schools by 5% if spending power on educational issues is based at the regional level. Therefore, the average global correlation is much higher than the correlation reported by Alesina et al. (1999) in the U.S., where an increase in ethnic fractionalization by one standard deviation implies a decrease in the share of educational spending by 2%.[6] I conduct placebo tests using nonpublic amenities and show that ethnic fractionalization does not have a differential correlation with the supply of restaurants depending on whether they are located in a decentralized country. I also conduct placebo tests based on the heterogeneity in the decentralization of education systems, showing that my findings are not attributable to decentralization alone. I observe significant correlations only if the level of government under study has spending power, not if spending power is located at lower or higher levels of government. I further show that the negative association of ethnic fractionalization becomes more prominent with increasing exposure to decentralization. The negative association of ethnic fractionalization is smaller if subnational governments handle less revenue, if decentralization took hold only recently and if the school financing is largely provided by the government.

My findings are robust to a large number of robustness tests, such as the use of different indicators of decentralization and ethnic heterogeneity, as well as controlling for regional urbanization and development. I further develop two new indicators to account for the degree of completeness of the OSM data. These indicators allow me to correct the data at the regional level or to control for the degree of completeness in cross-country analysis. When comparing the corrected data with official data of a subset of countries for which subnational data exist, I observe country-level correlations greater than 90%. Using official and OSM data to study the determinants of the degree of completeness, I find that completeness is primarily driven by national fixed effects and that regional development plays only a minor role. The indicators of OSM completeness and national fixed effects explain between 85% and 95% of the variation in observable completeness. Therefore, the risk of estimation bias arising from OSM data collection is mitigated via the use of indicators of OSM completeness and national fixed effects. The main findings of the paper, however, do not depend on the use of these controls.

My findings contribute to the economic literature examining the potential effect of social heterogeneity in the context of the political economy of government spending.[7] Section 2.1 provides more details regarding the literature on social heterogenicity and the provision of public goods. I confirm the existence of a negative correlation between ethnic fractionalization and local public spending on a global scale. Furthermore, my findings provide evidence consistent with the theory that globally, the dominant mechanism is the failure of policymakers to collectively provide public goods. My findings also add to the general literature on the costs of social heterogeneity; for some recent contributions, see Bluhm and Thomsson (2020), De Luca et al. (2018) or Montalvo and Reynal-Querol (2020).

My findings also contribute to the economic literature concerning the costs and benefits of decentralization of government power. There is a long-standing discussion in the economics literature regarding the gains and costs of decentralization that dates back to the work conducted by Tiebout (1956), Musgrave (1959) and Oates (1972). Since their initial arguments, many scholars have identified various moderators that affect the outcomes of decentralization, such as the level of national development (Lessmann, 2012), the freedom of the press (Lessmann and Markwardt, 2010), the level of government tiers (Fan et al., 2009), and the quality of the government (Neyapti, 2006).[8] My findings contribute another layer to this literature by showing that the ethnic fractionalization of subnational regions determines whether decentralization leads to a reduction in the ability to provide public goods at the local level.

My study also contributes to the literature concerning the evaluation

[4] I do not use Google maps data because my analyses require storing individual map features, which is prohibited by the usage rights set forth by Google. Furthermore, Google maps data are fraught with similar, if not more complicated issues, such as the potential bias in data collection, for example, the bias in urban areas, that can arise from the commercial interests of Alphabet Inc.

[5] See Appendix B for a related estimate.

[6] In Appendix B, I show that comparing increases in budget and the number of schools is reasonable, as both are highly correlated. In the U.S., at the school-district level, this correlation is approximately 70%.

[7] See Section 2 for more details regarding the general literature. My findings also add to the general literature on the costs of social heterogeneity; for some recent contributions, see Bluhm and Thomsson (2020) or De Luca et al. (2018).

[8] For a more detailed summary of the literature concerning the impact of decentralization, see Martinez-Vazquez et al. (2017).

of volunteered geocoded data in geography. How to best measure the quality of volunteered geocoded data, such as the data available on OSM and other sources, remains an open question in the field of geography. Thus far, only a few studies have evaluated the various aspects of the quality of OSM data. Senaratne et al. (2017) recently summarized this literature. I add to this line of research by providing the first globally available set of indicators of subnational mapping completeness. I also test the reliability of these indicators in a subset of countries at different stages of economic and OSM development.

2. Previous findings and hypotheses

2.1. Previous findings on public spending and social heterogeneity

The idea that the provision of public goods may be hampered by social heterogeneities was a frequent topic of discussion in the public finance literature of the late 1990s, with seminal contributions by Easterly and Levine (1997) and Alesina et al. (1999). The fundamental criticism of these researchers is that the costs of a collective action such as spending on public goods might depend on the social heterogeneity of the groups involved. High social heterogeneity may therefore lead to an inefficient supply of public goods. In the case of the U.S., Alesina et al. (1999) and Alesina and La Ferrara (2000), for example, find evidence for the underprovision of productive public goods in regions with high levels of social heterogeneity, as measured by ethnic fragmentation.[9]

The link between ethnolinguistic fractionalization and the supply of public goods, such as education and health care, has also been found in other countries. The vast majority of these studies relied on cross-regional data on specific countries and public goods (e.g., Dayton-Johnson (2000) (water supply in Mexico); Miguel and Gugerty (2005) (education in Kenya); Khwaja (2009) (infrastructure in Pakistan) or Díaz-Cayeros et al. (2014) (a range of public goods in Mexico)). Some studies use national-level outcomes and study cross-country variation (e.g., Baqir (2002), Alesina and Zhuravskaya (2011) or Desmet et al. (2020)).

A small subset of studies relies on individual-level data from lab experiments and survey data. These studies also confirmed that socially heterogeneous groups have a greater tendency to mistrust one another and to fail in the provision of public goods (e.g., Glaeser et al. (2000), Bernhard et al. (2006) or Habyarimana et al. (2007)). Using a combination of survey and census data, Gershman and Rivera (2018) find that deep-rooted subnational ethnolinguistic and religious diversity is associated with less local public good provision in Sub-Saharan Africa.

Only a few studies have examined the potential effect of political institutions on the link between fragmentation and regional public good supplies. An alternative to utilizing cross-country variation in political institutions is to use variation within a country over time. Miguel (2004), for example, found a positive significant correlation of nation building with regional education spending in ethnically heterogeneous regions in Kenya and Tanzania between 1996 and 2002. Glennerster et al. (2013) found no significant correlation of ethnic fragmentation with regional public good supplies using data for regions in Sierra Leone before and after the civil war. Cinnirella and Schueler (2016) detect a significant positive correlation of centralization with educational spending in linguistically fragmented eastern border regions of Prussia between 1886 and 1896. Alesina et al. (2019) report a significant negative correlation of deforestation with the administrative reforms that reduced the ethnic diversity of regions in Indonesia between 2000 and 2012.

There are various limitations to the findings in the previous literature. Studies are usually limited to specific countries or small sets of countries, limiting the external validity of the findings. Furthermore, most studies can only rudimentarily account for omitted country-level variables. Furthermore, most studies cannot demonstrate the mechanism behind the correlation of ethnic fractionalization with public spending.

My present work overcomes the limitations of the previous literature, as I study public spending across subnational units and across countries, thus allowing me to account for country-specific effects. Based on a global dataset, my findings help shed light on the link between ethnic fractionalization and local public spending in developed and less-developed countries within one unified framework. Most importantly, by using variation in ethnic fractionalization and the location of spending power, my design allows me to narrow down the mechanism behind the correlation of ethnic fractionalization with public spending.

2.2. Hypothesis

The relationship between social heterogeneity and public spending is complex, and there are many potential mechanisms that can influence this relationship. This is partly because public spending is linked to economic development, which itself is affected by social heterogeneity in various ways (Montalvo and Reynal-Querol, 2020). The focus in this paper is on the potential social heterogeneity effects on the supply of public goods, resulting from the political economy of public spending. Based on the literature, I derived two opposing mechanisms that offer different explanations for how social heterogeneity affects local public spending through the political economy of public spending. (1) With increasing social heterogeneity, politicians may increasingly aim for special treatments of their peer group. For example, for educational spending, this might translate into more schools: Instead of one school for all groups within a region, there might be segregated schools for each group. This likely translates into higher overall spending due to the loss via economies of scale. (2) The opposite might be the case if budget restrictions affect collective action to decide public spending. In that case, social heterogeneity might lead to collective action failure, as different factions might not agree on a common policy.[10] Most importantly, both mechanisms are not mutually exclusive. The previous empirical literature suggests that the dominant mechanism seems to be collective action failure (see Section 2.1). Hence, I expect that social heterogeneity induces a policy failure that leads to lower levels of public spending. For this effect to occur at a regional level, spending power on public goods needs to be in the hands of regional politicians. My working hypotheses therefore is that the joint effect of regional social heterogeneity and regional spending power on regional public spending is negative.

To test my hypothesis, I use the variation of public spending with first-level subnational regions across countries. I assume that the agenda of regional policymakers is closely aligned with the interests of their respective peer groups. Hence, I assume that social heterogeneity within a region is an indicator of heterogeneity among regional policy makers. Furthermore, I assume that the number of amenities is a reliable proxy for related public spending.[11] My data needs are therefore threefold: subnational indicators of social heterogeneity, autonomy related to public spending and public amenity endowment.

[9] I can replicate their findings for the U.S. using OSM data; see Appendix B for details.

[10] For a more detailed discussion on both mechanisms, see Alesina et al. (1999).

[11] In Appendix B, I show that, at least in the U.S., the number of schools seems to be informative of the total educational spending per school district.

3. Data

3.1. A global map of public amenities

3.1.1. Data collection

I use the number of amenities of regions as reported in the OSM project as a proxy for public spending. The OSM project is a free, editable map of the world built largely from scratch by volunteers and released with an open-content license. By the end of 2017, the project had more than 4 million registered mappers, with an average of 40,000 people contributing data to the project per week.[12] The OSM project is the largest existing dataset of volunteered geographic information. The incredible success of the project is attributable to several factors, which have been well documented and discussed; see, e.g., Senaratne et al. (2017). One factor is that untrained people, regardless of their expertise and background, have been able to add geographic information since the start of the project.[13] This has allowed the OSM project to substantially increase its coverage in recent years, particularly in less-developed parts of the world. Different mappers and programmers associated with OSM have created visualizations that beautifully illustrate this growth.[14]

Data on the OSM project are provided by referencing latitude/longitude nodes, liens, or polygons and attaching attributes to these objects in the form of tags (e.g., "amenity" = "yes" and "building" = "pub"). I built the dataset using this information. I extracted all polygons, multipolygon relationships, liens and points and their locations, which carry tags associated with the various amenities under study from the OSM project, until the end of 2017. For example, to identify schools, I use the tags "amenity" = "school" or "building" = "school". Table 9 in Appendix C summarizes all tags used.

I use these data to calculate the main dependent variables of my analysis, which are the number of schools, hospitals, libraries and police stations recorded within first-level administrative regions as defined by the GADM dataset. Within the main analysis, I focus on the number of schools. This has several advantages: (1) for schools, data to cross-validate OSM records are available for a subset of countries (see Appendix A for details), (2) schools as proxies for educational spending allow for the closest comparison to previous findings (e.g., Alesina et al. (1999), Alesina and La Ferrara (2000), Miguel and Gugerty (2005) or Desmet et al. (2020)), and (3) there exist data to measure the fiscal decentralization of educational systems.

3.1.2. General data quality issues and initial cleaning of the raw data

Using volunteered geocoded information generally has some drawbacks. Senaratne et al. (2017) summarized the current strand of the geography literature on the various quality issues associated with volunteered geocoded information. When examining economic geography, some issues are less important than other issues. For example, topological consistency (e.g., whether objects overlap) and positional accuracy (e.g., whether objects are half a meter further south or north) are not of high importance for applications that typically interest economists. However, there are other issues, such as thematic and semantic accuracy, that require further discussion.

It is well known that tags are not consistently used in the OSM project since people are free to define new tags. To address this problem, the OSM project has set guidelines on how and where to tag common objects, such as public amenities. The selection of tags I use to identify different amenities is based on these guidelines. Beyond the wording used in the different tags, they can be placed on different objects; for example, sometimes only the wall of a school is tagged with "building" = "school", and sometimes the relationship between various objects that form the school is tagged with "amenity" = "school". To avoid the resulting double counting (e.g., each school yard wall being counted as a separate school), I merge all objects with the same tag within a 100-m radius into one observation.[15]

One final general issue worth mentioning at this point is that in the OSM project, tags defining whether an amenity is public or non-public are rare. While this may not be a problem for police stations, it could be an issue for schools if we aim to study effects only relevant to one supply form. I will come back to this issue in relation to my analysis at a later point.

3.1.3. Completeness

A quality dimension critical for the application of OSM data in the study of economic geography is completeness. It is to be expected that, depending on the popularity of the project, the OSM data do not record all existing amenities. Various issues could determine the magnitude of this effect, such as the lack of Internet access or legal boundaries. In the case of China, for example, mapping by private individuals is illegal.

The descriptive statistics of the cleaned raw data can provide an initial impression of the data and the potential extent of missing sites. Fig. 1 provides a look at the data after the initial cleaning. The figure displays all of the schools in the OSM project by the end of 2017 as a 50-m radius dot. At first glance, it is encouraging to see the close resemblance of Fig. 1 to night light images and population density maps.

At closer look, however, one can spot some unusual patterns, such as the large number of schools in Uganda. One explanation for this finding might be that the Humanitarian OpenStreetMap Team (HOT)[16] has a large and successful project running in Uganda in response to the ongoing refugee crises. As shown in detail in, Appendix A despite the incredible increase in data volunteered in recent years, OSM data remain incomplete in many dimensions. In this sense, Uganda is most likely an outlier, with relatively more data than other countries in Africa. An example of a possible outlier with less data might be China or North Korea, where the number of schools appears to be very small. Simple descriptive statistics based on the cleaned raw data give weight to these concerns.[17] Descriptive statistics imply that there are 0.5 schools per 1000 citizens in the state of New York but just 0.02 schools per 1000 citizens in the province of Shanghai. Nevertheless, as Fig. 1 suggests, the distribution of schools across countries does not seem to be dramatically distorted. In China, for example, many schools are obviously missing, but the allocation shown still seems plausible; the greatest density of schools in the OSM data is in the heavily populated eastern regions of China. Therefore, the absence of certain items may be mostly driven by country-level effects.

3.1.4. Approximating the regional OSM completeness

A detailed discussion of how I derived a proxy for mapping completeness is documented in Appendix A. Here, I will only briefly summarize the essential ideas and concepts behind the proxy and direct the reader to Appendix A for a more elaborate documentation.

I assume that mapping occurs in two phases. Mapping within regions without any data in the OSM project starts by adding fundamental landmarks, such as roads. In the second phase, detailed data, such as socioeconomic features (e.g., schools, police stations, cinemas, and restaurants), are added. If so, then the probability (p^{I+II}) that a specific amenity is recorded in the OSM project is the product of the probability

[12] https://wiki.openstreetmap.org/wiki/Stats.

[13] To see this point demonstrated, go to (https://wiki.openstreetmap.org/wiki/Beginners%27_guide), and see how easy it is to add something.

[14] For example, see http://tyrasd.github.io/osm-node-density/#2/19.1/21.4/latest and.https://www.youtube.com/watch?v=AM2fMJedqAc.

[15] Obviously, this may create some errors because in densely populated regions, public amenities could be in such close proximity that they are counted as one although there are actually two or more amenities. However, changing the radius to 50 m does not change the results. For restaurants, the radius is reduced to 10 m.

[16] See https://www.hotosm.org/for more details.

[17] See Table 12 for more descriptive statistics from the raw data.

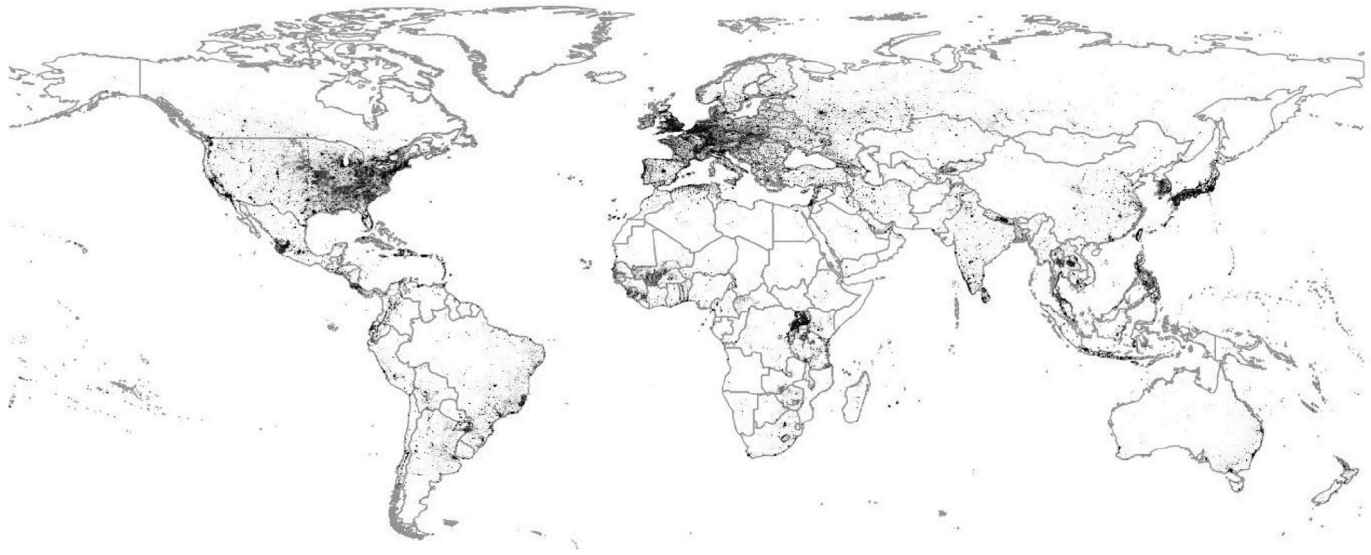

Fig. 1. Schools in raw OSM data.

that phase one (p^{I}) and two (p^{II}) have occurred.[18]

Based on these assumptions, deriving a regional, amenity-specific proxy for the degree of completeness of OSM data requires finding proxies for the two phases of mapping. The two proxies I propose are somewhat based on the same general idea, which is to select objects associated with a specific phase of mapping and compare their numbers in OSM to that in a potentially unbiased source. However, the challenge lies in finding representative objects for which also an unbiased source exists.

To derive a proxy for the completeness of the first mapping phase, I focus on one of the most fundamental landmarks in volunteer mapping: residential roads. To obtain an indicator of the existence of residential roads in the regions, I use the GHSL (2015) dataset. First, I define each area\pixel in the GHSL (~one km^2) with urban buildup and more than 100 inhabitants as a settled area. Second, I assume that at least one residential road must exist within such a settled area. Hence, if no residential roads are recorded in the OSM data in a settled area, it is likely that the first phase of mapping has not occurred in this location. My proxy for the completeness of the first phase of mapping is then simply the share of the total settled area of a region with residential road records. A proxy for the extent to which the second mapping phase has been realized in a region can be obtained based on a similar logic. However, in the second phase, the subgroup of objects of interest is fixed because the aim is to obtain an amenity-specific indicator of the completeness of the second phase of mapping. This is a problem because no good proxy is available for the true number of amenities in a location. Therefore, I erroneously assume that each area that has undergone the first mapping phase should contain at least one target amenity. I further assume that the degree of error induced by this assumption is country specific. My proxy for the completeness of the second phase of mapping of a region is then the share of the total settled area with records of OSM residential roads that has a record of at least one target amenity.

In Appendix A, I show that despite the obvious issues with the various assumptions made to derive the proxies for mapping completeness, they seem to perform very well in predicting deviations in countries/regions where official and OSM amenity data exist. I tested the robustness of the indicators in a subsample of countries, such as South Africa, the U.S. or Namibia, where official data on the allocation of schools are available. My findings demonstrate that in combination with country fixed effects, both indicators of mapping completeness can jointly explain more than 95% of the difference between the number of schools recorded in the OSM database and the number of schools in official registries. In fact, a large part of the variation can be explained entirely by country fixed effects, as speculated when discussing Fig. 1. In Appendix A, I describe the two indicators for mapping completeness in detail, present robustness tests and discuss how to best use both indicators to reduce a potential omitted variable bias arising in estimates using regional OSM data.

3.1.5. Regional ethnic heterogeneity

Among the various dimensions of social heterogeneity, ethnic heterogeneity has been shown to be widely important to various economic outcomes, such as growth or the likelihood of civil conflicts (Montalvo and Reynal-Querol, 2005). In line with conventions in literature, I use the following two indicators: ethnic fractionalization and polarization. Both indicators rely on the number of people belonging to different ethnicities in a country or, in this study, in regions within a country as a measure of ethnic fragmentation. The main difference between the two indicators is how the population weights contribute to the indicator. The general rule of thumb is that, in the case of the fractionalization indicator, large groups contribute more than their relative size to the indicator, while the opposite is the case for the polarization indicator.

Defining $\pi_{e,r}$ as the share of people belonging to group e in region r that hosts m ethnic groups, ethnic polarization is measured by

$$Ethnic\ Pola_{e,r} = 1 - \sum_{e=1}^{m} \left(\frac{1/2 - \pi_{e,r}}{1/2} \right)^2 \pi_{e,r} = 4 \sum_{e=1}^{m} \pi_{e,r}^2 \left(1 - \pi_{e,r}\right) \quad [1]$$

and ethnic fractionalization is measured by

$$Ethnic\ Frac_{e,r} = 1 - \sum_{e=1}^{m} \pi_{e,r}^2 = \sum_{e=1}^{m} \pi_{e,r} \left(1 - \pi_{e,r}\right) \quad [2]$$

Ethnic fractionalization has a very intuitive interpretation. The indicator measures the probability that two randomly selected individuals are not of the same ethnicity. In contrast, the polarization indicator

[18] Support for this model is derived not only from observations of the evolution of OSM data over time but also from the guidelines provided by the OSM wiki. Under the rubric mapping techniques (https://wiki.openstreetmap.org/wiki/Mapping_techniques), there is the following text: "Mapping is done in two steps: First, you need to know where things are, mainly the streets and ways. Then you need to know what there is, namely the POIs, street names and types. You can do these one after another, or both at the same time, but you can hardly do the what before the where."

measures how close the distribution of the ethnic groups is to a bipolar distribution. Hence, high values of the polarization indicator correspond to cases in which there is an ethnic majority that is challenged by a unified "large" minority. For an in-depth discussion of the origin and uses of both indicators, see Montalvo and Reynal-Querol (2005).

In the previous literature, ethnic fractionalization is the indicator of social heterogeneity that is most commonly used when studying collective action failure. Hence, the main analysis focuses primarily on ethnic fractionalization as a measure of social heterogeneity and uses ethnic polarization to test the robustness of the findings.[19]

The data of the population belonging to different ethnicities are derived from a combination of gridded population data from the 2015 GHSL dataset and the ethnic homeland boundary data from the Georeferencing of ethnic groups (GREG) dataset, provided by Weidmann et al. (2010). The GREG database reflects the distribution of ethnic groups worldwide in the 1960s and is based on a digitized version of the classic Soviet Atlas Narodov Mira. GREG documents the location of 928 ethnic groups in 8969 homelands. These homelands are projected to the current political boundaries of the first-level subnational administrative regions as defined by the GADM dataset. This approach creates 23,874 regional homelands within 3219 regions.[20] For 2658 of these regional homelands, GREG reports more than one ethnicity residing in the area. For these regions, it is not possible to assign their population to a specific ethnicity.[21] These multigroup homelands are spread across 1044 of the 3219 regions for which OSM data are available. Applying a strict exclusion criterion would therefore ultimately decrease the sample size by one-third. Furthermore, it is likely that regions that contain homelands in which multiple ethnicities reside are also regions with higher levels of ethnic heterogeneity. Excluding these regions from an analysis, therefore, might induce a sample selection effect. For the main estimates, I therefore do not restrict the sample. However, I provide robustness tests where I omit regions that have more than 10% of the population living in multi group homelands and where I omit all regions that contain multi group homelands. Furthermore, the results do not depend on how the population residing in the multigroup homelands is allocated to the different ethnicities when calculating the social heterogeneity indicators.[22]

Ethnic heterogeneity has thus far mostly been studied at the national level or the regional level within selected countries. Therefore, the question arises of whether there is a meaningful difference between regional and national ethnic heterogeneity. To visualize this difference, Fig. 2 displays the difference between national and regional ethnic fractionalization.[23] It is clear from Fig. 2 that there are substantial differences in the degree of regional ethnic fractionalization within countries. These differences can go in both directions; in Brazil, for example, most of the regions are more fractionalized than the overall country, while the opposite is the case for India, where the regions are much more homogeneous than the overall country.

Recently, the concept of measuring social heterogeneity in the context of political economy has been challenged by Desmet et al. (2012) and Desmet et al. (2020). They argue that traditional measures do not account for local interactions that may induce effects counteracting those of traditional measures. To account for these local effects, Desmet et al. (2020) construct an indicator of aggregated local language heterogeneity with the goal of complementing traditional measures. This additional measure is intended to capture how likely it is to actually encounter someone from a different language group in a person's day-to-day life. Using this new measure, they show that national local linguistic heterogeneity drives national public good spending upwards in contrast to overall linguistic heterogeneity. The mechanism they propose driving their findings is that antagonism toward other groups decreases with more frequent encounters with others, which in turn facilitates collective action. The focus of my research is to study the effect of ethnic heterogeneity within subnational regions in the context of the costs of decentralization. Within this framework, studying the additional effect of local interactions is intriguing but goes beyond the scope of this paper.[24]

3.1.6. Regional autonomy on public spending

The analysis focusing on the supply of schools needs to be based on a reliable indicator of autonomy regarding educational spending. I hand-collected information on whether spending power is located at the level of the central, regional and/or local government for as many countries as possible to build such an indicator. The regional level refers to the first subnational administrative units and local levels to anything beneath. The indicator is of simple design and indicates only whether a subnational unit has power over educational matters, without specifying the degree or type of power. To measure power over educational issues, I utilize various sources on the financing of education within countries. I assume that a subnational unit has power over educational matters if it provides a large share of the funding of the education system. In the case of OECD countries, the classification is based on the final funds after transfers devoted to education at the different levels of government, as reported in the Education at a Glance Indicators 2019. For other countries, the classification is based on the description of educational funding structures documented in different World Bank Public Finance Reviews and scientific papers. In cases where precise shares of funding were available, the cutoff that defines large used was 40%. As the indicator is based on the funding of education, it is a de facto measure of educational decentralization. Such measures are often tainted by various problems; see, for example, Lessmann (2009) for an in-depth discussion. The main reason for utilizing this measure is simply the availability of data. The indicator covers a total of 124 countries, of which 31 are classified as countries with educational power located at the regional level; see Table 59 in Appendix C for details.[25]

When studying other amenities, I have to rely on general indicators of decentralization, as deriving a specified indicator of decentralization, for example, of public safety spending, is not possible due to data limitations. The indicator used for the main analysis is the federalism indicator described by Treisman (2008), which is available for 155 countries, of which 21 are classified as federal. Using this indicator results in a substantial loss in precision. Being a federal country increases the likelihood that public spending is at the hand of first-level subnational regions but is not a guarantee. My indicator for educational

[19] Alesina et al. (1999), in their seminal paper, already discussed the effect of polarization. Given the available data, however, they only tested for the effect of fractionalization.

[20] To minimize measurement error, regional homelands with a population smaller than one are excluded.

[21] Gridded population data are taken from GHSL (2015), 1000-m resolution image.

[22] For the main specification, the assumption is that the first named group in a multiple group homeland is the dominant one, and the population of the homeland is added to the total population of this group. The results do not depend on this assumption. Allocating the population of multigroup homelands equally among the named groups or with the same shares as in the rest of the region yields the same results.

[23] The picture does not change when examining the level of regional fractionalization or polarization or the difference between national and regional polarization; see Figs. 5–7 and 6 in Appendix C. Note that in the figures, regions with ethnic homelands that have residents belonging to multiple ethnicities are not omitted.

[24] For such an analysis, the raw data of Desmet et al. (2012) would need to be used to reestimate the aggregated local heterogeneity indicators (developed by Desmet et al. (2020)) on the subnational level.

[25] In my final estimates, I can only utilize data on 121 countries, as observations had to be omitted due to singleton observations. This omission typically affects countries, such as Bosnia and Herzegovina, with few regions to begin with.

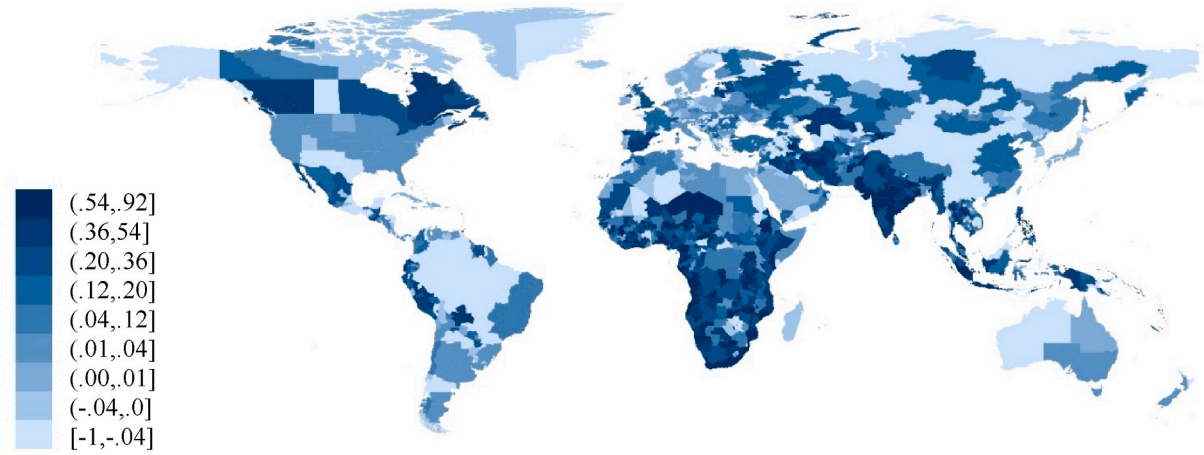

Fig. 2. National - Regional ethnic fractionalization.

spending shows, for example, that a larger number of nonfederal countries such as Poland or Norway have a decentralized education system. Furthermore, in some federal countries, education is decentralized to low levels. In the U.S., for example, decisions on spending are mainly located at the local school-district level, while the central government and the states have limited power on spending decisions.

3.1.7. *Descriptive statistics*

Table 15 in Appendix C presents summary statistics for the entirety of the data collected that is used in the main estimates.[26] Due to the availability of some of the data, especially the decentralization indicators, the number of observations varies considerably between the different variables. Here I only shortly highlight some key facts related to the data collected only for this paper; for more detail on, for example, the controls, please see Table 15 in Appendix C.

The number of amenities recorded at the region level differs considerably given the heterogeneity in the size of first administrative region, ranging from 0 to 21,073 in the cases of schools. Further, quite naturally, among the different amenities, the number of amenities recorded differs; usually, there is less need for police stations than for schools. An average region that hosts 164 schools only hosts approximately 30 hospitals, 17 libraries and 25 police stations. This difference, however, cannot fully be explained by a difference in the need for different amenities. This is most apparent when looking at the number of restaurants, which is still the second most recorded amenity. It seems reasonable to assume that there are more restaurants in the world than schools.

The frequency of missing data varies also considerably when looking at the proxies for mapping completeness. In some regions, the proxies indicate that 100% of the data is missing, and in others, zero.[27] Phase one mapping, on average, was completed at a 67% level, and phase one and two specifically related to schools at an 11% level. The latter highlights the potential need for controlling mapping completeness when estimates are based on OSM data. For a detailed discussion on the issue of missing information in the OMS data, please see Appendix A.

The indicator for the decentralization of decision power on educational matters indicates that 23% of the regions are autonomous, while only 16% are in a federal country. Furthermore, ethnic fractionalizations at the regional level seems to be smaller, with an average of 0.14, which is less than the one commonly observed at the national level (0.30).

To glance more generally at the data beyond the scope of this paper, it seems appealing to look at the number of schools per 1000 inhabitants. Doing so, I observe quite large differences between countries at different stages of development; see Table 11 in Appendix C for details. In part, this difference is likely driven by differences in OSM data collection. The analyses conducted in Appendix A suggest that amenity allocation within a country is much less distorted by missing data than amenity allocation across countries. This is also supported by the descriptive data. While the difference in means is quite large between countries at different development levels, the coefficient of variation is quite similar. What is interesting is that the variation in school provision is much smaller than the variation in development measured by the mean of light per region. Nevertheless, despite similar variation in development between regions, the variation in school provision in developed countries is the lowest.

4. Estimation approach

The units of observation in my study are first-level subnational regions. This design allows me to separate effects from general country-level effects. The two main estimation equations are

$$\ln\left(A_{OSM,i,r,j}\right) = \alpha + \beta_1 Hete_{r,j} + \zeta \mathbf{Z} + \mu_{i,j} \quad [3]$$

and

$$ln\left(A_{OSM,i,r,j}\right) = \alpha + \beta_1 Hete_{r,j} + \beta_2 Hete_{r,j} \times Auto + \zeta Z + \mu_{i,j} \quad [4]$$

where $A_{OSM,i,r,j}$ is the number of different public amenities i in subnational region r in country j, $Hete_{r,j}$ is a measure of social heterogeneity, $Auto$ is a measure of the degree of local autonomy, $\mu_{i,j}$ are country fixed effects, and $\boldsymbol{Z}$ is a vector of regional-level controls.

Studying this phenomenon in subnational regions lends itself to the test of the main hypothesis, since the expected policy failure should occur only if regional politicians have power over public spending. To test this, I estimate [3] for different subsamples of regions that are part of countries where spending power is located in first-level subnational regions and those where this is not the case. The main prediction is that β_1 is significant and negative only in regions with spending power. Using [4], I also test for the significance of the observed differences between the different samples. For schools, secondary data are available to a larger extent than other amenities. Therefore, the main analysis in Section 5.1 focuses on the log number of schools as the dependent variable because this allows me to run a multiple robustness test. To perform the central sample splits, I use the indicator for educational decentralization discussed in Section 3.2. Using general decentralization indicators, I expand my analysis to other amenities in Section 5.2. The

[26] Please see Table 10 in Appendix C and Section 3 for the definition and source of the variables used.

[27] I omit such regions in the main estimates.

main analysis uses ethnic fractionalization as a measure of social heterogeneity. In the robustness part of Section 5.1, I also present the main estimates related to the alternative use of a polarization indicator.

My estimation approach will not result in causal evidence for the discussed mechanism. My findings can only be seen as evidence in line with the theory of collective action failure in the presence of social heterogeneity. Nevertheless, several measures are taken to reduce the risk of the most plausible omitted variable or measurement biases. In addition to using a large set of controls, various sample splits based on theoretical moderators and placebos are discussed to support the claim that findings can indeed be traced back to a collective action failure among policymakers in charge of supplying public goods. The details of these tests are discussed alongside the main findings in 5.1 and 5.2.

The set of controls always includes country-level fixed effects μ_{ij}. Hence, the potentially identifying variation comes from the differences between regions within a country. The vector of regional-level standard controls $\boldsymbol{Z}$ contains the log of regional population, land area, night lights and the share of urban population. Larger regions might have a greater likelihood of being ethnically fractionalized (correlation 0.23). If so, the findings may reflect a simple size effect in regions that are a part of a federal state. Therefore, I always control the regional population size and the land area of a region. There is, furthermore, a considerable body of literature initiated by Easterly and Levine (1997), among others, arguing that ethnic heterogeneity is generally associated with less development. Hence, I control for the level of night light as a proxy for regional development to ensure that it is not an indirect development effect driving my main findings. A recent contribution by Eberle et al. (2020) highlights that ethnic heterogeneity is a driver of urbanization. Additionally, Nagaraj (2021) shows that mapping intensity in phase two is higher in urban areas than in rural regions given the same degree of phase one completeness. Furthermore, school size might differ systematically between rural and urban regions. I therefore control for the degree of regional urbanization.

There are considerable omissions in the OSM data, as revealed by the analysis in Appendix A. Therefore, using OSM amenity data to study the allocation of amenities across regions must be performed with caution. On a theoretical level, some could argue that the provision of data to the OSM project might be hampered by ethnic heterogeneity. The decision to add data to the project, however, lies with the individual mappers, and collective actions may indirectly affect only the data provision. One potential channel could be the infrastructure needed to add data, such as mobile networks, which are provided by collective action on a regional level. If this would drive my findings, the mechanism would still be the same, I would just focus on the wrong public goods. A further mechanism might be that ethnic conflicts in very heterogeneous regions prevent mappers from collecting data outside their homeland. However, if any of these concerns were valid, I should not find different results for different public goods, which I do (see Section 5.2). Furthermore, I should observe significant correlations with nonpublic goods such as restaurants as well, which I do not (see Sections 5.1 and 5.2). I also should not find different effects between regions facing different exposure of local autonomy, which I do (see Sections 5.1 and 5.2). In addition, the estimates discussed in Appendix A.3 show no evidence that regional ethnic fragmentation or the degree of decentralization impacts mapping completeness in countries for which official data on the allocation of schools across regions are available.[28] However when using my proxies for mapping completeness that are available globally, I see a significant and negative correlation of ethnic fragmentation with the completeness of the first phase and second phase of mapping p^{I+II}. This correlation seams further to be larger in decentralized educational systems; please see Table 14 discussed in Appendix A.3.[29] This would suggest that without accounting for mapping completeness, I might over estimate effects. I therefore focus in the main analysis on estimates based on using my proxy for the completeness of the first phase and second phase of mapping p^{I+II}. Nevertheless, to show that the findings do not depend on the assumptions associated with the different proxies for completeness, I also present estimates based on the raw data without controls and when controlling for the completeness of the first phase of mapping p^{I}.

5. Results

5.1. Main results

Table 1 reports the main results of the paper. The complete set of control variables is accounted for but not reported in all estimates. Control variables include country-fixed effects and the log of population, land area, night lights, the share of urban population and the proxy for completeness of phases one and two (ln (p^{I+II})) of mapping, as defined in Appendix A.[30] The sign of the covariates is within expectations for population, land area, and light, increasing the number of amenities. The opposite is the case for the share of the urban population. The proxy for the completeness of phases one and two of the mapping has the greatest explanatory power and is approximately .8. Second in line in explanatory power is population, followed by the level of night light. Overall, the explanatory variable and the covariates explain a large part of the variation in the number of amenities between regions within countries, with a within R^2 ranging around 90%.

The determinant variable in Table 1 is the log number of schools within first-level subnational regions.[31] Columns (1) and (4) report

Table 1
Main estimate.

Dep. Var.: ln (N Schools)	(1) Full sample	(2) No regional power	(3) Regional power	(4) Full sample
Ethnic Frac.	−0.105	−0.040	−0.267***	−0.040
	(0.073)	(0.089)	(0.095)	(0.089)
Regional education power X Ethnic Frac.				−0.228*
				(0.129)
# Countries	121	90	31	121
# Regions	1671	1236	435	1671
Within R-sq.	0.888	0.877	0.917	0.889

Note: The unit of observation is the first-level administrative region. The dependent variable is the log of the number of schools in OSM in active OSM areas. Ethnic Frac. Is regional ethnic fractionalization based on GREG and GHSL data. All estimates include country fixed effects, the log of the regional population, land area, night lights, the share of urban population and the log of the proxy for completeness of phases one and two (ln(p^{I+II})), as defined in Appendix A. Estimates are based around the indicator "Regional education power" that is equal to one if first-level administrative regions have a large say in educational decisions and zero otherwise. The sample used in columns (1) and (4) is based on the set of countries for which this indicator is available; in column (2), those where the indicator is zero; and in column (3), those where the indicator is one. In column (4), all regressors are interacted with the indicator, but the interactions with controls are not reported. Standard errors are reported in parentheses and are clustered at the country level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

[28] See Table 13 in the appendix.

[29] Estimates are based on the same sample and proxies for local autonomy and social heterogeneity that are used to derive the main results of the paper.

[30] Please note that to remain consistent with the theory underlying the approximation approach of the proxies for the completeness of the OSM data, I restrict the observations to those within active OSM areas.

[31] Note that regions with less than one observation in the OSM project are omitted. Table 16 in the appendix summarizes the main descriptive statistics for the baseline dataset.

estimates for regions for which data on the location of educational power is available. Column (3) reports estimates for regions in countries where educational power is located at the first-level subnational regions,[32] and column (2) for countries that do not have control over educational matters located at the regional level. Looking in column (1) at the full sample, I find a negative but insignificant association between ethnic fractionalization and the number of schools within a region. Column (3) shows that when educational power is located at the regional level, ethnic fractionalization has a significantly negative relationship with the number of schools within a region. The coefficient further suggests a sizable decrease of 5% in the number of schools, given an increase in ethnic fractionalization by a standard deviation (0.19). When focusing on regions in countries with no power over educational matters located at the regional level (column (2)), I find a substantially smaller correlation that is not significant, despite a larger sample. The difference between the estimates for both subgroups of regions is also significant (column (4)). This finding is in line with the expectation based on the theory that ethnic heterogeneity reduces the ability of policymakers to realize public spending. This effect can only occur in the presence of regional decision power. What is notably important, however, is that the theory suggests that fragmentation might also induce more group specific lobbying and paternalism, leading to the provision of more schools. In this sense, estimates have to be interpreted as a lower bound of the potential effect of failing to collectively supply public goods in the presence of ethnic heterogeneity. Unfortunately, the current data does not allow me to separate these two effects to obtain an estimate of the upper bound as well.

5.2. Robustness tests

The main estimate is based on various assumptions on the proxy for the number of amenities and the measures of ethnic heterogeneity. To reduce the associated risk of measurement-bias, several robustness tests have been conducted.

To account for a potential bias arising from the **collection of OSM data**, all estimates presented in Table 1 are replicated using the raw OSM data (Table 17, Appendix C)[33] and only the indicator for completeness of the first phase of mapping ($\ln(p^{I})$) (Table 18, Appendix C). In both sets of estimates, the results are qualitatively the same. The findings differ in two dimensions. First, the negative relationship between the number of schools and ethnic fractionalization if educational power is located at the regional level is significant at lower levels. The same is true for the difference in the relationship between regions with or without educational power. Second, coefficients are, on average, twice as large as in the main estimate. There are two possible reasons for the difference in coefficient magnitude. First, it is possible that if the degree of completeness is not considered, the effect sizes are overestimated, which could be the case if the degree of completeness is negatively affected by ethnic fractionalization and decentralization. The estimates presented in Appendix A.3 suggest this could be the case. They further suggest a bias roughly the size of the observed reduction in coefficient size when controlling for completeness of phases one and two ($\ln(p^{I+II})$); see Table 14 (Appendix C) for the related estimates. That the entire correlation is purely a result of measurement bias, however, is combatted by the placebo test on the number of restaurants discussed at the end of the next section.[34] Second, the effect size may be underestimated when I control for the completeness of mapping of phases one and two since the proxy relies on the assumption that areas that contain at least one amenity are representative of the region. Hence, the effect of ethnic fragmentation and decentralization on the number of amenities outside these cells may not be accounted for. The total effect is underestimated if this effect works in the same direction as in the representative cells. This issue is not the case when controlling only for the first phase of mapping, where coefficients seem to be systematically closer to estimates based on the raw data.

Using GREG data to measure ethnic heterogeneity comes with the issue that ethnic homeland maps overlap in some cases, creating **multi-group areas**. In the main estimate, the population residing in such multi-group homelands is added to the population of the first ethnicity named in the GREG data. To account for a potential bias arising from this assumption, all estimates presented in Table 1 are replicated using only the regions that contain no multi-group areas (Table 19, Appendix C) and regions where less than 10% of the total regional population reside in multi-group areas (Table 20, Appendix C). In both sets of estimates, the results are qualitatively the same. The only difference is that the observed difference in the estimates for regions with and without power on educational issues is not significant at conventional levels.

Finally, I address whether the findings depend on the measure of social heterogeneity used. Table 22 in Appendix C summarizes the baseline estimates analogous to Table 1. using **ethnic polarization** as an indicator of social heterogeneity. Comparing both sets of results reveals very few differences. The results based on ethnic polarization are slightly weaker, and the coefficients are somewhat smaller; otherwise, the results are very similar. The difference in coefficient size might be related to the positive effect that social heterogeneity might have on public spending via patronage. If the ability to realize patronage is linked to group size, then polarization is more likely to detect this effect than fractionalization. This is induced by the design of both heterogeneity measures. Polarization tends to be higher than fractionalization if a few groups are larger than most other groups. In regions with many ethnicities where a selected few are large enough to lobby successfully for group-specific policies, the collective action failure effect might be partly compensated by the patronage effect. This issue requires further analysis, which is beyond the scope of the current paper.

The findings are subjected to a large set of additional robustness tests.[35] Omitting all observations where less than 75% of phase one of mapping is realized does not change the main finding, despite a considerable loss of observations (Table 21 Appendix C). Capital regions might be unique for various reasons. Adding a dummy for capital regions does not change the main result.

5.3. Potential mechanisms

The data at hand does not allow for direct observation of the mechanism through which the provision of public goods is changed in the presence of social heterogeneity. A causal identification of the mechanism at play is, therefore, not possible. Yet, when the failure of collective actions causes the main findings, this should affect other observations as well. In what follows, I discuss whether the assumption of a causal effect is in line with other observations (i.e., correlations that should or should not exist, depending on whether the assumption is correct). These estimates cannot be seen as direct proof but as indirect evidence consistent with the theory.

Depending on the **prevalence of private education**, the effect

[32] Note that given the large set of controls, regions from two countries with educational power at the regional level are omitted, as they create singleton observations.

[33] Please note that the raw OSM data is not restricted to observations within active OSM areas.

[34] Additionally, for the subset of countries where I can measure the actual degree of missingness, I do not find a significant correlation of ethnic fractionalization with this variable. See Table 13 in Appendix C for more details.

[35] Notably, some potential omitted variables are already addressed by the country fixed effect included in all estimates.

might be smaller if schooling is predominantly not collectively provided by the government.[36] Hence, I should observe a larger correlation if the share of government contribution to education is higher. To test this, I replicate in Table 2 the estimates from Table 1, splitting the sample of countries with regional education power into countries with or without more than 90% governmental contribution to educational expenditures. The findings suggest significant and large differences between both groups. The association between ethnic fractionalization and the number of schools in a region is almost twice as high in regions with predominantly government-financed education than in the main estimate.

Suppose the effect of ethnic fractionalization is the consequence of a collective action failure among policymakers. In that case, it stands to reason that the effect should increase with **more prolonged exposure** to the policy failure. To test this, I replicate in Table 23 (Appendix C) the estimates from Table 1 using only regions with educational power and split them into those with or without power before 2000. For regions with more prolonged exposure to power on educational issues at the regional level, I observe a larger negative relationship then for regions that were less exposed.[37]

The ability to drastically change overall spending might also depend on the **degree of de facto decentralization**. With more spending realized at the regional level, the effects should be larger. To test this, I replicate in Table 24 (Appendix C) the estimates from Table 1, splitting the sample between countries where institutions controlled by first-level subnational regions provide less than 30% of the educational spending and the ones where this is not the case. The correlation between ethnic fractionalization and the supply of schools tends to increase when omitting countries with de facto weak decentralization of educational spending power. The correlation remains further significant despite the considerable reduction in sample size; however, the differences are not significant.[38]

I further conduct two **placebo tests** that assess alternative mechanisms that might drive the main finding. The first placebo test is linked to two alternative mechanisms that might indirectly induce an adverse effect of ethnic fractionalization on the observed number of schools, impaired development and OSM data collection. Ethnic fractionalization might reduce regional development and therefore the ability to spend on public amenities. Ethnic fractionalization might also reduce incentives for volunteer mapping. Consequently, I always try to control for both effects in the main estimates by using light as a proxy for regional development and by using the indicator of mapping completeness. To ensure further that the indirect effects of ethnic fractionalization do not drive the results, I use the number of restaurants to conduct a placebo test. Suppose the main results are induced by the indirect effects of ethnic fractionalization and not the collective action failure among policymakers. In that case, the **number of restaurants** should be affected similarly to the number of schools. For example, if mapping is hampered by ethnic fractionalization, we would expect to find a decrease in the number of both schools and restaurants in regions with high ethnic fractionalization. Therefore, Table 3 replicates the estimates of Table 1 using the log number of restaurants as a dependent variable. In all estimates, ethnic fractionalization negatively correlates with the number of restaurants, which, however, never reaches significance. There is also no significant difference in the correlation between regions with and without educational power.

The second placebo test aims to ensure that findings are not driven by indirect effects of decentralization or omitted variables that led to decentralization. The devolution of power to the regional level is collinear with other transitions, such as the shift toward more private education or a general increase in state capacity that moderates the effect of ethnic fractionalization. If this is the case, ethnic fractionalization should also reduce the supply of schools in regions in countries where **power over educational decisions** is decentralized but to a **lower level of subnational administration**, as in the U.S.[39] or Norway. To

Table 2
Main estimate (more than 90% government financing of schools).

Dep. Var.: ln (N Schools)	(1)	(2)	(3)	(4)
Ethnic Frac.	−0.220	0.060	−0.454***	0.060
	(0.141)	(0.083)	(0.063)	(0.078)
90% gov. edu.				−0.514***
X Ethnic Frac.				(0.098)
# Countries	10	5	5	10
# Regions	217	91	126	217
Within R-sq.	0.928	0.949	0.924	0.938

Note: The unit of observation is first-level administrative regions that have a large say in educational decisions. The dependent variable is the log of the number of schools in OSM in active OSM areas. Ethnic Frac. Is regional ethnic fractionalization based on GREG and GHSL data. All estimates include country fixed effects, the log of the regional population, land area, night light, the share of urban population and the log of the proxy for completeness of phases one and two ($\ln(p^{I+II})$), as defined in Appendix A. Estimates are based around the indicator "90% gov. edu." that is equal to one if more than 90% of schools are government financed and zero otherwise. The sample used in columns (1) and (4) is based on the set of countries for which this indicator is available; in column (2), those where the indicator is zero; and in column (3), those where the indicator is one. In column (4), all regressors are interacted with the indicator, but the interactions with controls are not reported. Standard errors are reported in parentheses and are clustered at the country level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 3
Main estimate (Placebo restaurants).

Dep. Var.: ln (N Restaurants)	(1)	(2)	(3)	(4)
Ethnic Frac.	0.027	0.066	−0.023	0.066
	(0.085)	(0.095)	(0.175)	(0.095)
Regional education power				−0.090
X Ethnic Frac.				(0.197)
# Countries	121	90	31	121
# Regions	1603	1192	411	1603
Within R-sq.	0.863	0.851	0.894	0.863

Note: The unit of observation is the first-level administrative region. The dependent variable is the log of the number of restaurants in OSM in active OSM areas. Ethnic Frac. Is regional ethnic fractionalization based on GREG and GHSL data. All estimates include country fixed effects, the log of the regional population, land area, night light, the share of urban population and the log of the proxy for completeness of phases one and two ($\ln(p^{I+II})$), as defined in Appendix A. Estimates are based around the indicator "Regional education power" that is equal to one if first-level administrative regions have a large say in educational decisions and zero otherwise. The sample used in columns (1) and (4) is based on the set of countries for which this indicator is available; in column (2), those where the indicator is zero; and in column (3), those where the indicator is one. In column (4), all regressors are interacted with the indicator, but the interactions with controls are not reported. Standard errors are reported in parentheses and are clustered at the country level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

[36] We now know from the work of Muralidharan and Kremer (2006) and others that the private sector plays a major role in the delivery of education in many countries. For the provision of private schools, the degree of collective actions needed, however, is potentially much smaller than that needed for public schools.

[37] Using the years since decentralization is not without problems. Decentralization is often occurring alongside other major reforms in the political systems. Hence, it is not so surprising that the findings here are not very stable; see also the discussion on federalization in the next section.

[38] Data on educational spending at subnational levels is very limited; hence, the sample split is based on regions in only 21 countries.

[39] Note that for the U.S., I find a significant negative correlation of ethnic fractionalization at the school district level. See Appendix B.

test this, Table 25 (Appendix C) presents the baseline estimate for regions in countries where power over education is decentralized but to a lower level than the first-level subnational regions. In all estimates, the correlation of ethnic fractionalization with school supply is much smaller than in the baseline and never significant.[40] These findings contradict the notion that factors collinear with decentralization moderate the effect of ethnic fractionalization observed in my main estimates.

5.4. A universal effect on public amenities

The collective action failure associated with social heterogeneity is likely more relevant for specific public goods. The theoretical argument and empirical findings of Alesina et al. (1999) indicate that the supply of productive public goods is mainly diminished by social heterogeneity. This section focuses on whether this is the case also from a global perspective and whether my findings can be extended to a broader set of public amenities. In what follows, I therefore have a closer look into the relationship of other amenities with ethnic heterogeneity and decentralization.

To extend the analysis to other amenities beyond schools, I have to sacrifice precision in my estimate, as I have to rely on general decentralization measures. There are no data available to classify the location of public spending power on other amenities, such as hospitals or police stations, across subnational units for a large set of countries. Therefore, I rely on Treisman's well-established indicator for general decentralization. In doing so, my estimates become noisier. Even though spending is delegated to the first administrative level in most federal countries, this is not always the case and, most importantly, not the case for all categories of spending. For example, even though the U.S. is a federal country, educational spending is dominated by local entities, not states. Irregularities such as these are likely to occur for other categories of spending in other countries as well. The general decentralization indicator by Treisman (2008) is therefore a noisy indicator of amenity-specific spending power at the regional level.

To establish a new benchmark, I first replicate, wherever possible, the main analyses for the log number of schools using the general measure of decentralization before turning to the other amenities. Table 4 summarizes the replication of the estimates from Table 1. Splitting the sample between regions in federal countries (column (2)) and regions in nonfederal countries (column (3)) again reveals a more negative association between ethnic fractionalization and the supply of schools. Even the coefficient sizes are similar. In contrast to the main estimate, the difference in coefficient estimates is not significant at conventional levels.

Table 26 to Table 30 (Appendix C) document the replication of the main robustness tests. As in the robustness tests of the main estimate, we observe larger coefficients and more significant estimates when not controlling for mapping completeness (Table 26) or when only considering phase one mapping (Table 27). Furthermore, the robustness tests on the assumption's made regarding the measure of ethnic heterogeneity (multi-group areas Tables 28 and 29appsec1) also deliver qualitatively similar findings as in the main estimates.

Table 32 to Table 34 (Appendix C) present the replication of the estimates aiming to narrow down the potential mechanism behind the observed relationship between ethnic fractionalization and the supply of schools. As for the main estimate, I observe significantly stronger correlations in countries with a higher overall share of public financing of education (Table 31). As in the main estimates, a higher percentage of overall education spending at the first administrative level is associated with larger coefficient estimates (Table 32). Looking at countries with prolonged exposure to decentralization, I find the opposite of the findings of the main specification (Table 33).[41] The placebo test (Table 34) shows, as in the main estimates, that there is no significant relationship between ethnic fractionalization and the supply of restaurants.

Overall, using a much noisier proxy for local spending power reveals a qualitatively very similar picture as in the main estimates.

Based on the available OSM data, I can study three alternative proxies for government spending. The number of libraries within a region provides an alternative measure of educational spending that can be classified as a productive public good similar to schools. The number of hospitals in a region can serve as a proxy for healthcare spending. This measure is not without problems because some hospitals may be private or partly private. In many countries, governments subsidize hospitals for their ambulance services to ensure countrywide emergency health care provision. Finally, I treat the number of police stations in regions as an indicator of public safety spending. The argument is that a higher police station density decreases response times. Alesina et al. (1999) argued that, in contrast to educational and health spending, the effect of social heterogeneity on expenditures for public safety is theoretically ambiguous, even when focusing only on the potential failure of collective action.

I conduct the same estimates when studying the other amenities as for schools. Table 35 to Table 58 in Appendix C present these estimates. To have a holistic few of the related main estimates and the central robustness estimates, Table 5 summarizes the sign and significance level of the central coefficients of these estimates. The coefficients referred to are the coefficient of ethnic fractionalization in the sample split between nonfederal (n-f.) and federal countries (fed.), and in the full sample (int.), the coefficient of the interaction of ethnic fractionalization and the federalism dummy Looking at hospitals, I see a very similar pattern even though weaker as when looking at schools, in the main estimates

Table 4
Main result and federalism.

Dep. Var.: ln (N Schools)	(1)	(2)	(3)	(4)
Ethnic Frac.	−0.149**	−0.107	−0.237**	−0.107
	(0.071)	(0.086)	(0.088)	(0.086)
Federal country				−0.130
X Ethnic Frac.				(0.122)
# Countries	149	128	21	149
# Regions	1866	1458	408	1866
Within R-sq.	0.879	0.860	0.933	0.880

Note: The unit of observation is the first-level administrative region. The dependent variable is the log of the number of schools in OSM in active OSM areas. Ethnic Frac. Is regional ethnic fractionalization based on GREG and GHSL data. All estimates include country fixed effects, the log of the regional population, land area, night light, the share of urban population and the log of the proxy for completeness of phases one and two ($\ln(p^{I} + {}^{II})$), as defined in Appendix A. Estimates are based around the indicator "Federal country" that is equal to one if regions belong to a federal state, as defined by Treisman, and zero otherwise. The sample used in columns (1) and (4) is based on the set of countries for which this indicator is available; in column (2), those countries where the indicator is zero; and in column (3), those where the indicator is one. In column (4), all regressors are interacted with the indicator, but the interactions with controls are not reported. Standard errors are reported in parentheses and are clustered at the country level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

[40] This placebo test is furthermore rather conservative, as the policy failure induced by ethnic fractionalization at the subnational regions below the first-level subnational regions under study should bias results toward a more negative effect.

[41] This might be driven by the substantially lower cut-off to define long exposure in the cases of federalism. The cut-off to split the sample in half is 1970 in contrast to 2000 that was used in the cases of educational decentralization.

Table 5
Result summary and federalism alternative proxies for government spending.

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dep. Var.:		ln(N Hospitals)			ln(N Libraries)			ln(N Police)		
Est. Approach		**n-f.**	**fed.**	**int.**	**n-f.**	**fed.**	**int.**	**n-f.**	**fed.**	**int.**
Main estimates	+/-	-	-	-	-	-	-	-	-	-
	p		**						**	
raw OSM data	+/-	-	-	-	-	-	-	-	-	-
	p		***	***		**	**		***	**
p^I	+/-	-	-	-	+	-	-	-	-	-
	p		**	**					***	**
no multiple group regions	+/-	-	-	-	-	-	-	-	-	+
	p	**	***		*			***	*	
pop. share multiple group areas less than 10%	+/-	-	-	-	-	-	+	-	-	+
	p	*	**					**		
p^I>75%	+/-	-	-	-	-	-	-	-	-	+
	p							**		

Note: The unit of observation is the first-level administrative region. The dependent variable in columns (1-3) is the log of the number of hospitals in OSM in active OSM areas, in columns (4-6) that of libraries and in columns (7-9) that of polis stations. The table reports the sign (+/-), and significance level (p) of the central coefficients of the estimates of Tables 35-40, 43-48 and 51-56 in Appendix C. Coefficients referred to are the coefficient of ethnic fractionalization in the sample split between nonfederal (n-f.) and federal countries (fed.) and in the full sample (int.), the coefficient of the interaction of ethnic fractionalization and the federalism dummy. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Each row corresponds to a different estimate specification beginning with the main specification followed by estimates based on: the raw OSM data; using only the proxy for the first phase of mapping; a sample that excludes multigroup homelands; a sample of regions where the share of people living in multigroup homelands is less than 10%; and a sample of regions with more than 75% completion of the first phase of mapping. Green cells indicate the alinement of estimates with the theory of collective action failure and does in red the opposite.

and the robustness tests. Things begin to differ when looking at libraries. The coefficients size and relationships are similar in most cases but often insignificant.[42] The most deviations can be observed when studying the supply of police stations. Here both the sign and size of the estimate coefficients differ more widely. In some cases, the larger coefficients are even observed in nonfederal countries. The same is true for standard errors. They do not systematically increase with a decreasing sample size as in the cases of libraries. Hence, the picture for police stations is less clear than the other estimates.

My findings overall depict a very similar pattern as Alesina et al. (1999), except they are on a global level and are not restricted to the U.S. In both works, educational spending measured by the number of schools and libraries seems to decrease with ethnic fractionalization. In contrast to the findings for the U.S., I find a similar systematic pattern for healthcare spending measured by the number of hospitals. For both types of spending, the negative correlation with ethnic heterogeneity emerges only in decentralized countries, which is in line with the hypothesis of the dominant effect of a collective action failure among policymakers responsible for public spending. As is the case for the U.S., the same cannot be stated for safety spending measured by the number of police stations. The negative relationship of ethnic fractionalization on the number of police stations does not seem to depend systematically on local spending power.

[42] Overall, the estimate's precision is smaller when studying other amenities. This might be partly the result of a reduction in sample size or the nature of amenities. Naturally, the number of alternative amenities is smaller; e.g., we don't need as many hospitals as we need schools. This gives missing data, on average, more weight, and estimates become noisier due to OSM data collection issues.

6. Conclusion

This paper provides a first global view of the potential link between decentralization and social heterogeneity in the provision of regional public goods. The estimates indicate that increasing local autonomy may hamper the provision of public goods in regions that face high levels of social heterogeneity. This finding is in line with the theory of collective action failure and social heterogeneity. The correlation is also sizable. It implies that an increase in ethnic fractionalization by a standard deviation may decrease the supply of schools in a region by 5% if educational spending is decentralized to the regional level. I find similar results for the number of libraries and hospitals but not for restaurants or police stations.

The analysis is based on a new dataset derived from the OSM project, which contains the global locations of various public amenities associated with public goods that are typically provided mainly by the state. Well-known accuracy problems associated with using volunteered geocode data are addressed by developing a new method that accounts for the completeness of the data within first-level subnational regions by cross-referencing OSM settlement indicators with indicators derived from satellite data. The new approach minimizes the risk of potential biases due to omitted variables creating systematic missing data in the OSM data. The quality of the approach is tested by correcting the OSM raw data and comparing the corrected data with official data of a subset of countries. The observed correlation between the corrected OSM data and the official data is typically greater than 90%. The main findings of the paper hold when the raw OSM data are used and when different technical details of the algorithms used to clean the raw data or account for the possibility of systematically missing data are altered.

The findings are robust to a large set of robustness tests based on a broad set of controls and to alternative indicators for public goods, social

heterogeneity and decentralization. The placebo test indicates that the supply of regional nonpublic amenities, such as restaurants, is not affected by the joint occurrence of social heterogeneity and decentralization. Examining the data shows no indication that regional social heterogeneity might be the driver of decentralization or that the provision of public goods might induce social heterogeneity or decentralization. Furthermore, placebo tests using the heterogeneity in the location of spending power at different subnational levels of government indicate that findings are not linked to the status of being decentralized. If lower tiers of government hold the main spending power, social heterogeneity does not impact educational spending at the regional level. Furthermore, I find that more prolonged and more profound decentralization is associated with a larger correlation between ethnic fractionalization and the number of schools. Overall, all of the robustness tests indicate that it is plausible that the findings document the effect of social heterogeneity and decentralization on the provision of regional public goods. Nevertheless, given the nature of the available data, a hard claim of a causal effect cannot be made.

The findings shed light on a potential and previously unknown further dark side of the growing worldwide trend of government decentralization from national to regional levels. Increasing local autonomy might increase, on average, the effectiveness of government spending within the regions of a country. However, in some cases, the opposite might be true because power is given to a layer of government that is too socially heterogeneous to execute collective actions. This finding might explain how decentralization can lead to increases in regional disparities (e.g., Rodríguez-Pose and Ezcurra (2009), Lessmann (2012) or Lessmann and Seidel (2017)). One possible conclusion that some might draw from this finding is that decentralization should be accompanied by administrative reforms that decrease social heterogeneity within regions. However, such a policy might increase separatist tendencies; hence, further study should be undertaken before such policies are recommended. Furthermore, other costs and end benefits of decentralization and administrative reforms must be factored in when deciding on such measures. Finally, it is not clear a priori in which cases higher public spending is welfare enhancing in heterogeneous societies. In this area, more theoretical work is needed.

The dataset generated for this study, along with the proposed approach to account for missing OSM data, provides a variety of opportunities for possible further research. The dataset offers a way to examine other aspects of the political economy driving the provision of public goods via public amenities. For example, it could be used to explore political favoritism or to measure the extent to which political leaders use the provision of public goods to benefit favored regions and populations. Such examinations might, for example, help to understand the mechanism underlying the finding that favoritism impacts growth (Hodler and Raschky, 2014; De Luca, Hodler, Raschky and Valsecchi, 2018). These questions, which go beyond the scope of this paper, remain open for future research.

Author statement

I am the single author of this paper and solely responsible for its content. On the paper draft, I have received advice from multiple sources; all have been paid tribute to in the first footnote of the paper.

Data availability

Data will be made available on request. The amenity data and the proxies for mapping completeness are available directly via my website https://sites.google.com/site/researchandreseidel.

Supplementary data.

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jdeveco.2023.103113.

References

Alesina, A., La Ferrara, E., 2000. Participation in heterogeneous communities. Q. J. Econ. 115 (3), 847–904.

Alesina, A., Zhuravskaya, E., 2011. Segregation and the quality of government in a cross section of countries. Am. Econ. Rev. 101 (5), 1872–1911.

Alesina, A., Baqir, R., Easterly, W., 1999. Public goods and ethnic divisions. Q. J. Econ. 4 (1), 1243–1284.

Alesina, A., Gennaioli, C., Lovo, S., 2019. Public goods and ethnic diversity: evidence from deforestation in Indonesia. Economica 86 (341), 32–66.

Baqir, R., 2002. Social sector spending in a panel of countries. Workingpaper International Monetary Fund 2 (35).

Bernhard, H., Fischbacher, U., Fehr, E., 2006. Parochial altruism in humans. Nature 442 (7105), 912–915.

Blanes i Vidal, J., Kirchmaier, T., 2018. The effect of police response time on crime clearance rates. Rev. Econ. Stud. 85 (2), 855–891.

Bluhm, R., Thomsson, K., 2020. Holding on? Ethnic divisions, political institutions and the duration of economic declines. J. Dev. Econ. 144.

Buchmueller, T.C., Jacobson, M., Wold, C., 2006. How far to the hospital?: the effect of hospital closures on access to care. J. Health Econ. 25 (4), 740–761.

Burde, D., Linden, L.L., 2013. Bringing education to Afghan girls: a randomized ControlledTrial of village-based schools. Am. Econ. J. Appl. Econ. 5 (3), 27–40.

Cinnirella, F., Schueler, R., 2016. The Cost of Decentralization: Linguistic Polarization and the Provision of Education. CESifo Working Paper. *No. 5894*.

Dayton-Johnson, J., 2000. Determinants of collective action on the local commons: a model with evidence from Mexico. J. Dev. Econ. 62 (1), 181–208.

De Luca, G., Hodler, R., Raschky, P.A., Valsecchi, M., 2018. Ethnic favoritism: an axiom of politics? J. Dev. Econ. 132, 115–129.

Desmet, K., Gomes, J.F., Ortuño-Ortín, Ignacio, 2020. The geography of linguistic diversity and the provision of public goods. J. Dev. Econ. 143, 102384.

Desmet, K., Ortuño-Ortín, Ignacio, Wacziarg, R., 2012. The political economy of linguistic cleavages. J. Dev. Econ. 97 (2), 322–338.

Díaz-Cayeros, A., Magaloni, B., Ruiz-Euler, A., 2014. Traditional governance, citizen engagement, and local public goods: evidence from Mexico. World Dev. 53, 80–93.

Dražanová, L., 2019. Historical Index of Ethnic Fractionalization Dataset (HIEF). Harvard Dataverse, V1.0. https://doi.org/10.7910/DVN/4JQRCL.

Duflo, E., 2001. Schooling and labor market consequences of school construction in Indonesia:evidence from an unusual policy experiment. Am. Econ. Rev. 91 (4), 795–813.

Easterly, W., Levine, R., 1997. Africa's growth tragedy: policies and ethnic divisions. Q. J. Econ. 112 (4), 1203–1250.

Eberle, U.J., Henderson, J.V., Rohner, D., Schmidheiny, K., 2020. Ethnolinguistic diversity and urban agglomeration. Proc. Natl. Acad. Sci. USA 117 (28), 16250–16257.

Fan, C., Lin, C., Treisman, D., 2009. Political decentralization and corruption: evidence from around the world. J. Publ. Econ. 93 (1–2), 14–34.

Gershman, B., Rivera, D., 2018. Subnational diversity in sub-saharan Africa: insights from a new dataset. J. Dev. Econ. 133, 231–263.

Glaeser, E.L., Laibson, D.I., Scheinkman, J.A., Soutter, C.L., 2000. Measuring trust. Q. J. Econ. 115 (3), 811–846.

Glennerster, R., Miguel, E., Rothenberg, A.D., 2013. Collective action in diverse Sierra Leone communities. Econ. J. 123 (568), 285–316.

Habyarimana, J., Humphreys, M., Posner, D.N., Weinstein, J.M., 2007. Why does ethnic diversity undermine public goods provision? Am. Polit. Sci. Rev. 101 (4), 709–725.

Hodler, R., Raschky, P.A., 2014. Regional favoritism. Q. J. Econ. 129 (2), 995–1033.

Kazianga, H., Levy, D., Linden, L.L., Sloan, M., 2013. The effects of' girl-friendly' schools: evidence from the BRIGHT school construction program in Burkina Faso. Am. Econ. J. Appl. Econ. 5 (3), 41–62.

Khwaja, A.I., 2009. Can good projects succeed in bad communities? J. Publ. Econ. 93 (7), 899–916.

Lessmann, C., 2009. Fiscal decentralization and regional disparity: evidence from cross-section and panel data. Environ. Plann.: Econ. Space 41 (10), 2455–2473.

Lessmann, C., 2012. Regional inequality and decentralization: an empirical analysis. Environ. Plann. 44 (6), 1363–1388.

Lessmann, C., Markwardt, G., 2010. One size fits all? Decentralization, corruption, and the monitoring of bureaucrats. World Dev. 38 (4), 631–646.

Lessmann, C., Seidel, A., 2017. Regional inequality, convergence, and its determinants – a view from outer space. Eur. Econ. Rev. 92 (1), 110–132.

Martinez-Vazquez, J., Lago-Peñas, S., Sacchi, A., 2017. The impact of fiscal decentralization: a survey. J. Econ. Surv. 31 (4), 1095–1129.

Miguel, E., 2004. Tribe or nation? Nation building and public goods in Kenya versus Tanzania. World Polit. 56 (3), 327–362.

Miguel, E., Gugerty, M.K., 2005. Ethnic diversity, social sanctions, and public goods in Kenya. J. Publ. Econ. 89 (11), 2325–2368.

Montalvo, J.G., Reynal-Querol, M., 2005. Ethnic polarization, potential conflict, and civil wars. Am. Econ. Rev. 95 (3), 796–816.

Montalvo, J.G., Reynal-Querol, M., 2020. Ethnic Diversity and Growth: Revisiting the Evidence. . The Review of Economics and Statistics. https://doi.org/10.1162/rest_a_00901.

Muralidharan, K., Kremer, M., 2006. Public and Private Schools in Rural India, vol. 9. Harvard University, Department of Economics, Cambridge, MA, pp. 10–11.

Muralidharan, K., Prakash, N., 2017. Cycling to school: increasing secondary school enrollment for girls in India. Am. Econ. J. Appl. Econ. 9 (3), 321–350.

Musgrave, R., 1959. The Theory of Public Finance: A Study in Public Economy. McGraw-Hill, New York.

Nagaraj, A., 2021. Information Seeding and Knowledge Production in Online Communities: Evidence from OpenStreetMap. Management Science, Forthcoming.
Neyapti, B., 2006. Revenue decentralization and income distribution, 166 Econ. Lett. 92, 409.
Oates, W., 1972. Fiscal Federalism. Harcourt Brace Jovanovich, New York.
OECD, 2019. Current Trends in Decentralisation. I *Making Decentralisation Work: A Handbook For Policy-Makers*. OECD Publishing, Paris. https://doi.org/10.1787/f9117991-en.
Rodríguez-Pose, A., Ezcurra, R., 2009. Does decentralization matter for regional disparities? A cross-country analysis. J. Econ. Geogr. 10 (5), 619–644.
Senaratne, H., Mobasheri, A., Ali, A.L., Capineri, C., Haklay, M., 2017. A review of volunteered geographic information quality assessment methods. International Journal of Geographical Information Science, 31(1. 139-167. International Journal of Geographical Information Science 31 (1), 139–167.
Tiebout, C., 1956. A pure theory of local expenditures. J. Polit. Econ. 64, 416–424.
Treisman, D., 2008. Decentralization Dataset. Retrieved from. https://www.danieltreisman.org/workinprogress/.
Weidmann, N.B., Jan Ketil, Rød, Cederman, L.-E., 2010. Representing ethnic groups in space: a new dataset. J. Peace Res. 47 (4), 491–499.
Wilde, E.T., 2013. Do emergency medical system response times matter for health outcomes? Health Econ. 22 (7), 790–806.